Comparing the Accuracy of IPv4 and IPv6 Geolocation Databases

Jan-Jelle Kester
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
j.j.kester@student.utwente.nl

ABSTRACT

IPv4 geolocation has been around for some time, but IPv6 geolocation is relatively new. This raises the question as to how the accuracy of IPv6 geolocation compares to the accuracy of IPv4 geolocation. Since IPv6 geolocation is a reasonably new topic, litte research has been done on its accuracy. In this study the accuracy is measured using a ground truth consisting of very precise locations of known computer systems. Different IP geolocation databases were queried and the resulting locations were compared to the known locations for the collected IP addresses. It turns out that IPv6 geolocation is significantly less accurate than IPv4 geolocation. However, IPv6 geolocation shows potential to become at least as accurate as its IPv4 counterpart.

1. INTRODUCTION

This paper describes a method for measuring the accuracy of IP geolocation databases and discusses the results of performing these measurements on a select number of databases.

This section points out the need for this research and tries to establish a foundation on which to build the methodology as described in section 2. Section 3 contains a comparison of the IPv4 and IPv6 protocols from a geolocation standpoint. Section 4 shows which geolocation databases were used and describes the collected ground truth. The results, recommendations and conclusion can be found in sections 5, 7 and 6 respectively.

1.1 Motivation

IP geolocation databases contain information about the physical location of a computer system given its IP address. These databases are widely used for multiple purposes, ranging from location-aware advertising and web site analytics to intelligent routing and cyber crime research. Many people rely on IP geolocation which makes it a valuable technology. However, it is currently unknown to what level IPv6 geolocation results can be trusted.

The IPv6 protocol specification [3] was published in 1998. However, the protocol was never widely adopted by ISPs until recently, when the need for IPv6 became apparent due to the fact that all IPv4 address blocks had been allocated by the Internet Assigned Numbers Authority (IANA). This triggered the large scale deployment of this protocol to end users. Because IPv6 is not backwards compatible with IPv4 this rollout is not without obstacles. For example, systems only supporting IPv4 cannot connect to IPv6-only systems, and vice versa. Inherently, implementations of IPv4 and IPv6 will coexist on many devices. This coexistence allows for a fair comparison of IPv4 and IPv6 geolocation databases as both are tested with the exact same ground truth locations. This reasoning is explained in detail in section 2.

Because of the slow rollout of IPv6, many internet services were primarily focused on IPv4, including geolocation databases. However, IPv6 usage is on the rise. As of fall 2015, more than 7% of all Google users use IPv6, while this was less than 1% in 2012 [7]. Since this number will only grow in the upcoming years geolocation databases will need to have good support for IPv6 as many services rely on the location of visitors.

1.2 Research questions

Since IPv6 is newer and less widely used than IPv4, database creators have had less time and data to create accurate databases. This might result in less accurate geolocation of IPv6 addresses, which impacts applications using IP geolocation. With a growing number of users connecting to the internet using IPv6, geolocation using IPv6 addresses becomes more important. However, the state of IPv6 support of geolocation databases is currently unknown.

This study will focus on the current support of IPv6 by well-known geolocation databases. This main subject is divided into three research questions.

- Q1. What geolocation databases are available?
- Q2. What are major differences between IPv4 and IPv6 from a geolocation standpoint?
- Q3. What is the accuracy difference of geolocation databases when comparing IPv4 and IPv6 geolocation?

Questions Q1. and Q3. are targeted at determining the IPv6 support of IP geolocation databases while question Q2. aims to provide knowledge to possibly explain the observed difference(s) in the results.

1.3 Expected results

There are many IP geolocation databases available. However, the expectation is that not all IP geolocation databases provide IPv6 support. Furthermore, it is expected that all databases that support IPv6 are able to locate IPv6 addresses on a country level. Secondly, the average accuracy of the location of IPv6 addresses is expected to be less or

24^{th} *Twente Student Conference on IT,* January 22^{nd}, 2016, Enschede, The Netherlands.

equal than the average accuracy of IPv4 addresses. Finally, no database might be able to pinpoint the location of any IPv6 address more precisely than it pinpoints the corresponding IPv4 address.

1.4 Related work

In the past there have been studies on the accuracy of geolocation databases. Finding a ground truth has always been an issue for reliably measuring the accuracy of geolocation databases. Poese et al [14] used a complete routing table from a large ISP's backbone router as their ground truth. This resulted in 357 BGP prefixes with a known location of the advertising router inside the network of the ISP. This ground truth is likely to be correct on a city level.

Huffaker et al [8] used multiple datasets as ground truth. The first dataset contains the 398 nodes of the PlanetLab research network, which all have a precisely known geographic location. The second dataset contained 2,680 unique locations in France representing ADSL network regions, with in total over 6 million unique IP addresses. The third dataset consisted of DNS name to IP address mappings, complete with the heuristic to map the names to geographic locations, from a U.S. Tier 1 transit provider. This set contained 23,644 addresses. This combined ground truth is likely to be correct on a city level, but is focused on the U.S. and France since the major part of their ground truth consists of IP addresses from the last two sources.

The examples above show that researchers have been able to establish a ground truth for testing geolocation databases with IP addresses. However, these ground truths are biased to certain regions and/or do not have accurate locations for the IP addresses in the ground truth. Furthermore, these ground truths consisted of IPv4 addresses and are therefore not useful in this research. The methods for collecting these ground truths can be used with IPv6.

In 2012, Zander [20] showed that IPv6 support of geolocation databases was poor. In fact, on a country level the Regional Internet Registry (RIR) delegation data and the MaxMind GeoIP database were almost identical. One should note that at that time the MaxMind GeoIP IPv6 database was still under development.

While not researched extensively, a measurement based geolocation approach are likely to work equally well with IPv6 addresses as with IPv4 addresses [18]. This is expected as most measurement based approaches use routing data (e.g. BGP tables, traceroutes), which is not different when using IPv6 instead of IPv4. Examples are [4, 5, 6, 10, 11, 12, 17].

2. METHODOLOGY

Before the current level of IPv6 support of IP geolocation databases can be determined three steps need to be taken. First of all a list of suitable IP geolocation databases needs to be compiled (Q1.). Secondly a ground truth needs to be established to compare the database results against. Finally the database results must be compared with this ground truth to calculate the accuracy of the databases (Q3.).

2.1 Available databases

The geolocation databases used for testing are gathered by looking at previous researches using IP geolocation and internet searches. A database must support both IPv4 and IPv6 addresses as well as provide locations as latitude and longitude coordinates. This level of support is required for research question Q3. since this method allows for comparing IPv4 and IPv6 geolocation performance of one database instead of comparing multiple databases with eachother. Furthermore, the database must provide an API for querying IP addresses or the database must be in a readable format. Lastly, the database must be free to use or made available for research purposes.

2.2 Ground truth

Ground truth data is needed to calculate the accuracy of a geolocation database. The ground truth consists of pairs of an IPv4 and an IPv6 address belonging to the same system with a known geographical location. This allows for calculating the deviation of a database result from the real location.

As shown by Beverly et al. [1], finding pairs of IPv4 and IPv6 addresses which are in the same physical location is hard, so automatically collecting corresponding IPv6 addresses for a set of IPv4 addresses with a known physical location will not provide the necessary certainty about the physical location of the system listening to the address. Instead, a set of nodes was obtained from reliable data sources. The IP addresses and locations of 107 CAIDA Ark Monitors [2] and 15,610 RIPE Atlas Probes [16] were collected. Both CAIDE Ark and RIPE Atlas are research networks which means that these groups of machines are spread around the globe and used in research, including internet topology measurements. The physical location of the machines is relevant to these measurements so the GPS coordinates given for these machines are very likely to be precise. The locations of the Ark Monitors are, according to a CAIDA representative, normalized to the nearest airport or big city (e.g. the center of Amsterdam for any data center in the Amsterdam area). This means that the locations in this dataset are correct on a city level. The RIPE Atlas Probes are mainly located in offices and homes rather than in data centers. Every Atlas Probe owner manually submits the GPS coordinates of the node to RIPE by pinpointing the location on a digital map. The accuracy of these locations is therefore likely to be precise because owners will enter their building address instead of using a geolocation database to get the GPS coordinates.

Of all 15,717 collected nodes, only 3,206 nodes proved to be usable for this research (see section 4.2).

2.3 Database comparison

The collected geolocation databases are queried with the IPv4 and IPv6 addresses from the ground truth. Each individual result, consisting of the (ground truth) node, geolocation database, location of the IPv4 and IPv6 addresses and relevant distances is stored. Only results providing latitude and longitude coordinates are accepted since the public APIs of the used databases claim to support this. In case a database is not able to provide a result, this is stored as well. However, these results are omitted in the metrics to remove any unfair disadvantages due to the selection of IP addresses.

The following metrics will be gathered from these raw measurements:

- Number and percentage of incomplete measurements per database. A measurement is considered incomplete when a database returned no IPv4 and/or IPv6 location. In other words, a measurement is considered complete when both an IPv4 and an IPv6 location are returned by a database.
- Average (mean) deviation (per database) in kilometers between:

– The location in the ground truth and the IPv4 location according to the database.
– The location in the ground truth and the IPv6 location according to the database.
– The IPv4 location and the IPv6 location according to the database. This number will show any correlation between the IPv4 and IPv6 version of the database.

These deviations are calculated using the coordinates in the WGS84 geodetic system[19].

- Deviation spread: the percentage of measurements with a mean deviation with the ground truth within given distances, e.g., the percentage of measurements that were accurate within a range of a certain distance. These percentages will be calculated for the IPv4 and IPv6 versions of each database seperately to show the difference between IPv4 and IPv6 geolocation.

2.4 Automation

The measurements are performed using software specifically written for this research.

The platform is built as a web application using Python and Django (`https://www.djangoproject.com/`). Data is stored in a PostgreSQL database with PostGIS (geospatial) extensions (`http://postgis.net/`). This combination allows for fast development and easy data storage and representation. Celery, a Python distributed task runner (`http://www.celeryproject.org/`), is used as a task runner for the measurements. The geolocation databases are queried using the official Python packages where possible.

The software allows the user to populate the database with ground truth nodes. A script has been written to automatically add RIPE Atlas nodes using their API. Nodes can also be added manually. Furthermore, the software provides an interface for virtually any geolocation database. This interface has been implemented for the databases used in this research. After a task is started, the task runner uses this interface to query the different geolocation databases. Each individual measurement, which is a unique combination of a geolocation database and a ground truth node (with an IPv4 and an IPv6 address), is stored in the database. When the geolocation database queries are completed the results are displayed in a table and graphs are generated. These results, along with the available databases and the ground truth nodes, are accessible through the web interface.

The source code can be found on GitHub (`https://github.com/jjkester/geoip`). A running version of the software, including the results of the measurements, can be found at `https://geoip.jjkester.nl`. This website will be kept online for a limited amount of time.

3. IPV4 VERSUS IPV6

The designs of the IPv6 protocol and address spaces are different from their IPv4 counterparts (Q2.). This design difference can influence the way IP geolocation databases work with addresses. These differences might explain differences in the results of the conducted measurements in this study.

From a geolocation standpoint the way the address space is divided and allocated to customers is of most interest since databases contain a mapping between address blocks and their physical location. Mapping individual addresses to locations is possible but unrealistic with IPv4 and more so with IPv6 due to the vast amount of IPv6 addresses. This means that IP geolocation databases will have to fall back to mapping subnets or other blocks (e.g. a sequential set of 3 IP addresses) of IP addresses to locations. Since the allocation and assignment of IP addresses directly influences the location of the systems responding to these addresses, the way RIRs and system administrators allocate and assign addresses also influences the design and performance of IP geolocation databases.

Table 1. Different IP geolocation databases

	IPv6	Format	Free
DB-IP	yes	csv	no
IP2Location	yes	binary	yes
IPAddressLabs	no	http	no
IPinfo.io	yes	http	no
IPligence	no	http	no
MaxMind	yes	binary	yes

To explain this in detail I will discuss the IPv6 address allocation policy of RIPE, the RIR for Europe. For organizations, RIPE specifies a minimum allocation size of /32, which contains the same amount of addresses as the whole IPv4 address space [15]. This means that the number of rows in IPv6 geolocation databases will grow to be greater than their IPv4 counterpart, simply because there are more blocks to divide. Furthermore, organizations can divide their subnets and allocate smaller subnets to their customers.

The Dutch ISP XS4All, which is an early adopter of IPv6, providers every customer with an /48 block. This is in accordance with the IETF recommendations [13]. The IETF states that smaller blocks might also be acceptable. In a worst case scenario this would mean that a fictional IP geolocation database would map every /48 block to a location, which would make the database too large to be functional and to maintain properly. This means that IP geolocation database vendors need to optimize their databases which, when done inproperly, could mean less accurate results.

Due to the limited availability of IPv4 addresses it has become common to sell and reassign IPv4 address blocks. This means that the physical location of the systems responding to these addresses moves as well. IP geolocation databases have to account for this phenomenon by updating the records in the database. This is not (yet) an issue for IPv6 since there are plenty of IPv6 addresses available.

4. DATABASES AND GROUND TRUTH

The first part to measuring geolocation database accuracy is gathering IP geolocation databases to measure. Secondly, ground truth data needs to be collected so the accuracy of the database query results can be calculated.

4.1 Geolocation databases

A list of well known geolocation databases was compiled. This list was filtered for databases with IPv6 support, good availability (for example no rate limiting) and low cost. The availability and cost was important due to the limited time and resources available for this research. Table 1 shows the databases and their features.

The companies behind the non-free databases on the list were contacted to see if free access for research purposes was available. Both DB-IP and IPinfo.io replied with a positive answer. IPinfo.io was not used because of their statement that they use MaxMind's database [9], which

Table 2. Completeness of database query results.

	Complete	IPv4 only	IPv6 only	No result	Incomplete
DB-IP	2981	225	0	0	7%
IP2Location	3206	0	0	0	0%
MaxMind	2951	245	5	5	8%

Table 3. Average deviation of database query results. GT = ground truth.

	IPv4 vs. GT	IPv6 vs. GT	IPv4 vs. IPv6
DB-IP	1249.4 km	1508.2 km	1999.7 km
IP2Location	238.6 km	1892.4 km	2950.3 km
MaxMind	264.4 km	1098.7 km	1038.8 km

was already included in the research. This brings the final database list to DB-IP (IP address to location), IP2Location (DB5.Lite) and MaxMind (GeoLite2).

4.2 Collected ground truth

The collected ground truth can be split into two sections, usable nodes and unusable nodes. The usable nodes are the entries with sufficient information for this research, the unusable nodes are missing essential parts an cannot be used for a complete comparison. These usable nodes are systems with both a public IPv4 and IPv6 address and a known geographical location. Of the unusable nodes, 1,425 nodes did not have a known location, 5,805 nodes did only have an IPv4 address, 42 nodes did only have an IPv6 address and 6,664 nodes did not have any address. This last number can be explained by the fact that RIPE only publishes IP addresses which were recently used to connect to the RIPE servers. This could mean that a node is broken, no longer in use or very new (it has not been installed yet).

5. RESULTS

The measurement process resulted in a total of 9,618 measurements, each containing a database query result for both the IPv4 and IPv6 address of a ground truth node.

Table 2 shows the number of measurements per database and the number of incomplete measurements in case the database query did not provide a result.

In this table the row for IP2Location shows an interesting number. The strategy of this database provider seems to be to always return a result no matter the expected accuracy. This strategy could cause some results to be very poor. Whether this is true for this database will be discussed hereafter.

The other databases show that almost all incomplete measurements were due to the localization of the IPv6 address. Only MaxMind was not able to localize an IPv4 address in five cases.

For the remaining comparisons fully incomplete results are discarded. Where possible, partially incomplete results are used, for example when results for only IPv4 or IPv6 are calculated.

Table 3 contains average deviations between the location of the nodes according to the ground truth and their locations according to the databases. The first column contains the average distance between the ground truth locations and the IPv4 database locations, the second column contains the average distance between the ground truth locations and the IPv6 database locations and the last col-

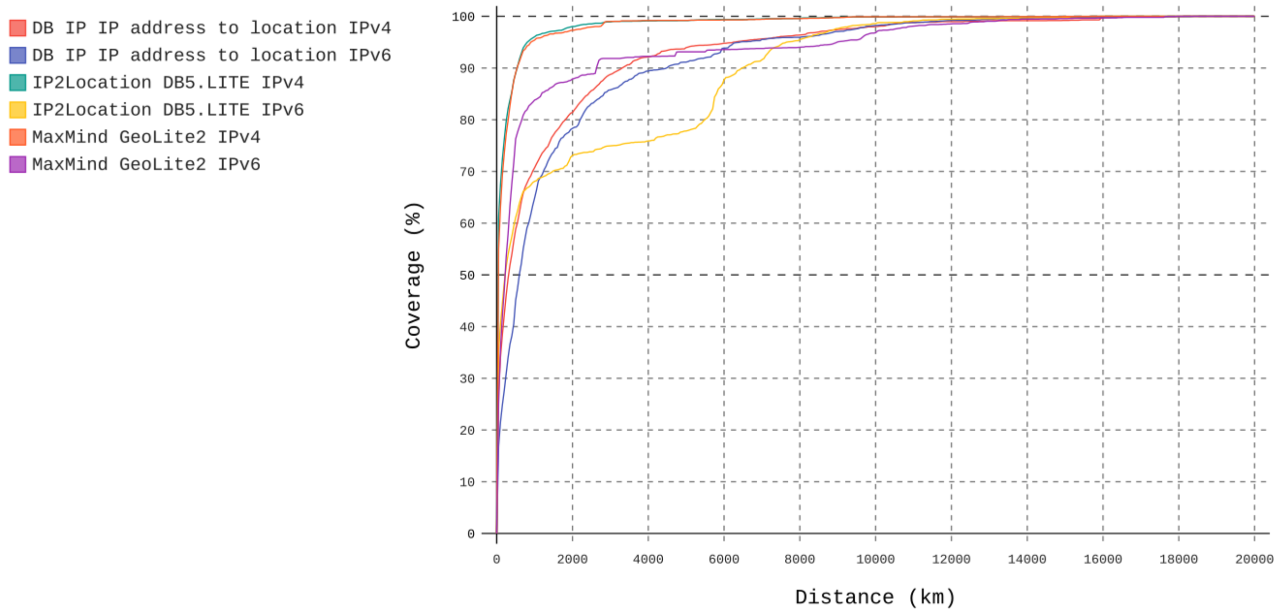

Figure 1. Geolocation database accuracy: percentage of results within a given range.

Table 4. Geolocation database accuracy: percentage of results within a given range

	< 1km	< 10km	< 25km	< 50km	< 100km	< 250km	< 500km	< 1000km
DB-IP (v4)	5%	16%	22%	27%	34%	46%	59%	71%
DB-IP (v6)	3%	9%	12%	16%	20%	29%	42%	60%
IP2Location (v4)	9%	43%	55%	61%	67%	80%	89%	96%
IP2Location (v6)	5%	24%	30%	35%	40%	52%	61%	68%
MaxMind (v4)	12%	38%	49%	55%	63%	77%	89%	95%
MaxMind (v6)	3%	15%	19%	23%	33%	49%	70%	77%

umn contains the average distance between the IPv4 and IPv6 database locations. This number can be larger than the IPv4 and IPv6 distances when the IPv4 and IPv6 addresses are located in opposite directions from the ground truth.

It is clear that DB-IP lags behind IP2Location and MaxMind when it comes to IPv4 geolocation. However, it is on par with the others when looking at the average deviation for IPv6 lookups. Furthermore, MaxMind performs on average the best of the three compared databases. It has by far the lowest deviation when it comes to IPv6 geolocation and it is almost as accurate as the best tested database for IPv4 geolocation. MaxMind also shows the lowest deviation between the IPv4 and IPv6 locations of a node. DB-IP and IP2Location perform respectively two and three times worse on this measurement.

The accuracy of each database is visualized in figure 1. The figure shows the percentage of the results, excluding the unknown locations, within the given distance from the respective ground truth locations. Certain distances are highlighted in table 4.

Figure 2 zooms in on the first 1,000 kilometers. This figure clearly shows that the average accuracy of IPv4 databases is better than their IPv6 counterparts. A general observation from this figure is that IPv4 geolocation performs a lot better than IPv6 geolocation. Also, the IPv4 performance of MaxMind and IP2Location is almost identical. This supports the data from table 3. DB-IP is the poorest performing database according to this figure. However, the IPv6 performance of DB-IP will eventually exceed the IPv6 performance of IP2Location according to table 3.

An analysis of the worst database results showed that in some cases distances of over 17,000 kilometers were found, both on IPv4 and IPv6. While these large distances on IPv6 were reported by all three databases, on IPv4 these numbers were mainly reported by DB-IP. This is shown in figure 3.

6. CONCLUSION

Before starting the research I expected that databases would be able to locate IPv6 addresses on a country level. Secondly, I expected the average IPv6 accuracy to be less than the average IPv4 accuracy.

Since the term 'country level' is not defined properly I assume 'country level' to be 'within a range of 500 kilometers'. This range is large enough to cover many European countries like The Netherlands, Estonia and Portugal and provides the user some idea to where a computer system is located in larger countries like the U.S. and China.

IP2Location and MaxMind can locate IPv6 addresses with

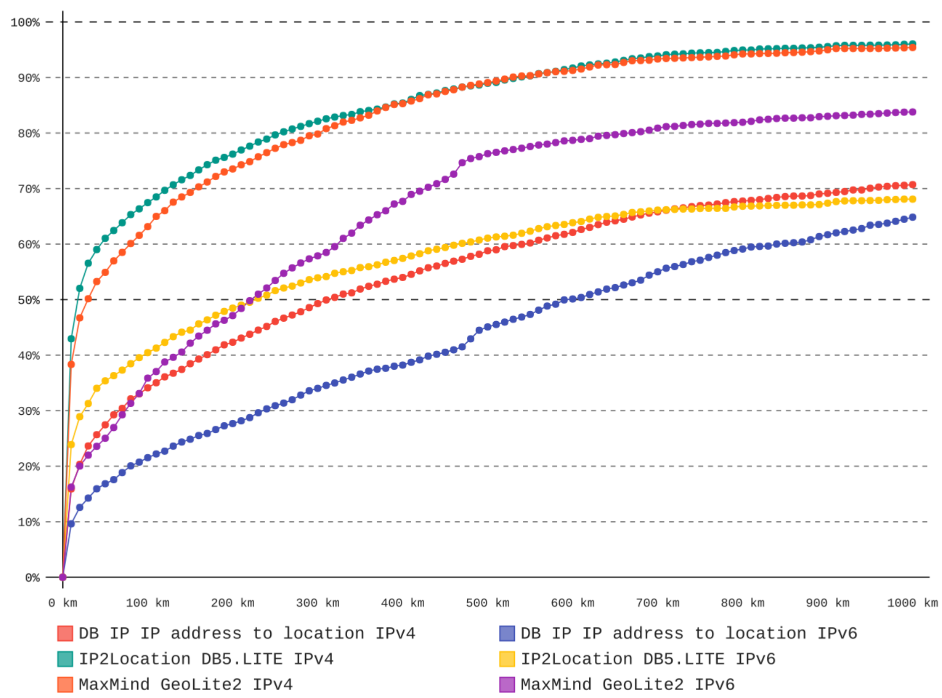

Figure 2. Geolocation database accuracy: percentage of results within a given range. Zoomed in on the most accurate results.

at least 60% accuracy within a range of 500 kilometers. DB-IP can do this with at most 40% accuracy. In conclusion, the IPv6 databases are not developed enough to properly locate IP addresses on a country level.

As expected, the average IPv6 accuracy is, for every tested database, less than the average IPv4 accuracy. This difference can be explained by the fact that IPv6 is not used enough to properly build a database for it. The research showed that IPv6 databases have the potential of being more accurate due to the structured way IPv6 addresses are assigned. This claim is also supported by the fact that IPv6 blocks are unlikely to get reassigned due to the amount of available address space.

Interestingly it turns out that all databases returned at least one result with a deviation from the ground truth data of about 15,000 kilometers and some databases return results with deviations over 18,000 kilometers, which is very close to the maximum possible deviation on earth. Therefore it is not possible to give any database a guaranteed maximum deviation, there always exists the possibility, although it is small, that the actual computer system is located at the other side of the world.

It is worth mentioning that on average MaxMind can geolocate IPv6 addresses more precise than DB-IP can geolocate IPv4 addresses. Furthermore, his research shows that, given the used ground truth and database versions, MaxMind is the best performing IPv6 database and DB-IP is the worst performing IPv4 database.

In conclusion the accuracy of IPv6 geolocation databases is relatively poor in comparison to their IPv4 counterparts. However, there is a lot of room for improving IPv6 geolocation. Based on the results of this research, MaxMind is the recommended database for general use, given that its IPv4 performance is on par with IP2Location, and its IPv6 performance is significantly better than the other databases.

7. FUTURE WORK

In studies on IP geolocation, and IP geolocation database performance in particular, the correctness of the ground truth has always been an issue. The ground truth in this study consists of location data from research network nodes. The locations of the RIPE Atlas nodes are manually entered by users. While this has the potential of being very accurate there is no real validation mechanism to check these locations. Also, the locations of the ground truth nodes can be biased to a country, continent, ISP, demographic, etc. This could mean that any conclusion is only valid for this bias (country, demographic, etc.) and might not hold in general. Therefore, future work should focus on improving the ground truth and/or proving the correctness of the ground truth data.

Furthermore, geolocation databases are updated regularly. Performing the same tasks a year from now with updated ground truth and databases may produce different results. A long running study into the development of the accuracy of the databases may show a correlation with the growing use of IPv6.

8. ACKNOWLEDGEMENTS

I would like to thank CAIDA for supplying me with the data of their research nodes and DB-IP for granting me access to their geolocation databases for the purpose of this research.

9. REFERENCES

[1] R. Beverly, A. Berger, N. Weaver, and L. Campbell. Inferring Internet server IPv4 and IPv6 address relationships. Technical Report, June 2013.

[2] Center for Applied Internet Data Analysis (CAIDA). CAIDA Monitors, the archipelago measurement infrastructure. http://www.caida.org/data/monitors/, Mar. 2015.

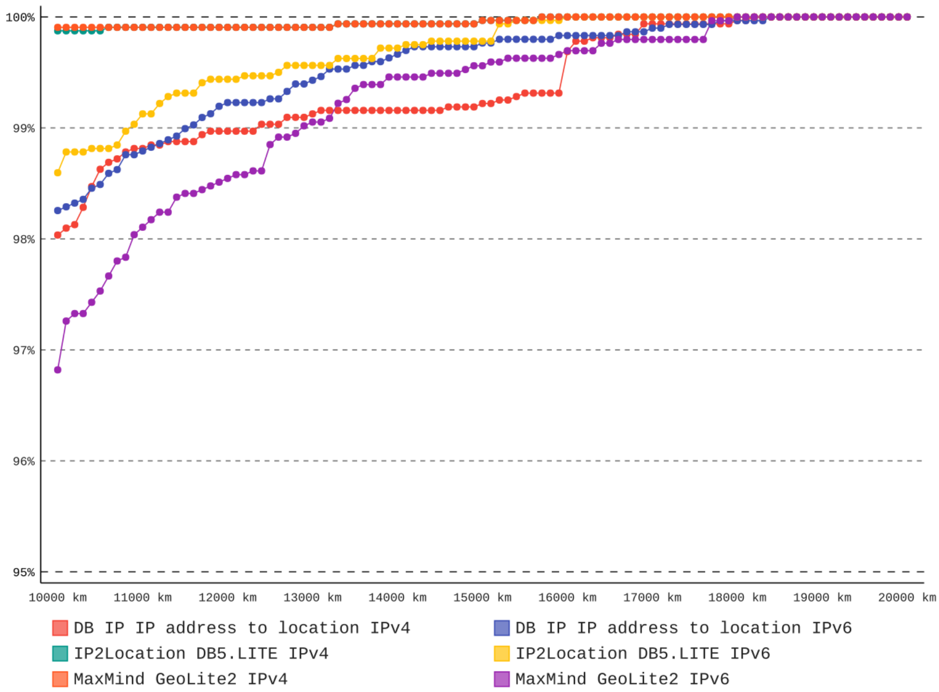

Figure 3. Geolocation database accuracy: percentage of results within a given range. Zoomed in on the least accurate results.

[3] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. `https://tools.ietf.org/html/rfc2460`, Dec. 1998.

[4] Z. Dong, R. D. W. Perera, R. Chandramouli, and K. P. Subbalakshmi. Network measurement based modeling and optimization for IP geolocation. *Computer Networks*, 56(1):85–98, Jan. 2012.

[5] B. Eriksson, P. Barford, B. Maggs, and R. Nowak. Posit: A Lightweight Approach for IP Geolocation. *SIGMETRICS Perform. Eval. Rev.*, 40(2):2–11, Oct. 2012.

[6] B. Eriksson, P. Barford, J. Sommers, and R. Nowak. A Learning-Based Approach for IP Geolocation. In A. Krishnamurthy and B. Plattner, editors, *Passive and Active Measurement*, number 6032 in Lecture Notes in Computer Science, pages 171–180. Springer Berlin Heidelberg, 2010.

[7] Google. Google IPv6 Adoption. `http://www.google.com/intl/nl/ipv6/statistics.html`, Oct. 2015.

[8] B. Huffaker, M. Fomenkov, and k. claffy. Geocompare, a comparison of public and commercial geolocation databases. Technical report, 2011.

[9] IPinfo.io. IP geolocation API documentation. `http://ipinfo.io/developers`, Oct. 2015.

[10] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe. Towards IP Geolocation Using Delay and Topology Measurements. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, IMC '06, pages 71–84, New York, NY, USA, 2006. ACM.

[11] S. Laki, P. Matray, P. Haga, T. Sebok, I. Csabai, and G. Vattay. Spotter: A model based active geolocation service. In *2011 Proceedings IEEE INFOCOM*, pages 3173–3181, Apr. 2011.

[12] H. Maziku, S. Shetty, K. J. Han, and T. Rogers. Enhancing the Classification Accuracy of IP Geolocation. Technical report, Oct. 2013.

[13] T. Narten, G. Huston, and L. Roberts. IPv6 address assignment to end sites. `https://tools.ietf.org/html/rfc6177`, Mar. 2011.

[14] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. IP Geolocation Databases: Unreliable? *SIGCOMM Comput. Commun. Rev.*, 41(2):53–56, Apr. 2011.

[15] RIPE NCC. IPv6 address allocation and assignment policy. `https://www.ripe.net/publications/docs/ripe-655`, Oct. 2015.

[16] RIPE NCC. RIPE Atlas Probes. `https://atlas.ripe.net/probes/`, Oct. 2015.

[17] Y. Shavitt and N. Zilberman. Improving IP geolocation by crawling the internet PoP level graph. In *IFIP Networking Conference, 2013*, pages 1–9, May 2013.

[18] T. V. H. Tran. *IPv6 geolocation using latency constraints*. Thesis, Monterey, California: Naval Postgraduate School, Mar. 2014.

[19] Wikipedia. World geodetic system. `https://en.wikipedia.org/wiki/World_Geodetic_System`, June 2015.

[20] S. Zander. On the Accuracy of IP Geolocation Based on IP Allocation Data. Technical report, 2012.