

A visualization tool for geographic information of NTP servers

Jonatan Schroeder*

University of British Columbia

ABSTRACT

Clock synchronization is an important and complex task in distributed systems. Understanding the state of NTP servers, partly responsible for this synchronization, may lead efforts and resources to focus on demanding regions. This work presents a visualization tool that allows an analyst to visualize and identify geographic regions that demand a greater number of NTP servers, as well as those with deficient servers.

1 INTRODUCTION

Clock synchronization is an important and complex task in distributed systems. The Network Time Protocol - NTP - is a protocol created to build and maintain the synchronization of computer clocks with the world real time. To do so, the NTP implements a dynamic logical network, presenting characteristics of a peer-to-peer network, which keeps itself self-organized, from the initial insertion of the computers, by means of frequent exchange of messages between peers. The NTP uses the Internet as a communication media for network nodes, and it is the standard protocol for synchronization of computer clocks connected to the Internet [6, 10].

The work presented in [10, 7] conducted a research about the synchronization of Internet hosts under the NTP. This research included data query, collection and analysis from several hundreds of thousands of NTP nodes. Many aspects that define the quality of time-keeping were collected. The results are available for download in [11].

The survey was taken in two rounds. In the first round, a NTP spider requested three kind of data in every NTP host: system information, peer list and monitor list. The peer list was used to continue the research, by adding new servers to be consulted. The NTP spider

started on August 30, 2005, querying an initial set of 263 public NTP servers listed on NTP Public Services Project. The collection finished on September 5, in the same year, discovering 1,278,834 unique IP addresses.

A second round was conducted to get more reliable data. This run started on September 20, 2005 and finished on September 30, 2005, discovering 11,895 new IP addresses, totaling 1,290,819 addresses. After preliminary analysis, some servers were discarded for insufficient information, and the research ended up with 147,251 complete response.

1.1 Available Data

There are several variables that were collected and are available, but only a small subset of variables is intended to be used; the remaining variables are not described in this proposal. In order to clarify the meaning of some of the variables, some definitions have to be described. The *stratum* of an NTP server is the level of the server in the NTP server. Servers with stratum 1 obtain time information from a reliable out source, such as an atomic clock. Servers with stratum 2 obtain information from a stratum 1 server, and so on. The *source server* of an NTP server is the server from which it obtains its time data, while the *root server* of an NTP server is the corresponding stratum 1 server from which the original time data for the server is obtained. The *dispersion* of a server is the maximum error to be considered in time data. This dispersion is mostly based on RTT (round-trip time) information between source and destination servers.

For each NTP server discovered in the survey, the following variables are available:

- *server*: the NTP server IP address;
- *proc*: text information about the server processor brand and model;
- *system*: text information about the operational system run in the server;
- *stratum*: level of the server in the topology;

*e-mail: jonatan@cs.ubc.ca

- *rootdelay*: calculated delay between server and its root server;
- *rootdispersion*: calculated dispersion between server and its root server;
- *peer*: identification of the server used as a source for time data (upper level in the hierarchy);
- *jitter*: server calculated jitter delay;
- *stability*: quality-related measurement of the server internal clock, i.e. the reliability of the time data in terms of time difference between two consecutive measurements.

Each NTP server with a stratum higher than 1 uses information from another NTP servers, and chooses the best data as a source. Each of these data collections in the server is called association. The available data also provides information about the associations for each server. The following variables are available for each association:

- *srcaddr*, *dstaddr*: the IP address of respectively the source server and the destination server;
- *stratum*: stratum information about the association;
- *rootdelay*: calculated delay between source server and its root server;
- *rootdispersion*: calculated dispersion between source server and its root server;
- *offset*: offset between source and destination servers;
- *delay*: calculated delay between source and destination servers;
- *dispersion*: calculated dispersion between source and destination servers;
- *jitter*: calculated jitter delay between source and destination servers.

1.2 Tasks

The papers that addressed the survey showed some results on the overall structure of the NTP network, such as delay overall distribution and distribution per stratum, overall dispersion and jitter [10, 7]. There was

no specific information about servers in a specific geographic region.

In this project I propose a solution to some specific tasks related to provide details on NTP servers in specific geographic locations, i.e., servers within a specific geographic region. The solution is proposed to fulfilled the following tasks:

- Show an overall visualization of the geographic topology of the NTP network, so that one can have a better understanding on how NTP servers are distributed worldwide;
- Find deficient NTP servers, i.e. servers with high delay or dispersion, or regions with low number of high quality NTP servers;
- Show an overall topology and find deficient NTP servers in a specific geographic region;

1.3 Report Outline

The rest of this report is organized as follows. Section 2 presents the proposed solution. Section 3 details some implementation issues, while section 4 describes some scenarios of use. Section 5 lists part of the related work. Finally, section 6 contains the conclusion and future work.

2 THE PROPOSED SOLUTION

In this project a new software was developed, focusing on the tasks described in section 1. This tool provides a geographic visualization in which both the number of servers and the distribution of servers according to a specific variable is displayed for the regions in the world, or the subregions in a specific region. The variable used for the analysis is initially set to the delay in the communication between the analyzed server and the root server that provides the time information for that server. The variable can be changed at any moment, as described below.

The visualization software proposed in this project has a main window containing two views, as shown in figure 1. The top area shows a representation of the region in focus as a map. For each subregion contained in the main region, a squared object is shown. This object has its area proportional to the number of servers in the region corresponding to the objects. A legend illustrating the number of servers corresponding to each size is shown in the right side of the map.

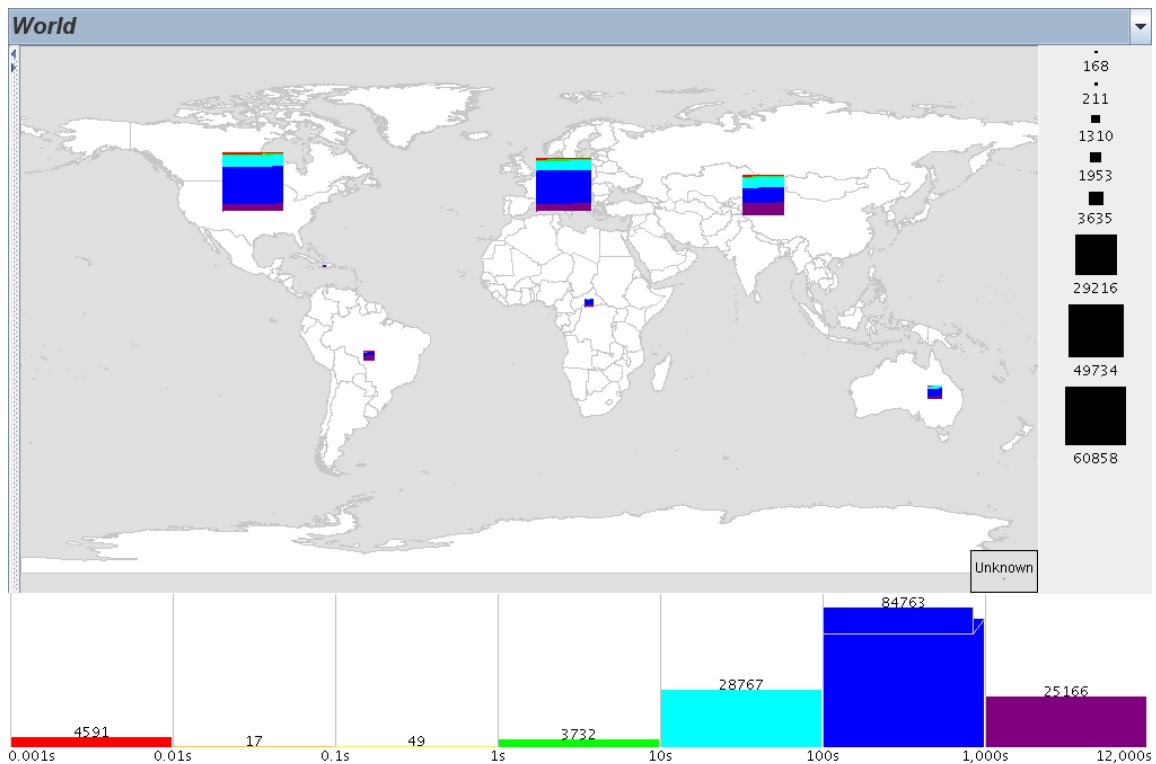


Figure 1: The main window of the software. The top part shows the active region (in this case the whole world), and squares representing the number of servers for each subregion. The bottom part shows a histogram of the servers in the region.

The regions comprehended in this project are divided in 5 levels. The first level corresponds to the entire world. The second level corresponds to continents, using the seven continents model (North America, South America, Central America, Europe, Asia, Africa and Oceania). The third level corresponds to countries. Countries whose territory is located in more than one continent are placed in the hierarchy under the continent in which most of their territory is located, e.g., Russia is placed in Asia, and Egypt in Africa. The fourth and fifth level are optional, and found only in some countries. These levels vary from country to country, but basically correspond to states or provinces and to counties, cities or towns.

The distribution of the delay of the servers is represented in the map view using color-coding in a rainbow-like range of colors. In each square, the proportion of servers with a specific range of delays is equivalent to the proportion of the area in the corresponding square with the color relative to the delay range.

The bottom area of the main window shows a histogram of all servers in the current region. In this his-

ogram, the horizontal axis represents the delay range, while the vertical axis represents the number of servers with the corresponding delay range. The horizontal axis will also use a color range corresponding to the same colors used in the geographic visualization, in order to ease the association between relative data in both visualizations.

The initial set of nodes used in the software comprehends the entire set of servers in the whole world. The user is able to select a different region by clicking on this region in the map view. When this action is performed, the map is zoomed to the selected region, and new squared objects are shown, now representing subregions of the new region. The histogram is also updated, representing the servers in the selected subregion. The total number of servers is still shown in the histogram, using a brighter color, to allow the user to visualize the relation between the region and the total set of servers. The result of this selection is shown in figure 2. If another subregion is selected, the process is repeated to the new region. To return to the upper level region, the user must click in the map view with

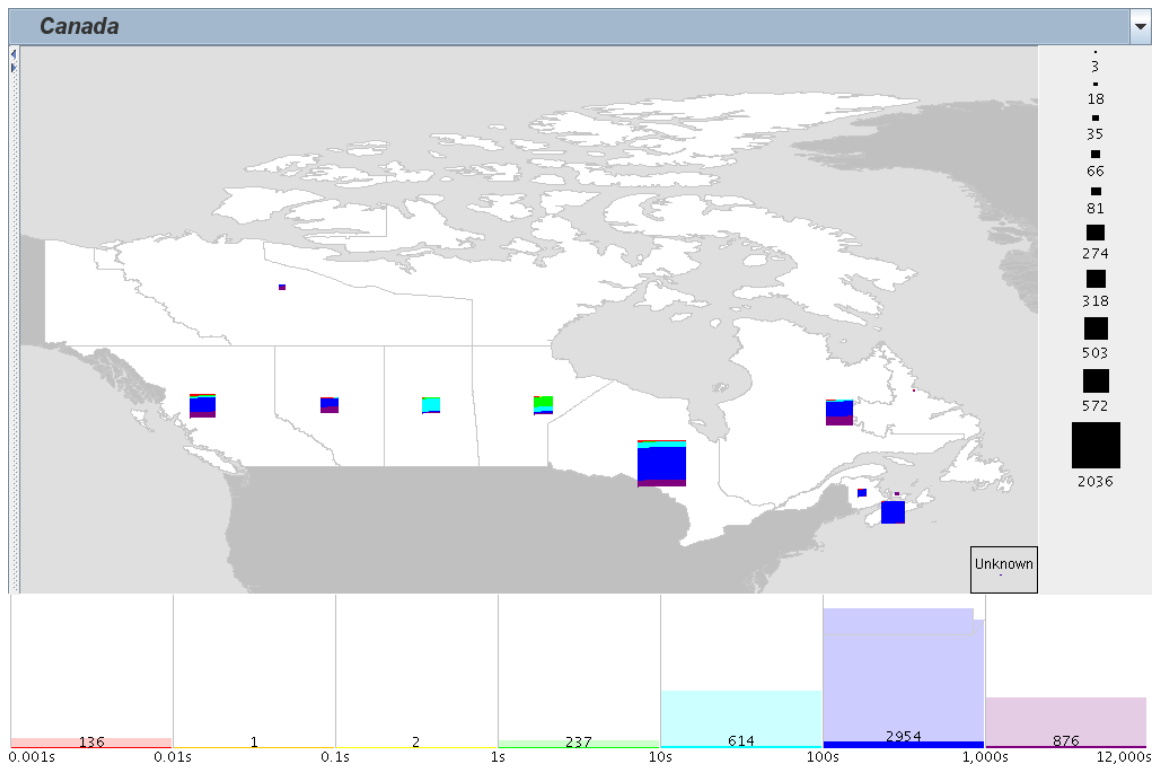


Figure 2: The results of the visualization after selecting the region 'North America' and after that the region 'Canada'. In this window, only servers in this region are shown in both views.

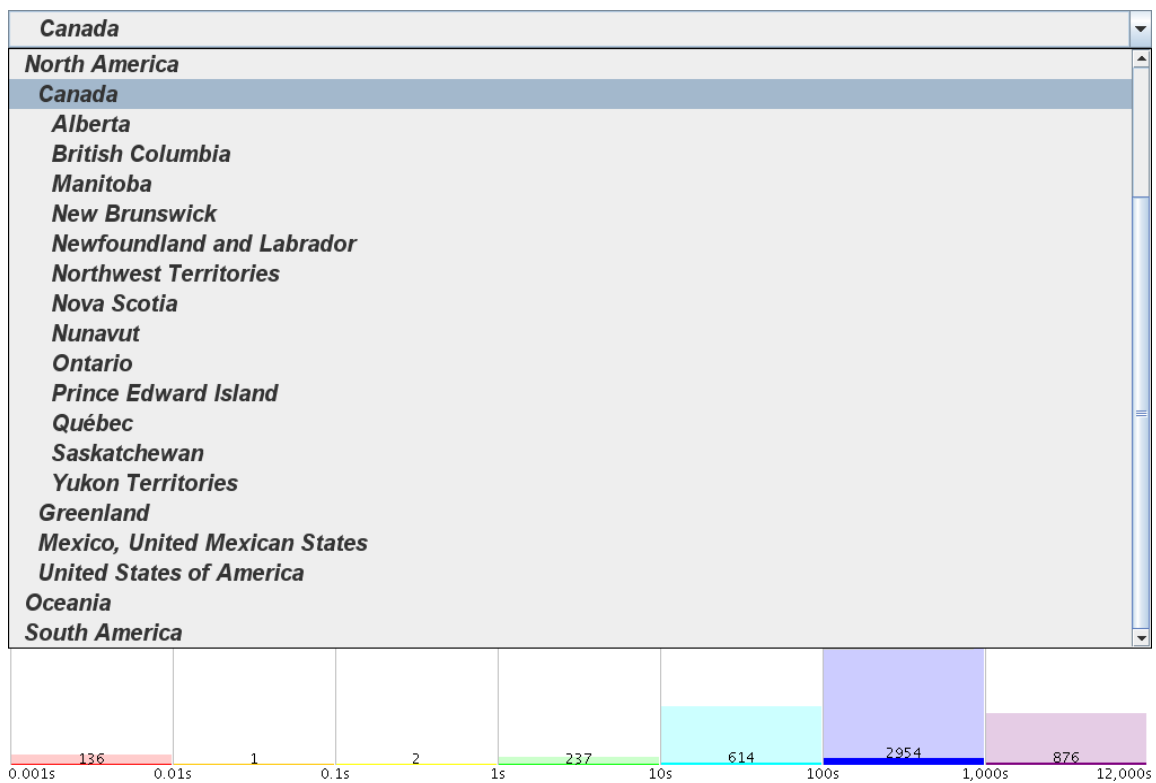


Figure 3: The selection of the subregion by name, using the combo box object available in the very top. This list includes the selected region, all its subregions, and all regions in the same level or above of the selected region.

the right button of the mouse.

Another resource for the selection of a specific region is available at the very top of the window. A drop-down combo box shows a list of all subregions in the selected region, as well as regions in the same level or above. This object is illustrated in figure 3. After selecting a region, both the map view and the histogram are updated to the set of servers in the selected region.

The user is also able to select servers within a delay range. This can be done by clicking the mouse on the representation of the desired range in the histogram view. The map view will be updated to show only servers with the specified delay range. Multiple sequential ranges can be selected at the same time by pressing the mouse button over the first range to be selected, and releasing it over the last range. The result of this selection is shown in figure 4. A mouse click using the right button of the mouse over the histogram will select the entire range of values.

Zoom and pan options are also available, and do not change the selection of servers. Zoom is available through the mouse wheel, while pan is available by dragging the mouse over the map.

On the left side of the map view, there is a hidden side panel, that can be opened with a click over the divider. This side panel is illustrated in figure 5. This panel shows information about the current map visualization. It also contains a component for the selection of a different variable for the analysis. With this object, the user is able to visualize servers based on the delay of the communication between the server and its root server (variable `ROOTDELAY`); the dispersion, or maximum error of the time information provided by the server (variable `ROOTDISPERSION`); the stratum, or the hierarchical level of the server in the NTP topology (variable `STRATUM`); or the stability, which is the measure of the quality of the local clock in the server (variable `STABILITY`). After the selection of a new measure, the map view is updated to show color-coding related to the new variable. The histogram is also updated, and the range intervals are changed to reflect the variations in the selected variable.

3 IMPLEMENTATION DETAILS

The tool described in this paper was developed in Java, using simple drawing methods and some Swing objects for the visualization [4]. Maps, color and size codings, histograms and legends were drawn using Java2D drawing methods.

Information about geographic location of specific IP addresses were obtained using GeoLiteCity [5]. This tool has location information about IP ranges in the levels of countries, states and cities. The free version of this tool provides 98% accuracy at the country level, and more than 70% accuracy at the city level for hosts in the United States, with similar accuracy in some other countries. The tool provides a datafile and a Java library to access the file.

A script written in Java was developed to read the survey files and generate a simpler file with all information that was actually supposed to be used in the solution. This script also obtained geographic information for each server, like latitude/longitude coordinates and country/region, using GeoLiteCity database access library, and included this information in the resulting file. The visualization tool used the resulting file instead of the original files, to reduce the latency of loading all server information.

Maps used in the geographic visualization were drawn using Java2D polygons, based on geographic boundary data available in specialized sites. The specific GIS information used in this project was obtained in two repositories. An overall information about boundaries in a country level was obtained from [2], while subnational boundaries, like states, provinces and cities, were obtained from [3]. The latter was not used in the country level due to the high level of detail, which demanded a higher processing time and memory usage for map drawing. These data provided boundaries in terms of longitude and latitude of each point of the polygon that describes a region. These latitude and longitude points were then converted to a specific location on screen and drawn as Java2D polygons. The continent level division data was obtained from [13]. In this level, no new boundaries were used, only the association between continent and country. Some adjustments had to be done in some files, due to some differences in country representations, like different ISO country codes for some countries.

There were some animation techniques employed in the development the software. The animation used for zooming and panning was implemented using slow-in slow-out technique, in which the speed of the animation is slower in the beginning and the end of the animation. A sin-curve was used to identify the step of animation in each moment of time, like the one shown in figure 6. To fasten the animation drawing, the map view was saved in an internal image, and only the image

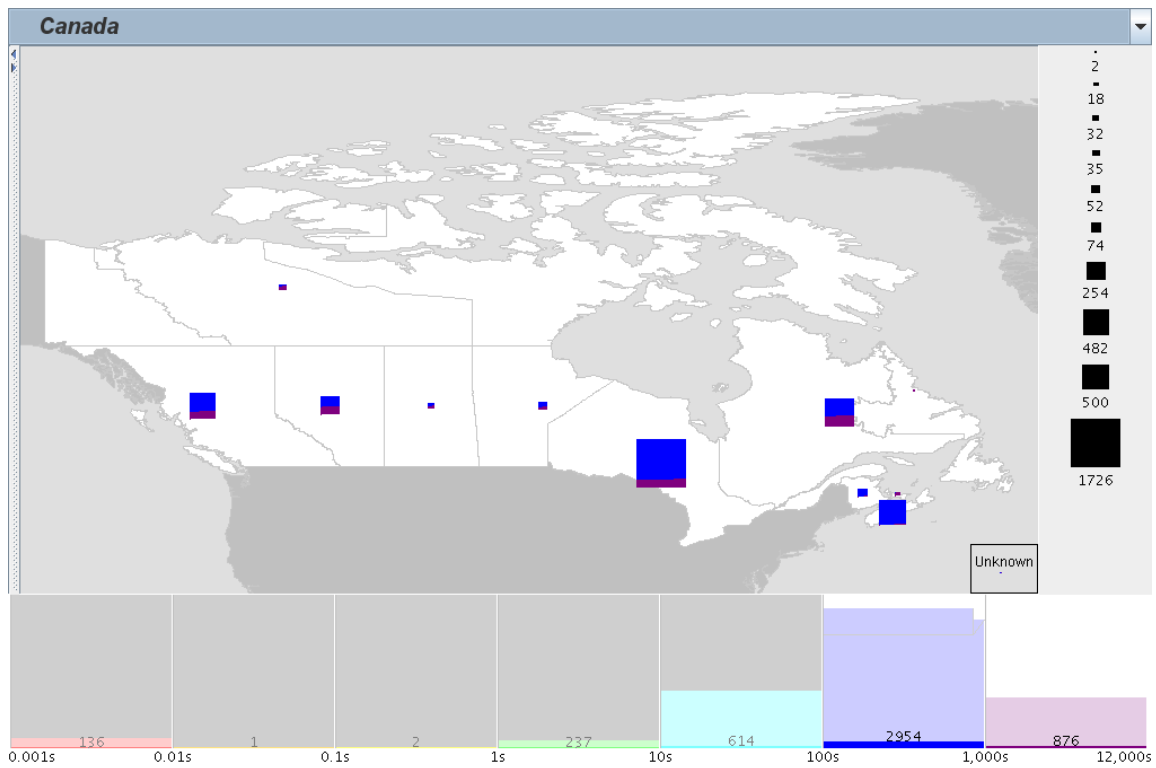


Figure 4: The results of the visualization after selecting a specific range of delays in the histogram view. In this window, only servers in this region are shown in both views.

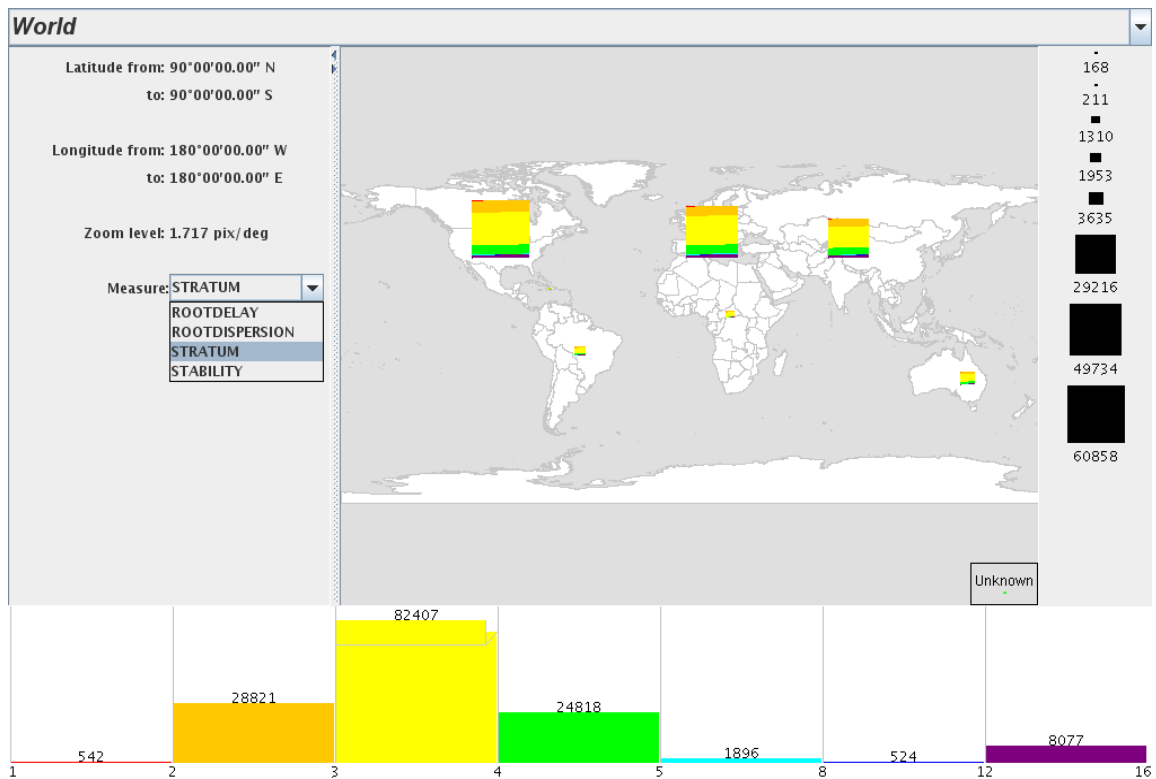


Figure 5: The side panel, available on the left side of the map. This panel shows information regarding the map visualization, and has components for the selection of a different metric for evaluation.




Figure 6: A sin curve, used for smooth animation effect.

was shown during the animation. Once the animation was completed, the screen was redrawn.

The representation of the number of servers based on squares was developed based on the count of servers for each variable range. The square is drawn pixel by pixel, and each pixel represents a specific number of servers. The color of the pixel is defined by the range in which the server is located. The pixels are drawn in order of the analyzed variable, from the first line to the last line of the square, and from left to right in each line. In order to adapt the size of each square to different situations, like regions with many subregions or with few subregions, the area of the square was defined as being:

$$\frac{n_i \times a}{k \times nr \times ns}$$

In this equation, n_i is the number of servers for the region being represented by the square, a is the total area of the view in pixels, nr is the number of regions being currently represented in the view, ns is the number of servers in the active region (the parent region corresponding to the entire current view) and k is a constant.

The range values for the histogram were manually defined for each variable, and have a fixed number of ranges. The intervals do not change for different regions or levels, so that the user can make a comparison between the region in focus and the whole set of servers. A possible improvement for future releases would be the automation of the detection of these intervals, as well as the recomputation of this automated detection when a range of values is selected.

4 SCENARIOS OF USE

As a scenario of use, suppose a user is looking for servers in Canada, trying to identify regions that lack reliable NTP servers. The intention of this user is to identify in what location in Canada a new NTP server will be most useful. After starting the software described in this paper, the user will click in the region 'North America', and then in the region 'Canada'. The resulting window is shown in figure 2. On this window, the user will be able to see how many servers are there in each province of Canada, and make a comparison to see which regions have a low number of servers. This information, combined with some background knowledge of the analyst, such as that the new server should be installed in the southern part of the country, can guide the user to deduce which region should receive more servers.

As part of the same scenario, the user may also notice that some regions, like Ontario, have a large number of servers with a considerably high delay. The user, then, can select the delay ranges that are considered high, and view only this servers in the map view. The result is shown in figure 4. The user is now able to identify in what regions to employ resources that could lead to a reduction of the general delay of the servers.

Another analyst may be interested in looking for servers with a specific stratum range. This analysis may help a server maintainer to improve the stratum of his own server or servers, by choosing a better alternative for the NTP associations used for the communication. For example, suppose the analyst knows his server has a high stratum, over level 10, and wants to change it to a level lower than 5. The user has to find a proper association with stratum at most equal to 4. By changing the used variable in the side panel, as shown in figure 5, the user will have an overall visualization of the servers and the distribution of servers per stratum. By filtering the desired stratum range and selecting a proper region, the user will be able to make a proper choice of where to look for a new association.

5 RELATED WORK

The data used in this work was obtained in a NTP survey described in [10]. Some results are shown in this work and in [7]. Most of the results are related to servers in general, without any specification of geographic location or network. These papers present data using cumulative graphs like PDF (probability distribution function), CDF (cumulative distribution function)

and CCDF (complementary CDF), as well as comparing several criteria in the set of servers.

Some information visualization techniques were employed in this project. The idea of keeping coordinated and multiple views is explored in many papers on CMV conference, such as [8]. Some techniques for smooth zooming and panning are described in papers like [1, 12]. The use of multiple colors for high-frequency color change is described by Rogowitz and Treinish in [9].

6 CONCLUSION

This project presented a tool for a geographic visualization and analysis of the location and distribution of NTP servers in the world or in specific geographic regions. This tool provides resources for the visualization and search of deficient servers, as well as the visualization of regions with lack of servers.

This tool can be improved by some future work. Many topics discussed previously were not implemented due to time. One of these topics is the visualization of details for specific servers, based either on a IP-based search or on the results of the described visualization. Beside that, some additional filters, such as an IP range filter for the identification of networks and autonomous systems in the world, could be implemented.

Some other improvements may include: the simultaneous selection of ranges in multiple ranges, like low delay and low dispersion; the automation of the number and size of the ranges in the histogram, both in the initial selection and a refinement after the selection of a specific range; the reduction of the latency, currently high in some regions due to the size of the dataset; and some usability issues, like hints and help windows.

REFERENCES

- [1] C. Appert and J.-D. Fekete. Orthozoom scroller: 1d multi-scale navigation. In *Proceedings of the Special Interest Group on Computer-Human Interaction (SIGCHI)*, pages 21–30, 2006.
- [2] E. Data and Maps. Countries 2002, 2003. <http://gis.esri.com/>.
- [3] R. Hijmans, N. Garcia, J. Kapoor, A. Rala, A. Maunahan, and J. Wiecek. GADM, Global Administrative Areas (Boundaries). <http://biogeo.berkeley.edu/gadm/>.
- [4] Java Technology. <http://java.sun.com>.
- [5] MaxMind LLC. GeoLite City. <http://www.maxmind.com/app/geolitecity>.
- [6] D. L. Mills. Internet time synchronization: The network time protocol. In *Zhonghua Yang and T. Anthony Marsland (Eds.), Global States and Time in Distributed Systems, IEEE Computer Society Press*. 1994.
- [7] C. D. Murta, P. R. Torres Jr., and P. Mohapatra. Characterizing quality of time and topology in a time synchronization network. In *49th*

- IEEE Global Telecommunications Conference, IEEE GLOBECOM*, San Francisco, CA, Nov. 2006.
- [8] J. C. Roberts. Coordinated and Multiple Views in Exploratory Visualization. In *Proc. Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV)*, 2007.
- [9] B. E. Rogowitz and L. A. Treinish. How not to lie with visualization. *Computers In Physics*, 10(3):268–273, May/June 1996.
- [10] P. R. Torres Jr. Caracterização da Rede de Sincronização na Internet. Master's thesis, Universidade Federal do Paraná - UFPR, Curitiba, Feb. 2007.
- [11] P. R. Torres Jr. and C. D. Murta. NTP Survey 2005. <http://www.ntpsurvey.arauc.br>.
- [12] J. J. van Wijk and W. A. Nuij. Smooth and efficient zooming and panning. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 15–22, 2003.
- [13] Wikipedia. List of countries by continent (data file) — Wikipedia, The Free Encyclopedia, 2004. [Online; last updated Oct 24, 2007].