

Chapter 1

Network Latency

Nowadays, people spend great time using electronic gadgets such as computers, smartphones or tablets. As time goes by, all of these devices seem useless without Internet connectivity. Users always want to improve their experience while using Internet services. They acquire Internet connections with the highest bandwidth possible; they buy the phone which promises the fastest connectivity by using 4G; and they sign up contracts for residential Internet through optical fibre. All of these expecting to have the best experience. However, there is a hidden element that can affect user experience that most users do not consider which is **network latency**.

1.1 Prelude

The latency of a network is the sum of the latency added by every node that takes part in a network communication. Each network node has an inherent latency, specified by its manufacturer. The latency of a network node —such as a router, a switch, a hub— is the time elapsed from the moment when it receives a packet and when it sends a response. The path between two computers in a network is dynamic and normally implies several hops through intermediate nodes. Therefore, the latency of a connection will depend on the path that it takes. The range of network latency in Internet are normally in the order of milliseconds. Since the latency is a base delay in each connection, protocols that open many connections are more affected. In consequence, services such as music, video and online games which typically split data into a bunch of small packets are more sensible to high latency.

Motivation

When a node of the network fails, the communication paths that pass by the node will increase their latency (the impact depends on the nature of the failure). From the point of view of the company which offers the service, there is a need for monitoring how is latency levels among the nodes of the network. In case of failure they can react and take corrective actions. Visual tools are extensively used for monitoring services. They allow operators to spot failures at a glance, or in other cases to be aware how close is the latency to the threshold where user experience start to compromise.

Case study

In our case study we analysed a proxy service that operates 145 nodes distributed in different cities around the world. By using Roassal we developed a tool for visualizing the status of the service. We called the tool Network Latency Visual Analyser (NLVA). The tool envisions two kind of users: a company that manages a network and clients of that company. The tool offers different views of the dataset which helps the company to be aware of failures (to spot where possible failures occurred and what potential users can be affected); to take actions for improving the service by highlighting nodes that are geographically close but that have a high latency—in comparison with the rest of them— allowing the company to possibly change the service provider between those nodes; and to spot for each node where are the best endpoints that offer the lowest latency, which can be of help to end users—e.g. a gamer who wants to know to what server he should connect to minimize the lag.

1.2 Network Latency Visual Analysis in a Nutshell

When developing NLVA we did not want to start from the scratch. We thought that MapBuilder example which comes with Roassal offers many of the features that we wanted to have in this tool. However, we decide not to directly modify it, or to replicate its code. We reused it by extending some parts of it to fit our purpose. NLVA implementation is hold in the NetworkLatency package, which includes six classes:

- NLCityBuilder
- NLMapBuilder
- NLPopup

- NLRRoute
- NLRRouteBuilder
- NLVisualise

All of these classes except by NLVisualise are extending a Roassal counterpart. We follow this naming convention `RTClassName` → `NLClassName`. These classes add or modify features of the original MapBuilder implementation. The remaining class NLVisualise is the entry point to run the visualization. It is the main class that holds methods to retrieve the dataset, to trigger ping commands in the local machine, and to build the visualization.

Resources

NLVA comes with a dataset that contains a snapshot of the latency between servers of a network. It also includes the location (lat/long) of the cities highlighted in the visualization. Besides, it holds IP addresses that are physically located in each city (used when testing user local connectivity). These data were collected from the following sites:

Resource	URL
Network Latency dataset	https://wonderproxy.com
IP Location	http://www.iplocation.net
Geolocation	http://www.latlong.net

Visual Queries

The following queries are examples showing that a visualization of the dataset provides a quick notion of the answer, that can be followed by a detailed analysis of the latency measures. At the end, it can help to exclude candidates from the set of possible answers narrowing the space of solutions. As an exercise the reader can surely find more queries that can be supported by the different views included in the tool.

1. Where a company using this network should store content for delivering to users in Paris? (*Hint: Best*)
2. Do you notice some possible failure in the network that can be impacting latency? (*Hint: Worst*)
3. Are cities with better connectivity distributed equally in both hemispheres? (*Hint: Average*)

4. What are the cities where the company should invest to improve overall network latency? (*Hint: CloseHigh*)
5. If you want to provide a service in both the US and Europe (using this network), where would be the best two cities to locate a host? (*Hint: FarLow*)
6. If I am in the Paris to which continents should I privilege connections?. How *my service* would be affected if a move it from Paris to New York? (*Hint: Left-click on a city*)
7. How is network latency in general from my city? (*Hint: Right-click on your city if listed or in another region of a country*)
8. Would I improve my connectivity using the service from which the dataset comes (over my actual provider)? (*Hint: Comparing the results of left and right click if your city is included in the visualization*)

1.3 How to load NLVA

Firstly, you have to download the latest version of Moose from <http://www.moosetechnology.org/#install>. Secondly, once in Moose you have to open a Workspace (World Menu -> Workspace). Finally, you can copy&paste and *Do It*:

```
Gofer new smalltalkhubUser: 'merino'
  project: 'NetworkLatency';
  package: 'ConfigurationOfNetworkLatency';
  load.
(Smalltalk at: #ConfigurationOfNetworkLatency) loadDefault
```

Afterwards, you can launch the tool by sending the message:

```
NLVisualise new open
```

1.4 Visualizing latency in a proxy service

NLVA shows a visualization of a map of the world highlighting 145 cities spread over 70 countries (purple circle). NLVA comes with several views of the data shown in a menu at the top of the window that helps to move from different views such as 'Best', 'Worst', 'MST' (Minimal Spanning Tree), 'Close High' and 'Far Low'. The default view is 'Best' which shows for each

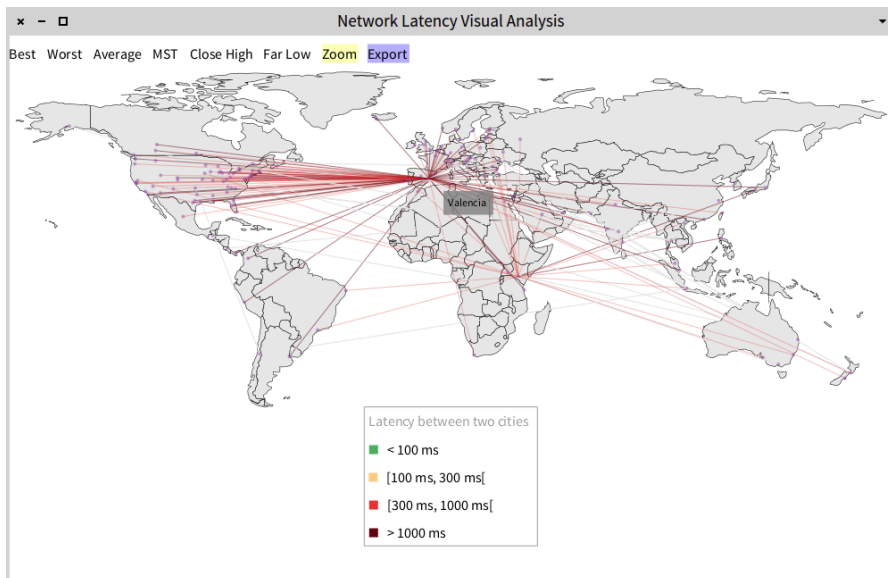


Figure 1.1: Worst latency for each city. Note that Valencia and Nairobi concentrate the most of the worst endpoint for connections.

city where is the city with which it has the best connectivity (in term of latency). The color of the edge between the cities represents the latency of the connection. Green edges are the best connections; orange ones means an average quality, red ones are worse and dark red edges mean the worst latencies (the ranges are split in $[0,100\text{ms}[$, $[100\text{ms}, 300\text{ms}[$, $[300\text{ms}, 1000\text{ms}[$ and greater than 1000ms).

There is also another visualization which is not triggered from the menu. By left-clicking in any city (purple circles) the user can see the latency from that city to the rest of the cities in the world (included in the dataset).

These visualizations show data which comes from a private company (see **Resources**) that offers a Internet Proxy service. The company exposes the status of their service in its website from which the data was manually collected. The data can be inaccurate (the snapshot was taken at the end of 2014), so its only purpose is to be a proof-of-concept for showing how network latency data can be analysed in a visual manner.

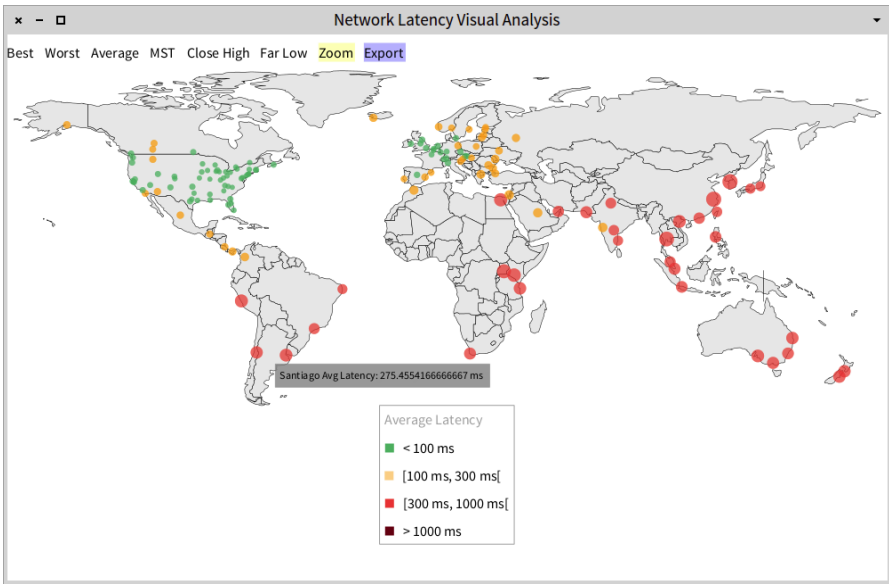


Figure 1.2: Average latency for each city.

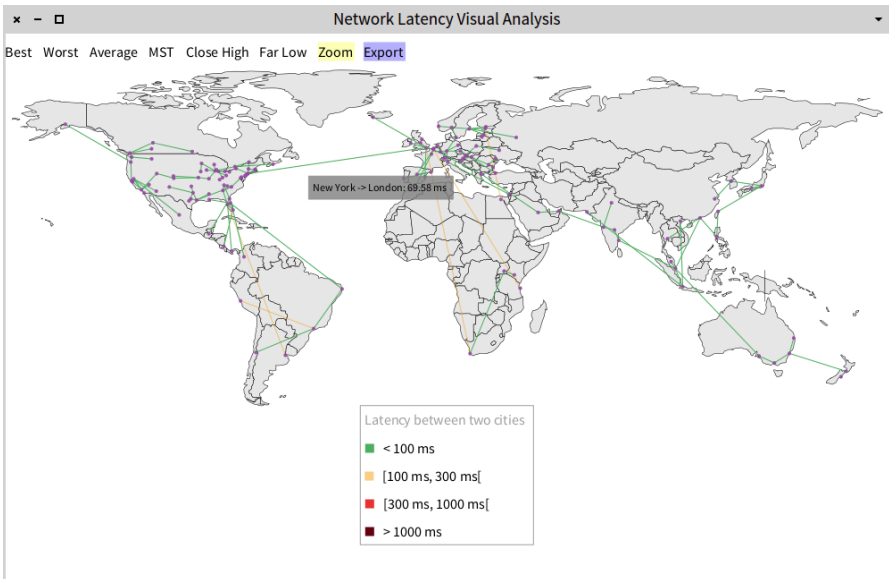


Figure 1.3: MST shows what probably are physical wires connecting cities.

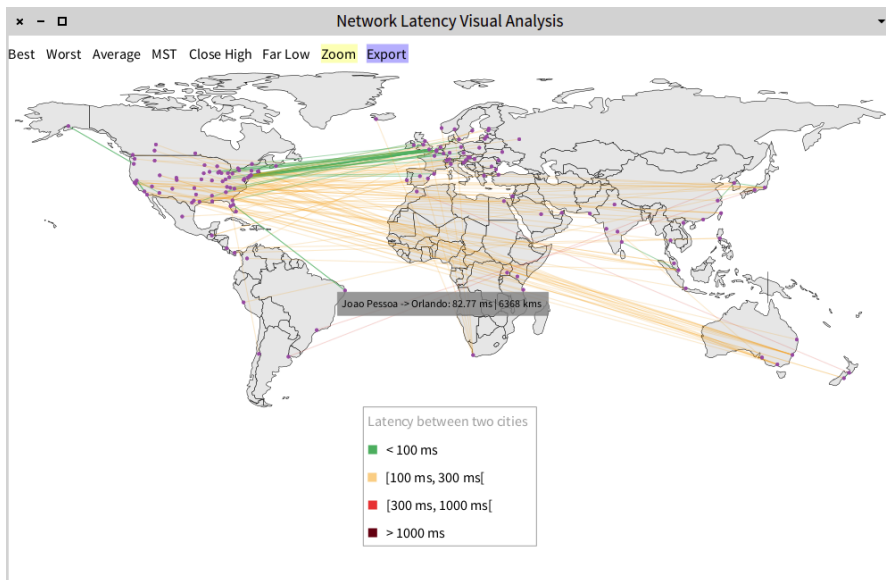


Figure 1.4: FarLow. Far cities that show low latency. Most of east cost cities of the US are the best endpoints for European connections. A couple of unexpected results Joao Pessoa-Orlando and Portland-Anchorage.

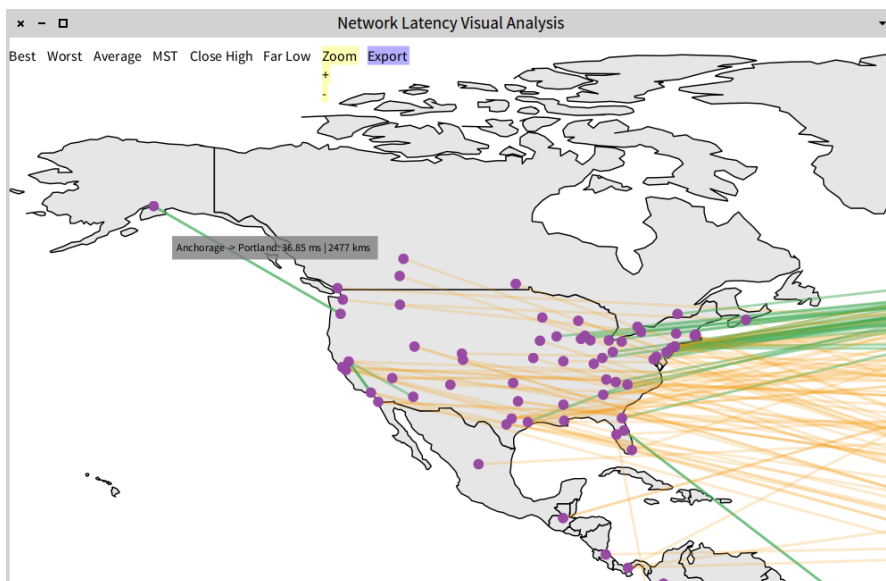


Figure 1.5: Zooming-in on the Portland-Anchorage latency.

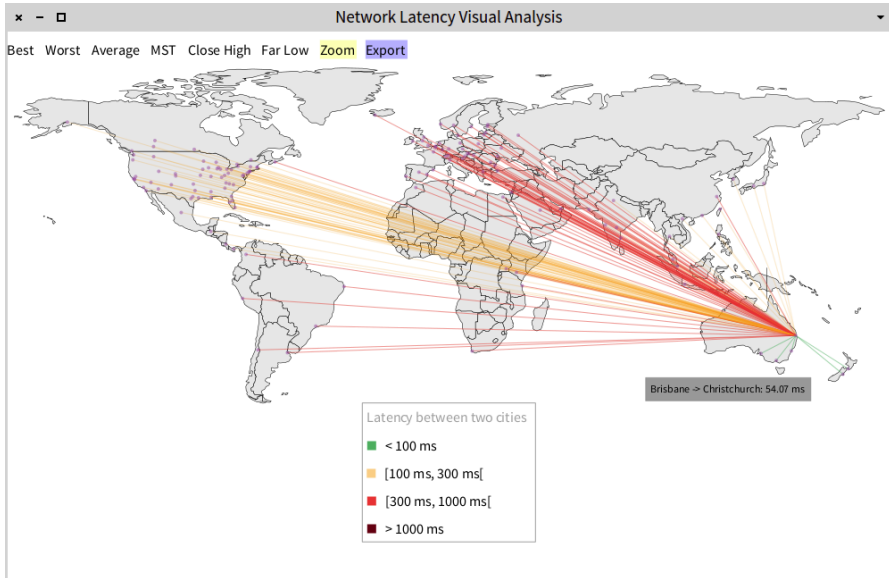


Figure 1.6: Latencies from Brisbane.

1.5 Visualizing your own connectivity

NLVA also includes a feature to allow users to test their own connectivity. To run it you have to right-click in the map approximately where is your location (the city can be already included in the map or not). A warning will popup asking you if you really want to run the test. It takes just a couple of seconds. It will show edges from your location to the rest of the cities included in the visualization (specifying their latency). An interesting analysis is to compare those results with the integrated data that comes with tool —caveat: it can only be done for the cities which comes with the visualization (left and right-click).

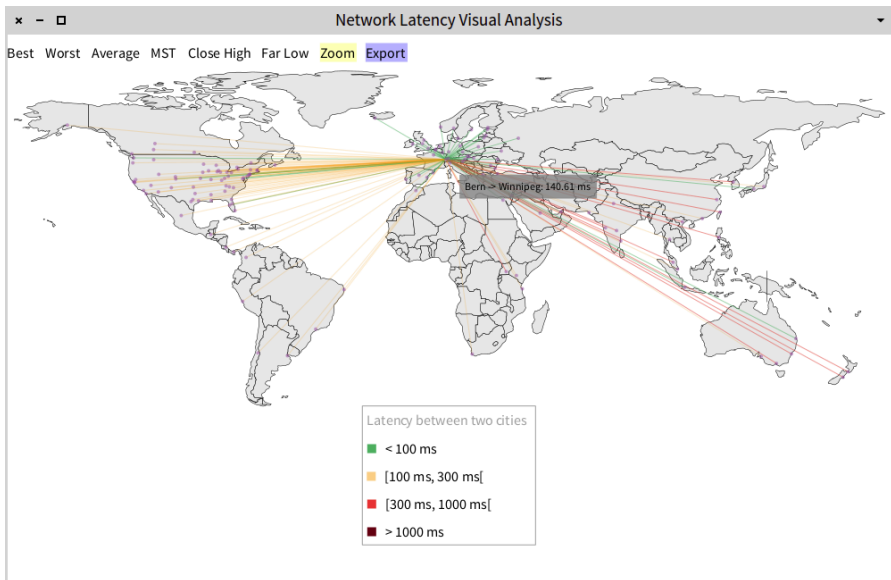


Figure 1.7: Latencies from "Bern", my actual location.