Analyzing Global and Local Latencies around the world

COMP 631 Final Project Report

Rohit Kumar Gupta and Sahil Narang

10th December 2014

Objective

Goal of the project was to study and analyze the patterns observed during datagram transfers through various locations around the world. Datagrams take different routes to reach a specific destination and our goal was to observe the variation in their latencies and check if there is any behavioral pattern for traces on weekdays and weekends. Also we studied the dependence of latency with respect to distance of the destination from source and for different countries. Additionally, we observed that there are at times transatlantic hops in some 'local' pings and we reported the observation and its impact it has over the network.

Introduction

Every OS comes with a traceroute tool which shows a list of router hops for a specific destination IP. It also tells where the traceroute stops, drops packets, or where the latency goes up a lot.

Traceroute steps are listed below and shown in Figure 1:

i. Launch a probe packet towards destination, with a TTL of 1
ii. At every router hop TTL is decremented by 1
iii. As TTL hits 0, packet is dropped, routers sends ICMP TTL exceed packet to source
iv. Source receives this ICMP message, displays a traceroute 'hop'
v. Repeat step 1 with TTL incremented by 1 every time until..
vi. Destination host receives probe and returns ICMP destination unreachable
vii. Source stops upon receiving ICMP destination unreachable.

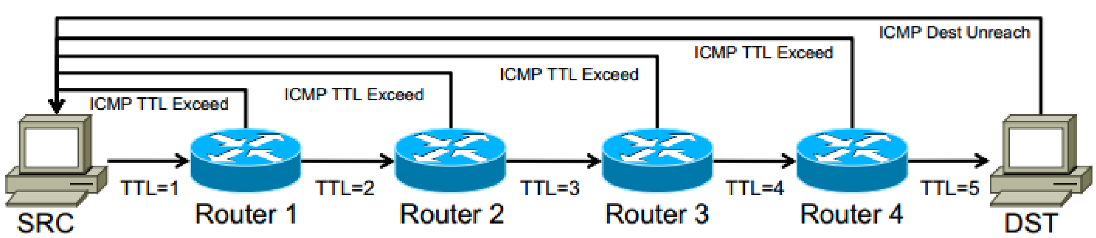

Fig 1: Traceroute Steps

Default for traceroute is set 3 probes per hop and 60 max hops. Traceroute latency is calculated by

RTT = Timestamp(ICMP) – Timestamp(probe launch)

There are three main types of network induced latency

i. Serialization Delay : delay caused by transmission through routers/switches in packet sized chunks
ii. Queuing Delay : time spent in a router's queue waiting for transmission (due to congestion)
iii. Propagation Delay : time spent 'in flight' over the transmission medium.

Interpreting DNS in a traceroute includes getting physical router location and in our case we used an open source tool, CuRL, to get the exact location of an IP addresss. Below is a sample output of CuRL command "***curl ipinfo.io/64.57.21.225***"

{ "ip": "64.57.21.225",

"hostname": "ae-4.586.asbn0.tr-cps.internet2.edu",

"city": "Ann Arbor",

"region": "Michigan",

"country": "US",

"loc": "42.2734,-83.7133",

"org": "AS11164 Internet2",

"postal": "48104"}

Using the above command we estimated the location of the hops in our traceroute output and estimated the latencies between source (i.e. Chapel Hill) and all these intermediate locations along with our final destinations.

Approach

In our experiments we selected 25 Globally distributed servers (spread over around 10 countries with 2-3 servers in each country) as shown in Fig 2. Traceroute tool was run every hour as a cronjob for approximately 20 days thereby collecting 10,200 traces for each server. For fair analysis between latencies of weekdays and weekends we took around 3000 traces from both weekends and weekdays individually. For observing transatlantic hops data from 5 different servers were collected every hour which span over 12 days of our observation window giving us around 1440 traces.

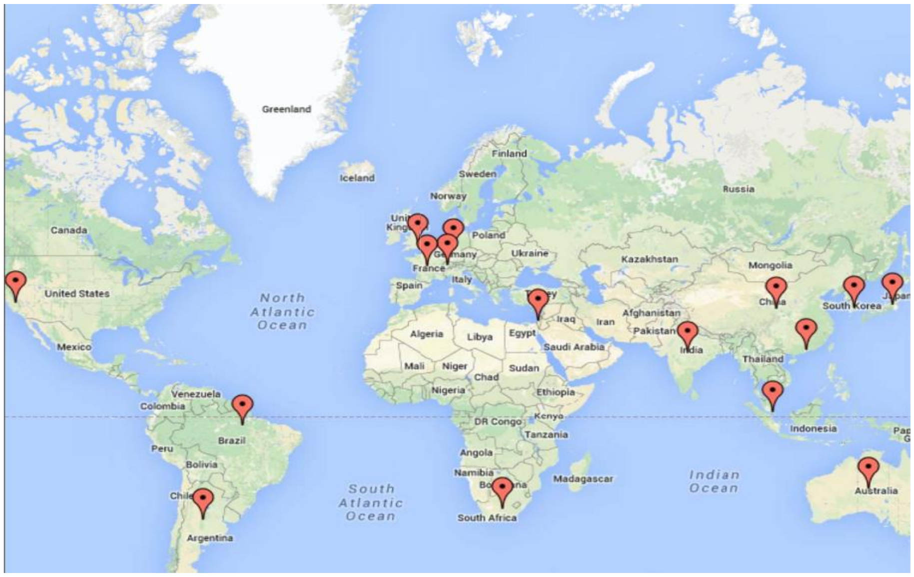

Fig 2: Geographical locations of our test servers

A sample of our processed data is shown in Fig 3 where IP's location was also obtained for each hop.

```
val(:,:,1) =

  Columns 1 through 5

    '152.2.31.1'       'Chapel Hill'    'North Carolina'    'US'    '5.314'
    '152.2.255.209'    'Chapel Hill'    'North Carolina'    'US'    '1.546'
    '128.109.1.89'     'Dobson'         'North Carolina'    'US'    '6.340'
    '128.109.1.94'     'Dobson'         'North Carolina'    'US'    '9.960'
    '4.26.6.193'       'ul'             'ul'                'US'    '13.999'
    '4.69.149.16'      'ul'             'ul'                'US'    '15.606'
    '4.69.149.208'     'ul'             'ul'                'US'    '25.947'
    '216.6.87.65'      'Wilmington'     'Delaware'          'US'    '23.919'
    '64.86.252.197'    'Wilmington'     'Delaware'          'US'    '66.790'
    '180.87.15.25'     'Pune'           'Maharashtra'       'IN'    '253.779'
                  []               []                  []     []          []
                  []               []                  []     []          []
                  []               []                  []     []          []

  Columns 6 through 10

    '152.2.31.1'       'Chapel Hill'    'North Carolina'    'US'    '5.319'
    '152.2.255.209'    'Chapel Hill'    'North Carolina'    'US'    '1.559'
    '128.109.1.89'     'Dobson'         'North Carolina'    'US'    '6.938'
    '128.109.1.101'    'Dobson'         'North Carolina'    'US'    '9.843'
    '4.26.6.193'       'ul'             'ul'                'US'    '14.531'
    '4.69.149.144'     'ul'             'ul'                'US'    '14.815'
    '4.69.149.208'     'ul'             'ul'                'US'    '19.452'
    '216.6.87.65'      'Wilmington'     'Delaware'          'US'    '23.938'
    '64.86.252.197'    'Wilmington'     'Delaware'          'US'    '70.713'
    '180.87.15.25'     'Pune'           'Maharashtra'       'IN'    '250.480'
                  []               []                  []     []          []
                  []               []                  []     []          []
                  []               []                  []     []          []

  Columns 11 through 15

    '152.2.31.1'       'Chapel Hill'    'North Carolina'    'US'    '5.338'
    '152.2.255.209'    'Chapel Hill'    'North Carolina'    'US'    '1.561'
    '128.109.1.89'     'Dobson'         'North Carolina'    'US'    '6.940'
    '128.109.1.101'    'Dobson'         'North Carolina'    'US'    '9.938'
    '4.26.6.193'       'ul'             'ul'                'US'    '14.005'
    '4.69.149.144'     'ul'             'ul'                'US'    '14.791'
    '4.69.149.80'      'ul'             'ul'                'US'    '14.769'
    '216.6.87.65'      'Wilmington'     'Delaware'          'US'    '23.921'
    '64.86.252.197'    'Wilmington'     'Delaware'          'US'    '71.326'
    '180.87.15.25'     'Pune'           'Maharashtra'       'IN'    '245.666'
                  []               []                  []     []          []
                  []               []                  []     []          []
                  []               []                  []     []          []
```

Fig 3: Sample of Processed Data

Results and Observations

In our observation we have shown variation of latency w.r.t time, distance, countries for both weekdays and weekends. Also we have shown transatlantic hops and its behavior and distribution over our limited observation time.

Variation of latency w.r.t time (Fig 4):

- No discernable pattern was observed but it averaged around a difference of approximately 10 seconds. As expected weekends had higher latency on average
- Few anomalies (for Brazil and China) ; maybe due to already huge traffic load or congestion in their networks

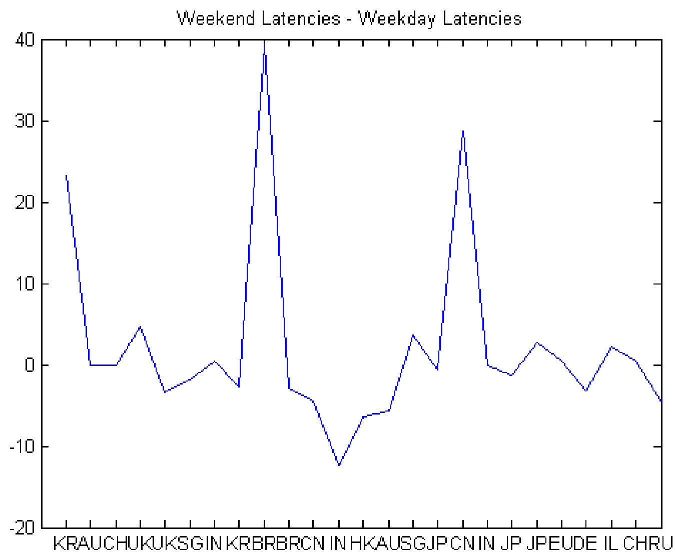

Fig 4: Difference in latency between weekend and weekday

Now we'll give some statistics for weekday and weekend latencies for 3 countries:

(server in Brazil) :

Weekday Latencies : 175 ± 21 ms

Weekend Latencies : 215 ± 152 ms

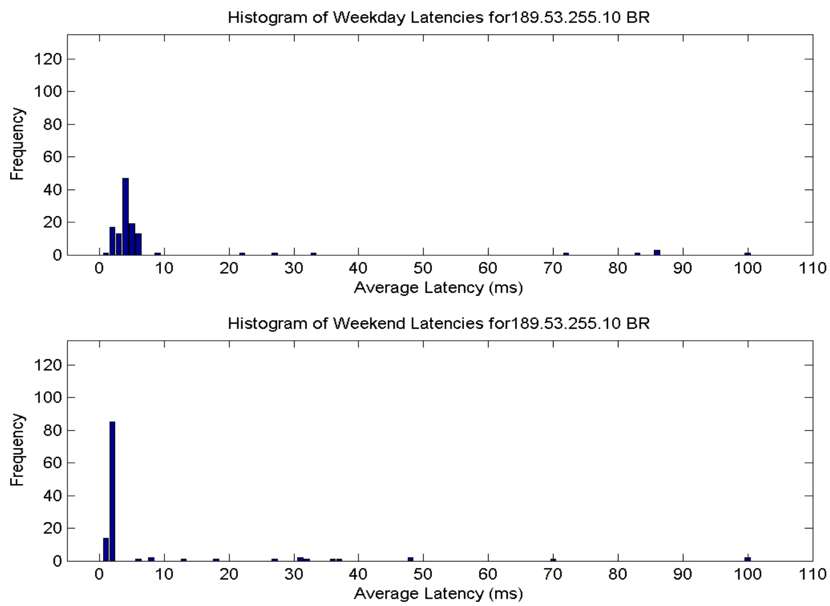

Fig 5: Histogram for the latencies in Brazil Server for Weekday and Weekend

(server in Korea) :

Weekday Latencies : 201 ± 13 ms

Weekend Latencies : 224 ± 28ms

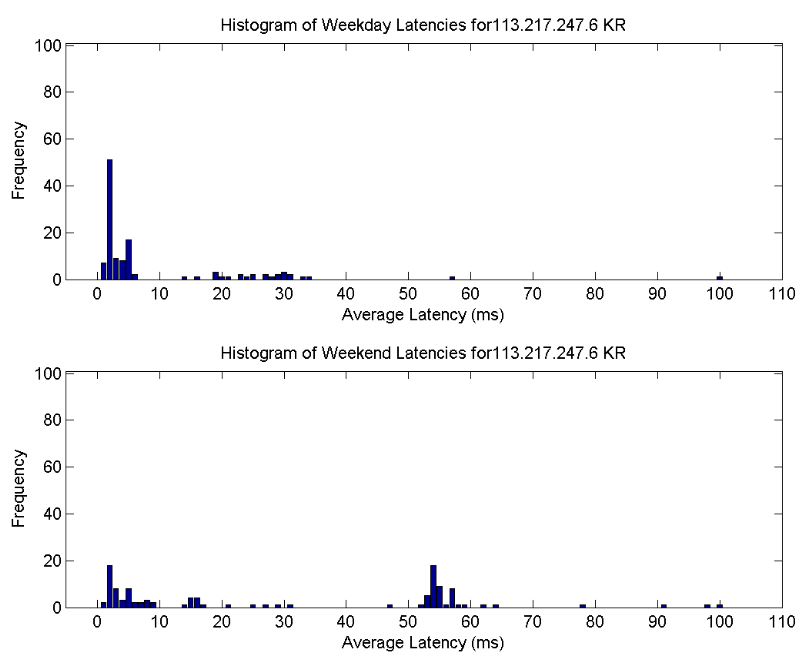

Fig 6: Histogram for the latencies in Korea Server for Weekday and Weekend

(server in Japan) :

Weekday Latencies : 277 ± 56 ms

Weekend Latencies : 306 ± 67 ms

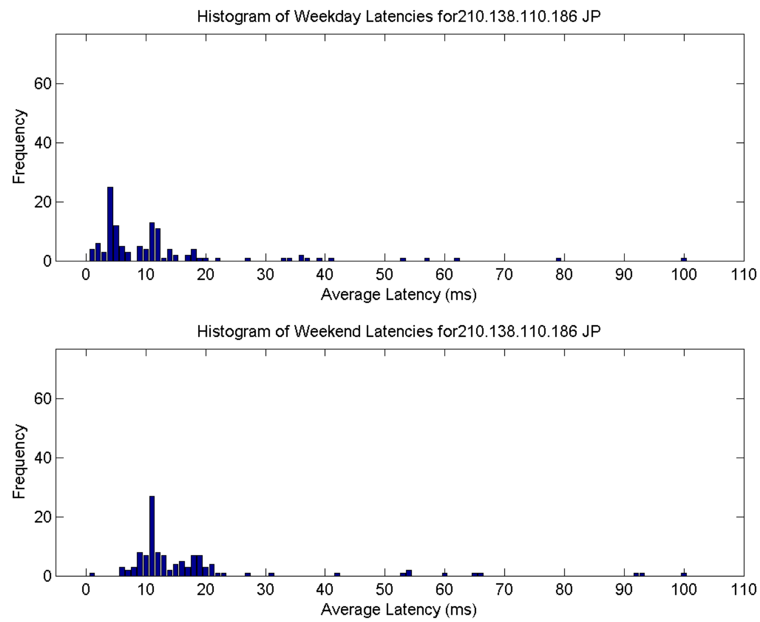

Fig 7: Histogram for the latencies in Japan Server for Weekday and Weekend

<u>Latency vs Distance</u>

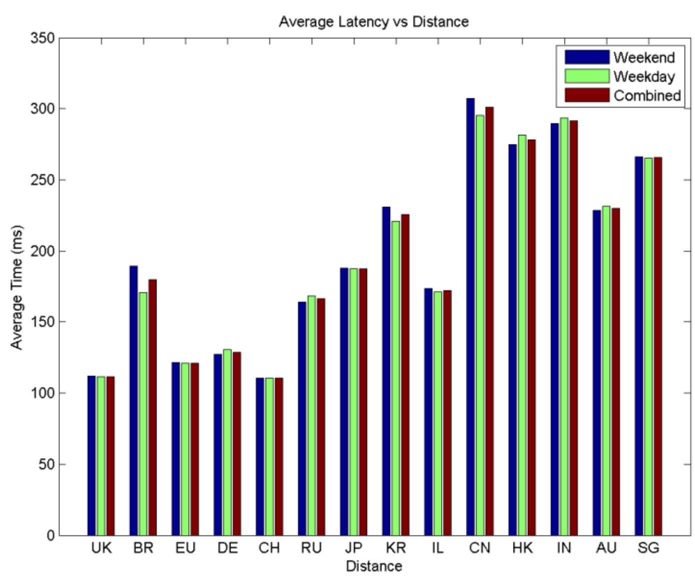

Fig 8: Plot of Latency for different Countries sorted in order of Distance for US

Latency w.r.t Countries

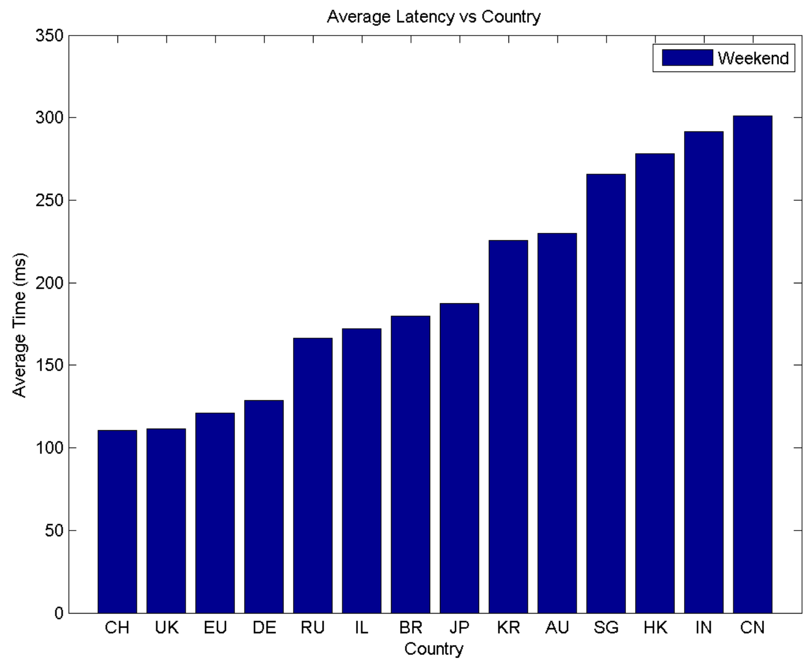

Fig 9: Plot of Average Latency with Countries

Transatlantic Hops

We observed that for some local pings, the datagram was sometimes surprisingly going to some other continent before coming to the American continent which isn't expected as it would be costly (or atleast seems so in first look). Below table shows the number of times we observed such event for these 5 big websites in US.

Table 1: Number of transatlantic hops observed for probes to 5 big US websites

Amazon		**Facebook**		**Google**		**Yahoo**		**Youtube**	
Non-US	**Total**	**Non-US**	**Total**	**Non-US**	**Total**	**Non-US**	**Total**	**Non-US**	**Total**
0	1440	416	1440	195	1440	0	1440	222	1440

Below is one sample of a probe to facebook server in California and same thing is shown in a figure 10.

curl ipinfo.io/128.109.9.102 {

"ip": "128.109.9.102",

"hostname": "No Hostname",

"city": null,

"region": "North Carolina",

"country": "US",

"loc": "35.2271,-80.8431",

"org": "AS81 MCNC"

}classroom(54)% curl ipinfo.io/64.57.21.225

{ "ip": "64.57.21.225",

"hostname": "ae-4.586.asbn0.tr-cps.internet2.edu",

"city": "Ann Arbor",

"region": "Michigan",

"country": "US",

"loc": "42.2734,-83.7133",

"org": "AS11164 Internet2",

"postal": "48104"

}classroom(55)% curl ipinfo.io/206.126.236.183

{ "ip": "206.126.236.183",

"hostname": "dc2.br01.iad1.tfbnw.net",

"city": "Redwood City",

"region": "California",

"country": "US",

```
  "loc": "37.5331,-122.2471",
  "postal": "94065"
}classroom(56)% curl ipinfo.io/31.13.24.12
{  "ip": "31.13.24.12",
  "hostname": "be3.bb02.iad3.tfbnw.net",
  "city": null,
  "region": null,
  "country": "IE",
  "loc": "53.0000,-8.0000",
  "org": "AS32934 Facebook, Inc."
}classroom(57)% curl ipinfo.io/74.119.78.31
{  "ip": "74.119.78.31",
  "hostname": "ae52.dr01.ash3.tfbnw.net",
  "city": "Palo Alto",
  "region": "California",
  "country": "US",
  "loc": "37.3762,-122.1826",
  "org": "AS32934 Facebook, Inc.",
  "postal": "94304"
}classroom(58)% curl ipinfo.io/31.13.29.171
{  "ip": "31.13.29.171",
  "hostname": "po126.msw01.10.iad1.tfbnw.net",
  "city": null,
  "region": null,
  "country": "IE",
```

"loc": "53.0000,-8.0000",

"org": "AS32934 Facebook, Inc."

}classroom(59)% curl ipinfo.io/31.13.69.80

{ "ip": "31.13.69.80",

"hostname": "edge-star-shv-10-iad1.facebook.com",

"city": "Ashburn",

"region": "Virginia",

"country": "US",

"loc": "39.0437,-77.4875",

"org": "AS32934 Facebook, Inc." }

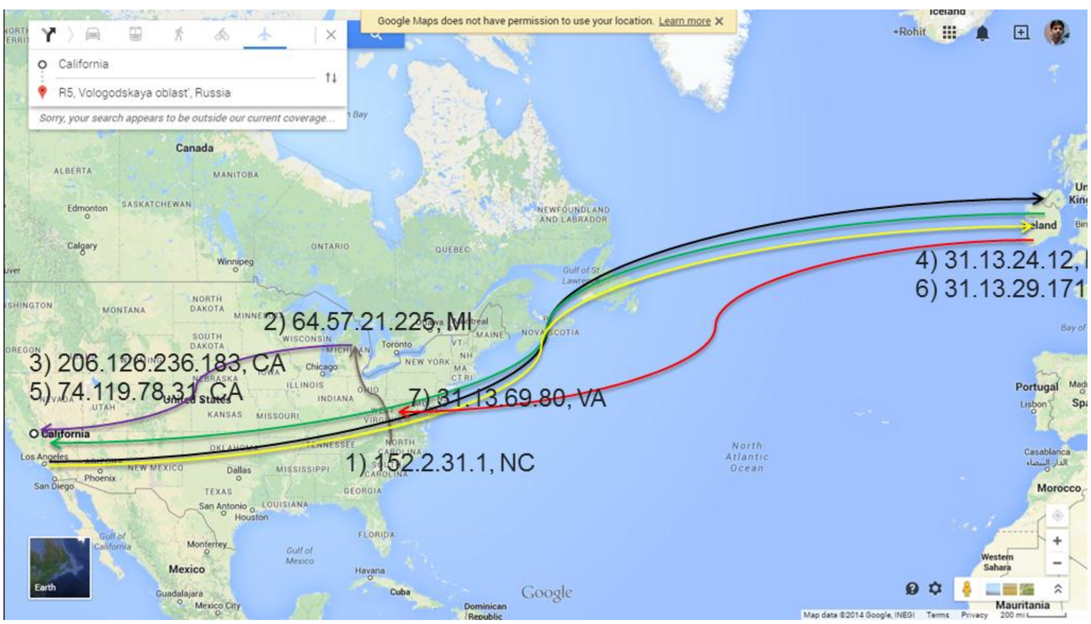

Fig 10: Color coded paths for a ping to facebook server in California which does a transatlantic hop in Ireland

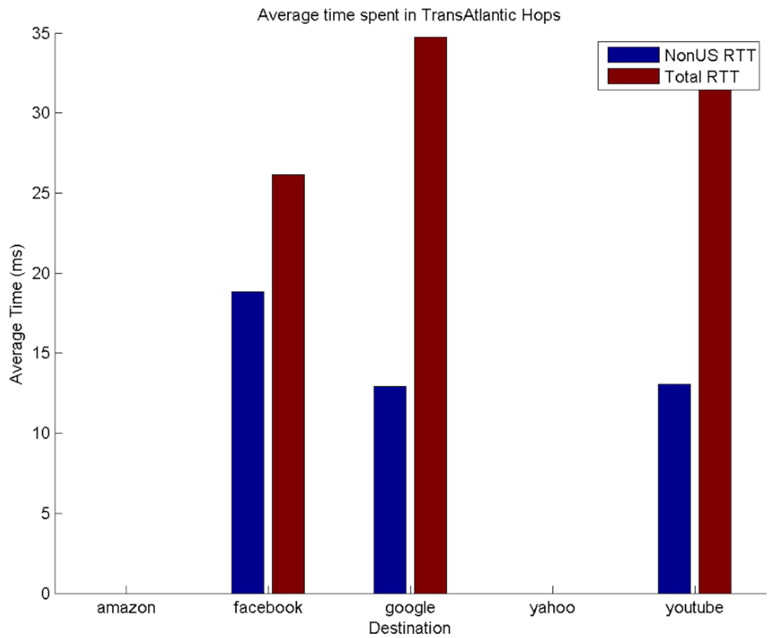

Fig 11: Histogram of time spent in transatlantic hops for US and NonUS RTTs

<u>Overall Mean and Variance</u>

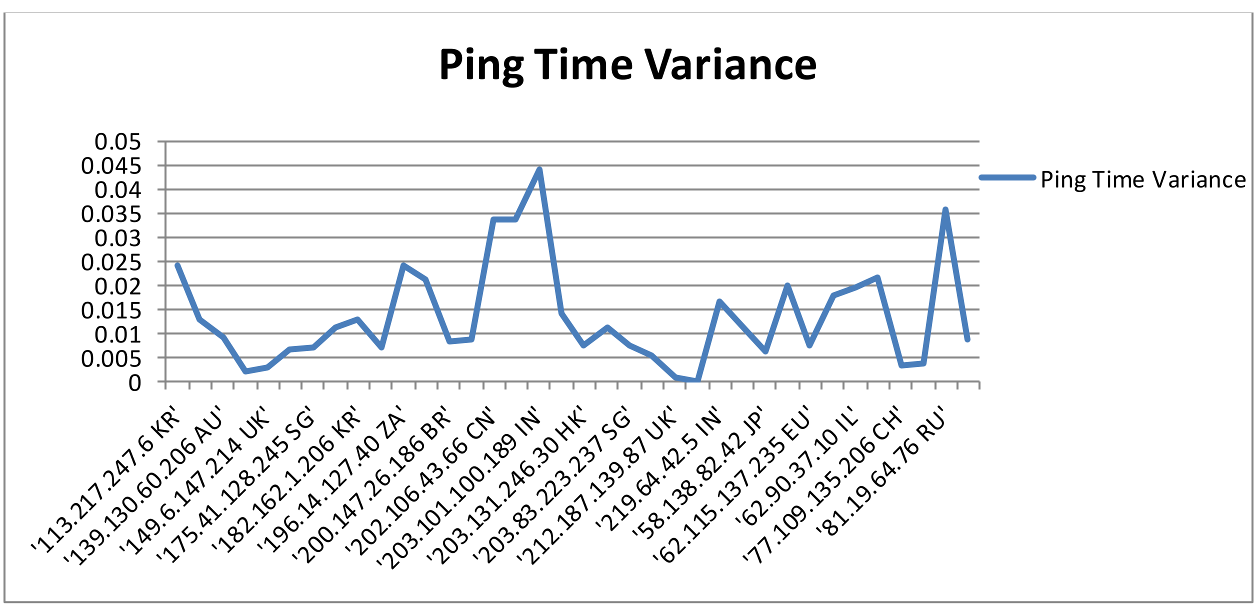

Fig 12: Ping Time Variance for all the test servers

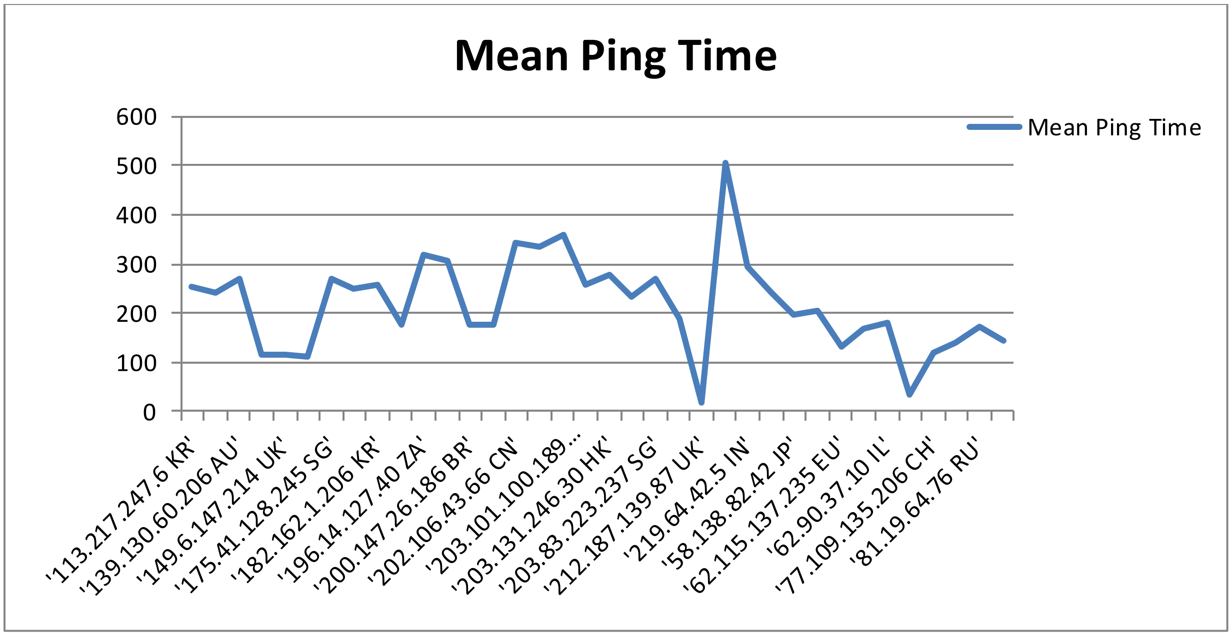

Fig 13: Mean of Ping Time for all the test servers

Conclusion

From our experiments we observed that Weekend Latencies were higher than Weekday Latencies for most countries as expected. Some anomalies were for Brazil and China but from our limited experiments we couldn't conclude the exact reasons. Some countries like China, India, HongKong have almost 2 to 3 times higher latencies than other countries which might be due to congestion in their networks as these countries have high population. We also saw that some local pings route via different continents which we termed as transatlantic hops which is seen for some major companies and reasons couldn't be inferred without doing further in-depth studies in this regard. But we saw that Non-US RTTs are at times smaller than US RTTs for some pings which might be one of the reasons for such a behavior. Moreover, using our analysis we can predict average latency for some specific countries (which we observed) for both weekday and weekend probes.