

PLANETLAB EUROPE AS GEOGRAPHICALLY-DISTRIBUTED TESTBED FOR SOFTWARE DEVELOPMENT AND EVALUATION

Dan KOMOSNY¹, Shaoning PANG², Jan PRUZINSKY¹, Pavol ILKO¹, Jakub POLASEK¹

¹Department of Telecommunications, Faculty of Electrical Engineering and Communications, Brno University of Technology, Technicka 12, 616 00 Brno, Czech Republic

²Centre of Computational Intelligence for Cyber-Security, Department of Computing, Unitec Institute of Technology, Auckland 1025, New Zealand

komosny@feec.vutbr.cz, ppang@unitec.ac.nz, xpruzi00@stud.feec.vutbr, xilkop00@stud.feec.vutbr, xpolas27@stud.feec.vutbr

DOI: 10.15598/aeec.v13i2.1245

Abstract. *In this paper, we analyse the use of PlanetLab Europe for development and evaluation of geographically-oriented Internet services. PlanetLab is a global research network with the main purpose to support development of new Internet services and protocols. PlanetLab is divided into several branches; one of them is PlanetLab Europe. PlanetLab Europe consists of about 350 nodes at 150 geographically different sites. The nodes are accessible by remote login, and the users can run their software on the nodes. In the paper, we study the PlanetLab's properties that are significant for its use as a geographically distributed testbed. This includes node position accuracy, services availability and stability. We find a considerable number of location inaccuracies and a number of services that cannot be considered as reliable. Based on the results we propose a simple approach to nodes selection in testbeds for geographically-oriented Internet services development and evaluation.*

Keywords

Delay, development, Europe, evaluation, location, PlanetLab, service, software.

1. Introduction

PlanetLab (www.planet-lab.org) [3], [14] is a network of the dedicated servers running services accessible by users. Its primary purpose is to support development and evaluation of new Internet technologies such as peer-to-peer systems, overlay routing, distributed storage, and topology mapping. The PlanetLab nodes are

grouped into sets, called *slices*. A node associated with a slice runs a Linux virtual machine (Linux distribution CentOS, www.centos.org), called *sliver*. A user can remotely login on a sliver and run their code on it. The users run their code in separate working environments by accessing different slivers on the same physical PlanetLab node. Various tools exist to simplify the management of the slices and to upload the users' code on the slivers, such as Plush, PIMan, Stork, and PLDeploy (www.planet-lab.org/tools).

PlanetLab consists of about 1100 nodes that are associated with over 500 sites distributed over the world. For each site, the hosting organization name, web page, number of running nodes, delivery address, and geographical position is provided. The location of the nodes is inherited from the site's position information.

PlanetLab is divided into four main autonomous branches based on the geographical distribution of the sites [20]. PlanetLab Central or USA (referred as PLC, www.planet-lab.org) is the main authority covering the nodes in the USA and, also, the nodes not assigned to any specific branch. The European nodes belong to PlanetLab Europe (PLE, www.planet-lab.eu), the nodes in Japan are associated with PlanetLab Japan (PLJ, www.planet-lab-jp.org), and the nodes in Korea are associated to Private PlanetLab Korea (PPK, www.planet-lab.kr). PlanetLab Europe currently consists of about 350 nodes at 150 sites and the geographical distribution of the PlanetLab Europe sites is shown in Fig. 1. The figure omits a site in Iceland for a better visibility of the nodes.

A PlanetLab site covers (run) one or more nodes. These nodes share some of the information provided for the PlanetLab site, such as its geographical location. Linux virtual machines (slivers) are run on each

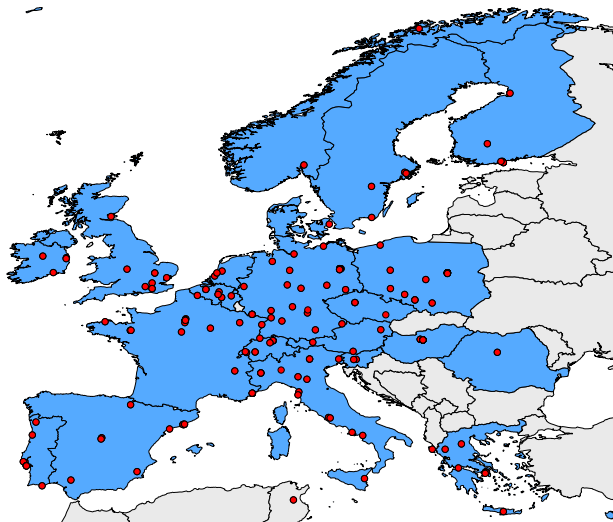


Fig. 1: PlanetLab Europe sites.

node to implement the PlanetLab services. The commonly used services are the remote login to a sliver using SSH (Secure Shell) and node echo request/reply (ping). PlanetLab creates a new sliver and runs it on a node when the node is assigned to a slice by a user as shown in Fig. 2.




Fig. 2: PlanetLab sites and nodes.

A slice is created by the PlanetLab administrators based on the user’s request. After establishing a slice, a user selects the nodes to be included in the slice based on the research purpose of the slice. Nodes from different sites can be added to a slice, and it is common that a node is assigned to a number of slices and, thus running a set of slivers at the same time. This scenario can be seen in Fig. 3. However, from the user’s (slice) point of view, a sliver equals a node. By assigning nodes to a slice, a user is able to operate the nodes by a remote login. The nodes assigned to a slice can be further grouped into subsets by assigning them different purpose-tags.

Due to remote access and geographical diversity of the nodes, PlanetLab is commonly used as a testbed for research and evaluation of geographically-oriented Internet services [6], [13]. However, correct and not-misleading results of such experiments and evaluations require trustful input data such as the nodes’ position (commonly referred as ground-truth data).

In this paper, we analyse the use of PlanetLab Europe as a geographically distributed testbed for Inter-

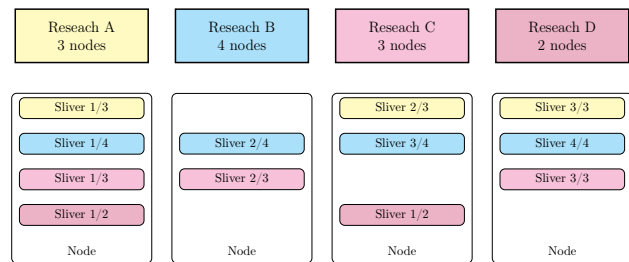


Fig. 3: PlanetLab slices, slivers, and nodes.

net research by studying the accuracy of the location information provided for the PlanetLab sites and, consequently, for the associated nodes. We also address the problem of the PlanetLab services availability and stability.

We used a number of techniques to evaluate the nodes’ location accuracy. These were the following:

- location’s uncertainty caused by the used number of digits for the coordinates (latitude and longitude),
- the distance between the site’s position (coordinates) and site’s postal address,
- the distance between the site’s position (coordinates) and the location (coordinates) provided by public location databases.

The PlanetLab services, we analysed, were echo request/reply (ping) and remote login (SSH). We inspected the availability and performance of these two services. We evaluated the availability of the services for a period of time to reflect the changes in time. For remote login, we additionally evaluated the SSH connection establishment delay since it is critical for Internet services that require timely actions taken on a large number of remote nodes.

The evaluations identified a considerable number of the location information inaccuracies and a number of services that cannot be considered as reliable. Based on the results we proposed a simple method to assist with the selection of the nodes to be used in the geographically-distributed testbeds.

The paper is structured as follows: the section ‘Motivation and related work describes in more detail the need for a proper selection of the PlanetLab nodes for the geographically-distributed testbeds. Examples of the geographical-oriented research based on PlanetLab are described. Related papers concerning the PlanetLab geographical and service properties are referred and discussed. In section ‘PlanetLab analysis description’ we talk about the techniques used in the paper and give an overview of the nodes covered by the analysis. The results are described in the section ‘Planet-

Lab analysis results'. First, we focus on the geographical properties, then we continue with the services performance and stability. The section 'PlanetLab nodes preferences' covers the proposal for the use of the PlanetLab nodes. Also, we present statistics for the nodes meeting the proposed criteria in this section. Finally, we conclude the paper.

2. Motivation and Related Work

The PlanetLab nodes have been used in a lot of Internet research. Research into geographically-oriented Internet services covers a broad variety of fields, such as estimation of IP node's location (geolocation), network topology design and discovery, assurance of local copyright and intellectual property laws of the resources available on-line, geographical organization of P2P VoIP networks, overhead traffic reduction by predicting the communication parameters, and prevention of Internet security attacks. Some of the papers based on PlanetLab are [4], [5], [8], [10], [11], [18], [19], [21]. PlanetLab information is also commonly used as groundtruth data for comparison of different geolocation tools and applications [1], [6], [16].

However, research papers reviewed by us did not inspect nor consider the location accuracy of PlanetLab nodes used. However, in order to achieve correct results and facilitate geographically-oriented Internet research, it is important to use error-free input data and services that are stable over time. We, therefore, address this gap in this paper. This problem has also been identified by other research teams, and we next present an overview of the related papers.

The ELTE's Location Survey (www.planet-lab.eu/node/220) focused on the position accuracy of the PlanetLab nodes. The finding was that a relatively large number of the PlanetLab nodes had very inaccurate location information with errors of varying magnitude. After some improvements, the result was that 90 % of the PlanetLab sites had a correct location, although the error distance limits to consider a position to be correct were not mentioned. The authors of paper [12] studied the geographical position of the PlanetLab nodes. They used a community-based geolocation database 'hostip.info' to retrieve the positions of the PlanetLab nodes, and they identified the distance between the coordinates obtained from the database and the coordinates provided for the PlanetLab sites. The conclusion was that, in some cases, the PlanetLab location information was unreliable to be used in geographically-oriented Internet experiments. Paper [2] brought recommendations for improving the PlanetLab's geographical

diversity. The approach described was to establish new PlanetLab sites to reflect the behaviour of the Internet. The authors suggested using the routing and topological data provided by Caida and Skitter to expand the PlanetLab's geographical diversity. Paper [17] studies the geographical distribution and services stability of the PlanetLab nodes. The authors proposed recommendations of the node selection for long-time research or network measurement purposes. In our previous work [9] we analysed the accuracy of PlanetLab Europe coordinates using geolocation databases. This paper extends work [9] by a new analysis of the location accuracy and, also, by a new study of the PlanetLab services performance and their changes in time.

3. Description of PlanetLab Analysis

In this section, we describe our analysis methodology. For our work, we used different information sources. We used information from the PlanetLab Europe's official site (www.planet-lab.eu) as well as information from outside PlanetLab, such as information from geolocation databases MaxMind, IP2Location and IPIgence. We developed a number of Linux shell scripts to communicate with the PlanetLab nodes and to implement the related geolocation procedures. For geocoding (transformation of postal address into coordinates), we used Google API.

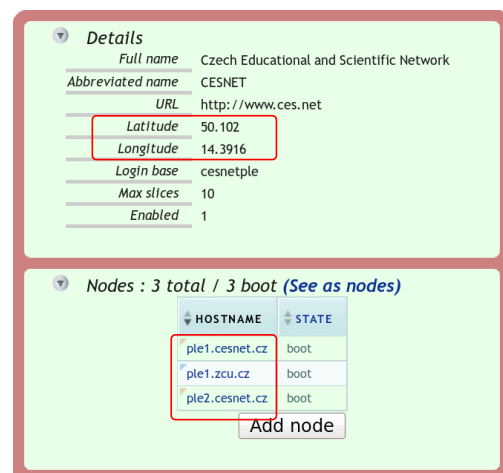


Fig. 4: Example of location information at PlanetLab Europe site.

Figure 4 shows the source of location information for the PlanetLab sites. It shows location data for a site named CESNET and located in Prague, Czech Republic. This site has three PlanetLab nodes associated (in the figure below the location data). All the nodes share the same data that is the hosting organization URL and

the geographical coordinates. We used URL to obtain the postal address of the organization. For the geocoding procedures we used API (The used code is based on The Google Geocoding API manual available at developers.google.com/maps/documentation/geocoding/) as shown in the command below:

```
wget -O siteLoc.json
"http://maps.googleapis.com/maps/api/
geocode/json?address=Zikova 4, 160 00
Praha 6, Czech Republic&sensor=false"
```

with this output:

```
"formatted_address" : "Zikova
1905/4, Czech Technical University in
Prague, 160 00 Prague 6-Dejvice, Czech
Republic",
"geometry" : {
"location" : {
"lat" : 50.1016729,
"lng" : 14.3907131
}
```

The listing demonstrates how we transformed the PlanetLab site postal addresses into geographical coordinates. The listing shows the coordinates obtained for CESNET site and shown in Fig. 4. The address ‘Zikova 1905/4, Czech Technical University in Prague, 160 00 Prague 6-Dejvice, Czech Republic’ was obtained from the URL specified for the site in the same figure (www.ces.net). This technique was used for absolute error distance evaluations.

In the analysis, we also studied the effect of coordinates rounding on location inaccuracy since less number of digits in the coordinate values means a greater area covered. For this purpose, we needed to calculate the distance uncertainty for each number of digits specified. In our analysis we used Eq. (1) for latitude uncertainty and Eq. (2) for longitude uncertainty:

$$LAT = \frac{1}{10^D} \cdot 55.64 \text{ (km)}, \quad (1)$$

$$LON = \frac{1}{10^D} \cdot 35.85 \text{ (km)}, \quad (2)$$

where *LAT* means latitude uncertainty, *LON* means longitude uncertainty, and *D* is number of digits after the decimal point in latitude and longitude coordinates respectively. If we found a PlanetLab site coordinates with different number of decimal digits for latitude and longitude values, we used the value with the lowest number of digits. If the number of decimal digits was equal, we used the latitude value since it produced a greater location uncertainty. Equation (1) and Eq. (2) use constant values of 55.64 km and 35.85 km taken from Tab. 1. The table shows distance uncertainty per one degree for different latitudes going through Europe.

We used values of latitude equal of a 50° since this goes through the area of central Europe.

In the analysis, we used location databases to find the geographical location for a given IP address. We used the demo pages of these databases: www.maxmind.com/en/home, www.ip2location.com/demo, and www.ipligence.com/. The demo pages allow to geolocate a certain number of IP address per day. The Maxmind demo service is based on the GooIP2 Precision product (www.maxmind.com/en/web_services) and allows up to 25 address per day to be geolocated. The limit for IP2Location demo service is 20 addresses per day and the product used is DB24 (www.ip2location.com/demo). The IPLigence demo service allows 50 queries per day and according to the information provided is based on the MAX product (www.ipligence.com/products/?max=#max).

We started to work with 367 PlanetLab nodes which seemed to be related to Europe. An overview of the nodes used is given in Tab. 2. After the first check of the nodes, we noticed some flaws in this set. We failed to resolve IP addresses for 19 domain names specified by PlanetLab (we call these nodes zombies). Next, we identified 14 nodes assigned to PlanetLab Europe (PLE), but geographically located outside Europe, for example Tunisia, Australia, Israel, and Thailand. On the other hand, we identified 29 additional PlanetLab nodes geographically located in Europe, but not assigned to PlanetLab Europe [15]. After leaving the nodes not geographically located in Europe and without an IP address, we obtained 334 nodes.

Tab. 2: Used nodes overview.

| | Node count |
|----------------------------|------------|
| Total | 367 |
| Without IP | 19 |
| Assigned to PLE, not in EU | 14 |
| Not assigned to PLE, in EU | 29 |

4. Results of PlanetLab Analysis

We divided our analysis into two parts. The first part is on the geographical properties of PlanetLab sites and the second part on the PlanetLab services.

4.1. Geographical Properties

The distribution of the PlanetLab nodes across the European countries is shown in Tab. 3. The locations used in the table were the coordinates of the PlanetLab sites provided by the PlanetLab Europe’s website. The ta-

Tab. 1: Latitudinal and longitudinal lengths for selected latitudes.

| Latitude | Lat. length / degree (km) | Long. length / degree (km) |
|----------|---------------------------|----------------------------|
| 30° | 110.885 | 96.486 |
| 45° | 111.132 | 78.847 |
| 50° | 111.229 | 71.696 |
| 60° | 111.412 | 55.800 |

ble does not show the nodes not assigned to PLE and nodes without an IP address (zombie nodes). Furthermore, the nodes belonging to a European country but not located in Europe are excluded. Such examples are the two nodes at Universite de La Reunion that are located on the Reunion island close to Madagascar. The island is politically part of France and, thus, the nodes are assigned to PlanetLab Europe.

The list shows that there are great differences in the nodes count in different European countries. The countries with the largest number of PlanetLab nodes are Germany, Spain, France, Great Britain, Italy, and Poland. For PlanetLab use as a geographically-distributed testbed, a uniform nodes distribution would be beneficial [2]. However, in practice, the nodes cover only specific areas. Based on the idea of the ITU ICT development index (The ICT development index is published by ITU - International Telecommunication Union, and is used to measure the information society status), which measures and compares the IT performance of different countries, we evaluated the support of the countries to the Internet research by providing resources in form of PlanetLab nodes. For each country, we calculated the ‘population/PlanetLab nodes’ ratio (Ppl/Plb) as shown in the Tab. 3. The result is that Switzerland, Cyprus, Finland, Hungary, and Slovenia are the most PlanetLab supportive countries.

Tab. 3: Nodes in countries.

| Country | Nodes | Popul. ($\times 10^3$) | Ppl/Plb ($\times 10^3$) |
|---------|-------|--------------------------|---------------------------|
| AUT | 6 | 8505 | 1417 |
| BEL | 14 | 11200 | 799 |
| CHE | 17 | 8100 | 477 |
| CYP | 3 | 900 | 289 |
| CZE | 5 | 10500 | 2102 |
| DEU | 55 | 80700 | 1468 |
| DNK | 2 | 5600 | 2814 |
| ESP | 25 | 46600 | 1864 |
| FIN | 12 | 5500 | 454 |
| FRA | 44 | 65900 | 1497 |
| GBR | 37 | 63700 | 1722 |
| HUN | 19 | 9900 | 520 |
| IRL | 8 | 4600 | 574 |
| ISL | 2 | 300 | 163 |
| ITA | 30 | 60000 | 2001 |
| NLD | 6 | 16800 | 2808 |
| NOR | 4 | 5100 | 1277 |
| POL | 32 | 38500 | 1203 |
| PRT | 16 | 10500 | 655 |
| ROU | 2 | 20100 | 10061 |
| SVN | 4 | 2100 | 516 |
| SWE | 10 | 9700 | 966 |

The countries in Tab. 3 are highlighted in Fig. 5. We divided the countries into four European parts (west, east, north, south) to indicate the diversity of PlanetLab nodes across Europe. Table 4 shows that the west part of Europe covers the majority of the nodes. A related paper [2] recommended a way to improve the PlanetLab’s unbalanced geographical diversity by adding new nodes/sites belonging to commercial organizations. These nodes should be assigned to a dedicated PlanetLab block to keep the routing distances small within the commercial and educational networks. Another suggested approach dealing with the PlanetLab’s diversity was to identify the existing GREN (Global Research and Educational Network) sites by examining routing data and motivate them to add the PlanetLab nodes.

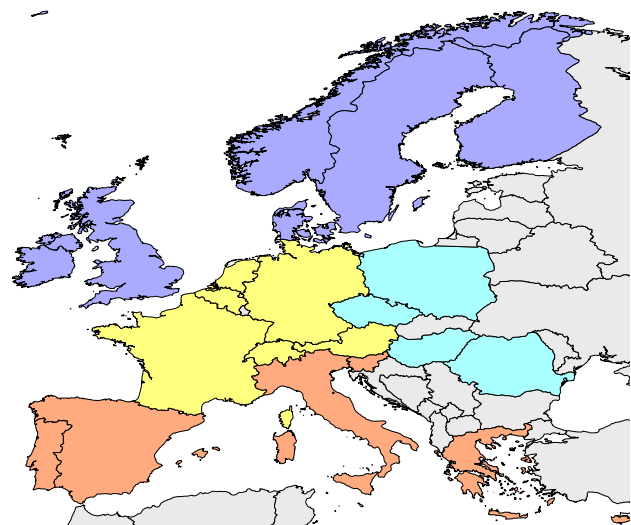


Fig. 5: Countries in four European parts with one or more PlanetLab sites.

Tab. 4: Nodes in Europe.

| Part of Europe | Node count |
|----------------|------------|
| East | 58 |
| West | 142 |
| North | 55 |
| South | 98 |

After the initial insight into PlanetLab Europe’s geographical properties, we continued with an analysis of the location information provided for the sites by the PlanetLab’s website [9]. We started with the position uncertainty caused by the number of the digits

used for specifying the coordinates (latitude and longitude). Some of the PlanetLab sites had the latitude and longitude specified using a low number of digits which resulted in erroneous locations, in some cases, of the tens of kilometres. The product of coordinates rounding is pointed out in Tab. 5. We observed that the rounding is not a major problem as it might have seemed. Approximately 90 % of the nodes have the coordinates' rounding error less than 5 km. However, attention should be taken not to use nodes with a possible rounding location error over 30 km, that is about 5 % of the nodes.

Tab. 5: Site location uncertainty due to coordinates rounding.

| Uncertainty (km) | Node count | Of all nodes (%) |
|------------------|------------|------------------|
| 55 | 8 | 2 |
| 35 | 8 | 2 |
| 6 | 16 | 4 |
| 4 | 10 | 3 |
| <1 | 325 | 89 |

The next information, we incorporated into our analysis, was the postal address of the organizations/sites running the PlanetLab nodes. We used this information to measure the difference between the provided coordinates and the site's postal address. We based this research on the idea that the organization's address and the location of the nodes associated should be close to each other in some extent to consider a node's location as trusted. Also, we were aware of the fact that an organization can reside in a different location than the nodes it owns. This could be the case of situating the servers in a dedicated server-room facility at another organization. However, we did not assume that this is common with the PlanetLab organizations since the vast majority are universities and large research centres. To obtain the distance between the nodes and their organization, we geocoded the postal address of each PlanetLab organization into coordinates. The result is shown in Fig. 6. The cumulative probability function shown indicates that about 70 % of the nodes are within 5 km distance from their organization's address. However, it can be seen that much greater distances were also measured.

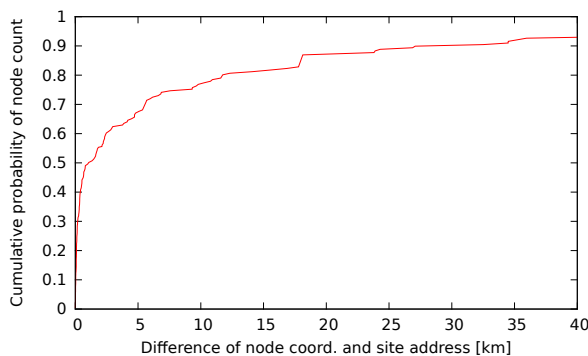


Fig. 6: Difference of site location and site postal address.

Up to this point, we used only information provided by the PlanetLab's website. In practice, a common way, to find a location of an Internet node, is to use a geolocation database. The geolocation databases maintain geographic locations for blocks of IP addresses. Upon a location request for an IP address, the corresponding block of IP addresses is found, and the searched node is assigned a position stored for this block. There is a broad variety of geolocation databases, some of them public or private. In our analysis, we used three public databases to find the positions of the PlanetLab nodes. Then we compared the locations (coordinates) from geolocation databases with the locations (coordinates) provided by PlanetLab. The databases used were:

- MaxMind (www.maxmind.com),
- IP2Location (www.ip2location.com),
- IPLigence (www.ipligence.com).

The results are plotted using the cumulative probability function in Fig. 7. The used databases gave similar results with Ip2location returning coordinates closest to the coordinates provided by PlanetLab. In terms of the distance difference between the PlanetLab location information and information from the databases, about 45 % of the nodes has difference lower than 5 km.

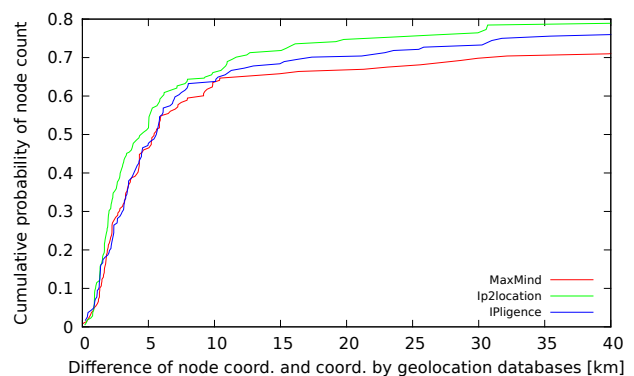


Fig. 7: Difference of site location and locations from geolocation databases.

The result of the location accuracy analysis is shown in Tab. 6. The column 'node count' displays the number of nodes which PlanetLab location information difference (inaccuracy for GPS rounding) was less than 5 km for the corresponding estimation method.

The probability function of the median of distance differences is shown in Fig. 8. The graph was generated using the location information sources shown in Tab. 6.

Tab. 6: Summary of geographical accuracy analysis.

| Location estimation method | Node count | Of all nodes (%) |
|------------------------------|------------|------------------|
| Coordinates rounding | 335 | 91 |
| Diff. to site postal address | 246 | 67 |
| Diff. to MaxMind | 160 | 44 |
| Diff. to IP2Location | 178 | 49 |
| Diff. to IPLigence | 164 | 45 |

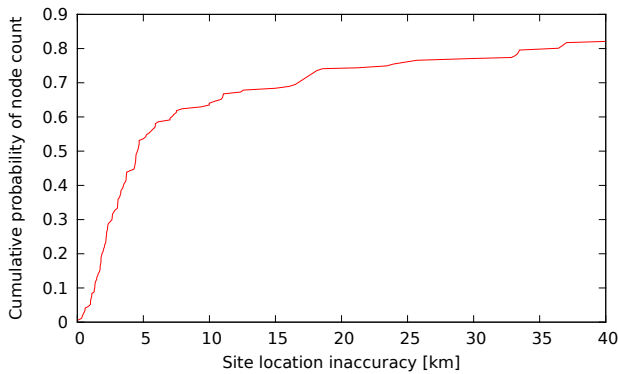


Fig. 8: Estimated site location inaccuracy.

4.2. Services Performance and Availability

In this analysis, we inspected the PlanetLab services used for geographically-oriented Internet research. We focused on echo-reply service and remote login as they are fundamental to any measurement-based geolocation service. We developed a number of BASH scripts, distributed them and ran them on the PlanetLab nodes. We used the standard ping and SSH (remote login) commands from the command line. We remotely connected to the PlanetLab nodes using a combination of private/public SSH keys, managed by the PlanetLab remote-login service. We evaluated the services for a period of time to reflect their performance and availability changes in time. In total, we performed around 400 measurements for each PlanetLab node at different times during a day.

Table 7 shows the node numbers for the SSH and echo-reply services. We considered a service as reliable if a node replied to at least of 95 % of the attempts over a period of three months. The numbers show a 'surprising' fact that only about half of the nodes were reliable for periodic latency measurements using the echo-reply service. Even a lower number of the nodes can be considered as reliable for a long-term use of the SSH service.

We linked ping and SSH services together (logical AND) and obtained that only 38 % of the nodes could be considered as reliable/stable for a longer-term use.

We also inspected how the performance of the services changed over time. For the evaluation period of

Tab. 7: Long-term service reliability.

| Service | Node count | Of all nodes (%) |
|--------------|------------|------------------|
| Ping | 198 | 54 |
| SSH | 153 | 42 |
| Ping and SSH | 140 | 38 |

three months, we plotted a graph shown in Fig. 9. The node numbers show that there is not a strong correlation between the ping and SSH service ($r = 0.64$). This indicates that the status of the both services should be checked independently. The graph also shows that about 250 nodes run the echo-reply service at a specific time with the number varying from 240 to 260. Regarding the SSH service, there were about 190 nodes available at a specific time with the number varying from 180 to 200 nodes.

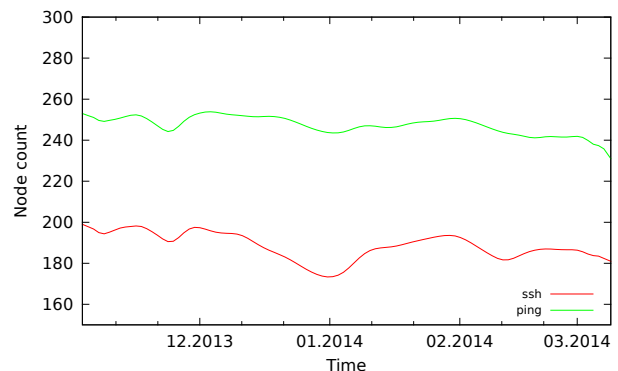


Fig. 9: Availability of ping and SSH services, changes in time.

When remotely accessing the PlanetLab nodes, we noticed another issue which should be definitely considered – the SSH connection establishment delay. We experienced some delays to be unexpectedly long. This significantly degrades the use of the nodes in time-sensitive Internet applications. We found this critical for the applications requiring a time-sensitive response from commands run on different machines, especially talking about timely measurements initiated from the PlanetLab nodes. In these cases, a number of seconds delay to connect a node plays an important role. We therefore measured the SSH connection delays using BASH scripts. The used script for the measuring SSH connection time to PlanetLab nodes is shown in the command below [7].

```

T="$(date +%s%N) "
ssh -o PreferredAuthentications=
publickey
-o PasswordAuthentication=no
-o ConnectTimeout=30
-o StrictHostKeyChecking=no
-o UserKnownHostsFile =/dev/null -l
sliceNameB -i privateKey user@IP exit
|| timeout =1
T="$(date +%s%N) -T) "
M="$( (T/1000000) ) "

```

Figure 10 plots a cumulative probability function of maximum, average, and minimum values for the SSH connection establishment delays. The graph covers the nodes we were able to at least once login into during the evaluation period – that was 204 nodes. Also, we defined a maximum delay of 30 seconds to consider a login to be successful. When a login was not established due to this limit, or the connection failed due to the other reasons, we set the delay to indefinite. This can be seen when inspecting the red line in the graph showing that the maximum connection delay was exceeded at least once for about 40 % of the nodes (81 nodes).

We observed that the average SSH connection delay for all the nodes was about 5 seconds and median about 4.5 seconds. We empirically suggest the long delays are caused by the reverse DNS lookup (determination of a domain name of the ssh client). We measured the time for the reverse DNS lookups on a local server and found that the time need was around 20 seconds.

5. Method for PlanetLab Nodes Selection

The results of the previous section show that definitely both:

- geographical information accuracy,
- service availability and performance,

of the PlanetLab nodes should be considered when the geographically-oriented Internet research is involved. Consequently, in this section, we combine the analysis results to reveal the complex information about the PlanetLab nodes. We propose a simple method to assist with the PlanetLab Europe nodes selection to be used in testbeds for geographically-oriented Internet research and evaluations.

We start with the geographical aspects. We consider location information from PlanetLab to be correct if it is within a 5 km distance difference from the median values of all the location accuracy results for a node.

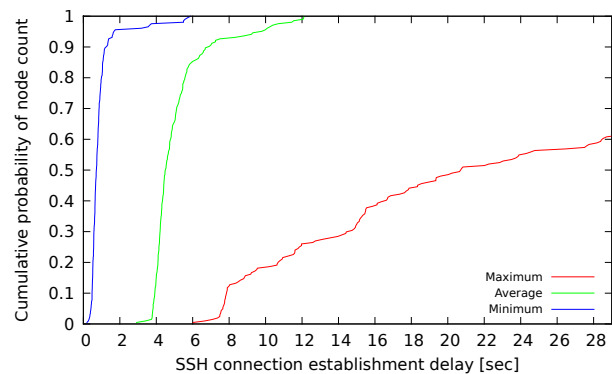


Fig. 10: Probability of node count vs SSH connection establishment delays.

We used the median of all location results for a node to filter out the extreme values. We would like to note that other distance difference thresholds can be set satisfying the needs of the specific application.

For the services, we keep our previous assumption to consider a service reliable to be available for at least of 95 % of the attempts over a specified period of time. We consider this as an important property since when producing correct outputs for longer evaluation periods, the same nodes should be used. We also involved the SSH establishment delay into our proposal. We empirically set the delay limit to be less than 5 seconds to filter out the nodes that might cause problems in time-sensitive applications.

Table 8 unveils the details about the parameters. We got 211 nodes with the correct location information provided. That was 60 % of the nodes geographical located in Europe. From this set, 83 nodes (40 %) provided reliable services. Additionally, these 83 nodes were assigned to the correct European PlanetLab branch (PLE), had IP address and the domain name could be resolved. Finally, from this set of nodes, we got 56 nodes (67 %) with an acceptable average SSH login delay.

6. Conclusion

PlanetLab as a geographically distributed testbed has some issues which should be considered prior to its use. Based on our experiences of using PlanetLab and the related work, we observed that the location information along with the services stability plays an important role in Internet research experiments. In the paper, we evaluated the accuracy of the location information provided for the PlanetLab Europe nodes. The results show that many nodes' location could not be trusted as correct input data (groundtruth) for experiments. Our other consideration was the availability and stability of

Tab. 8: Nodes preference parameters.

| Parameter | Node count | Percentage |
|------------------------------------|------------|-------------|
| Location diff. - median <5 (km) | 211 | 60 % of 353 |
| Service availability >95 (%) | 83 | 40 % of 211 |
| Connection delay - average <5 sec) | 56 | 67 % of 83 |

the PlanetLab services. We observed many nodes with the services unavailable or very delayed.

Based on the results we proposed a simple way for the selection of the PlanetLab nodes to be used in geographically-distributed testbeds. The paper can be useful for the readers and the PlanetLab users in the following: i) to know the PlanetLab nodes' problematic parameters to be checked before their use, ii) to know the PlanetLab node numbers available with the required location accuracy and service performance, and iii) to apply a simple method for assigning the PlanetLab nodes' preference for experiments.

Acknowledgment

Research described in this paper was financed by the National Sustainability Program under grant LO1401. For the research, infrastructure of the SIX Center was used. The work was also partially funded by the National Institute of Information and Communications Technology (NICT) Japan through a commissioned research project to the centre of Computational Intelligence for Cyber-Security (CICS) at Unitec Institute of Technology, New Zealand.

References

- [1] ARIF, M. J., S. KARUNASEKERA and S. KULKARNI. GeoWeight: internet host geolocation based on a probability model for latency measurements. In: *Proceedings of the Thirty-Third Australasian Conferenc on Computer Science*. Darlinghurst: Australian Computer Society Inc., 2010, pp. 89–98. ISBN 978-1-920682-83-5.
- [2] BANERJEE, S., T. G. GRIFFIN and M. PIAS. The Interdomain Connectivity of PlanetLab Nodes. Berlin: Springer, 2004, vol. 3015, pp. 73–82. ISSN 0302-9743.
- [3] BAVIER, A., M. BOWMAN, B. CHUN, D. CULLER, S. KARLIN, S. MUIR, L. PETERSON, T. ROSCOE, T. SPALINK, M. WAWRZONIAK. Operating Systems Support for Planetary-Scale Network Services. In: *Proceedings of the 1st conference on Symposium on Networked Systems Design and Implementation*. Berkeley: Usenix association, 2004, pp. 19–19.
- [4] BURGET, R., D. KOMOSNY and K. GANESHAN. Topology aware feedback transmission for real-time control protocol. *Journal of Network and Computer Applications*. 2012, vol. 35, no. 2, pp. 723–730. ISSN 1084-8045. DOI: 10.1016/j.jnca.2011.11.005.
- [5] HEGR, T., L. BOHAC, Z. KOCUR, M. VOZNAK and P. CHLUMSKY. Methodology of the direct measurement of the switching latency. *Przeglad Elektrotechniczny*. 2013, vol. 89, no. 7, pp. 59–63. ISSN 0033-2097.
- [6] HUFFAKER, B., M. FOMENKOV and K. C. CLAFFY. Geocompare: a comparison of public and commercial geolocation databases - Technical Report. In: *CAIDA: Center for Applied Internet Data Analysis* [online]. 2011. Available at: <http://www.caida.org/publications/papers/2011/geocompare-tr/geocompare-tr.pdf>
- [7] ILKO, P. *Geographic location of PlanetLab servers*. Brno, 2014. Master thesis. Brno University of Technology. Supervisor ass. prof. Ing. Dan Komosny, Ph.D.
- [8] KOMOSNY, D., J. BALEJ, M. SIMEK, P. DOLEZEL, H. SATHU and R. SHUKLA, Cable Length Based Geolocalisation. *Przeglad Elektrotechniczny*. 2012, vol. 7a, iss. 1, pp. 26–32. ISSN 0033-2097.
- [9] KOMOSNY, D., J. PRUZINSKY, P. ILKO, J. POLASEK and K. OGUZHAN. On Geographic Coordinates of PlanetLab Europe. In: *37th International Conference on Telecommunications and Signal Processing (TSP)*. Berlin: IEEE, 2014, pp. 229–233. ISBN 978-80-214-4983-1.
- [10] MORAVEK, P., D. KOMOSNY, R. BURGET, J. SVEDA, T. HANDL and L. JAROSOVA. Study and performance of localization methods in IP based networks: Vivaldi algorithm. *Journal of Network and Computer Applications*. 2011, vol. 34, iss. 1, pp. 351–367. ISSN 1095-8592. DOI: 10.1016/j.jnca.2010.06.014.
- [11] MORAVEK, P., D. KOMOSNY, M. SIMEK, J. SVEDA and T. HANDL. Vivaldi and other localization methods. In: *Proceedings of the International Conference on Telecommunications and Signal Processing (TSP)*. Dunakiliti: IEEE, 2009, pp. 214–218. ISBN: 978-963-06-7716-5.

- [12] OLBRICH, M., F. NADOLNI, F. IDZIKOWSKY and H. WOESNER. Measurements of path characteristics in PlanetLab. TKN-09-005, 2009. Technical University Berlin.
- [13] OneLab. Use case 1: Geographically distributed tested. Datasheet, 2009. UPMC Paris Universit as.
- [14] PETERSON, L., A. BAVIER, M. FIUCZYNSKI and S. MUIR. Experiences building PlanetLab. In: *Proceedings of the 7th USENIX Symp. on Operating Systems Design and Implementation (OSDI)*. Berkely: USENIX Association. 2006, pp. 351–366. ISBN 1-931971-47-1.
- [15] POLASEK, J. *Visualization of PlanetLab nodes location*. Brno. 2014. Master thesis. Brno University of Technology. Supervisor ass. prof. Ing. Dan Komosny, Ph.D.
- [16] SHAVITT, Y. and N. ZILBERMAN. A Geolocation Databases Study. *IEEE Journal on Selected Areas in Communications*. 2011, vol. 29, no. 10, pp. 2044–2056. ISSN 0733-8716. DOI: 10.1109/JSAC.2011.111214.
- [17] TANG, L., Y. CHEN, F. LI, H. ZHANG and J. LI. Empirical Study on the Evolution of PlanetLab. In: *Sixth International Conference on Networking (ICN'07)*. Martinique: IEEE, 2007, pp. 64–64. ISBN 0-7695-2805-8. DOI: 10.1109/ICN.2007.40.
- [18] VOZNAK, M., M. HALAS, F. REZAC and L. KAPICAK. Delay variation model for RTP flows in network with priority queueing. In: *10th WSEAS International Conference on SIGNAL PROCESSING, ROBOTICS and AUTOMATION (ISPRA '11)*. Cambridge: WSEAS, 2011, pp. 344-349. ISBN 978-960-474-276-9.
- [19] VOZNAK, M., J. ROZHON, H. A. ILGIN. Network Delay Variation Model Consisting of Sources with Poisson's Probability Distribution. In: *Proceedings of the 15th WSEAS international conference on Computers*. Corfu: World Scientific and Engineering Academy and Society (WSEAS), 2011, pp. 240-244. ISBN 978-1-61804-019-0.
- [20] KIM, W., A. ROOPAKALU, K. Y. LI and V. S. PAI. Understanding and characterizing PlanetLab resource usage for federated network testbeds. In: *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference - IMC '11*. New York: ACM Press, 2011, pp. 515–522. ISBN 978-1-4503-1013-0. DOI: 10.1145/2068816.2068864.
- [21] XIN, C. *Locating relay nodes for P2P VoIP applications: A PlanetLab-based experimental study*. Burnaby, Canada. 2008. Master thesis. Simon Fraser University. Supervisor Dr. Jiangchuan Liu.

About Authors

Dan KOMOSNY is an associate professor at Department of Telecommunications, Brno University of Technology, Czech Republic. He received his Ph.D. degree in Teleinformatics. His research is focused on geolocation in IP and sensor networks. He leads courses dealing with communication in IP networks and network operating systems.

Shaoning PANG is an associate professor at Department of Computing, Faculty of Creative Industries and Business, Unitec Institute of Technology, New Zealand. He obtained his Ph.D. in intelligent systems, from Shanghai Jiao Tong University, China (2000). In 2003, he completed his postdoctoral research at Pohang University of Science and Technology (POSTECH), South Korea.

Jan PRUZINSKY is an M.Sc. student at Department of Telecommunications, Brno University of Technology.

Pavol ILKO is a B.Sc. student at Department of Telecommunications, Brno University of Technology.

Jakub POLASEK is a B.Sc. student at Department of Telecommunications, Brno University of Technology.