Triangulation Using RTT

A Constraint-Based Geolocation Redux

Samir Menon
Stanford University

ABSTRACT

We replicated the result from Gueye et al., "Constraint-Based Geolocation of Internet Hosts", SIGCOMM IMC 2004 using modern data. We built our replication entirely from the source material and new data; we never had access to the original data or code. Our implementation and reproduction instructions are available at https://github.com/menonsamir/CBG-replication.

The central finding of the paper is confirmed: a series of RTT measurements to hosts with known locations can produce a relatively accurate estimate of geolocation. In our tests in the United States and Western Europe, we achieve median errors of less than 70km.

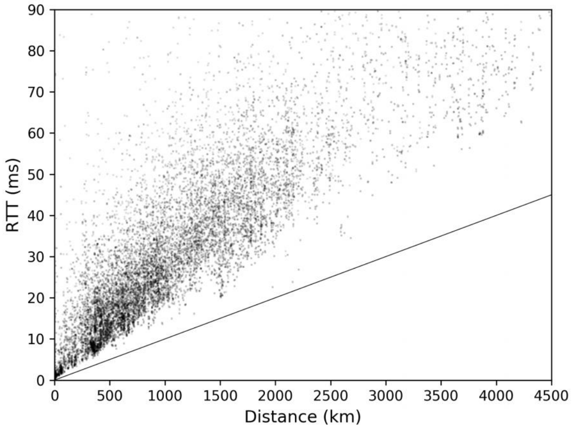

Figure 1: Scatter plot of geographic distance and network delay.

1 CURRENT LANDSCAPE & BACKGROUND

A variety of techniques exist to geographically locate an internet host. Today, the most common one is IP geolocation, a technique that uses large volumes of traffic data to map blocks of the IP address space to the geographic areas where hosts are located. As has been studied, these databases require large amounts of data over long periods of time to stay useful [3][5]. However, as the internet has grown, the rise of applications used by large amounts of people has made the collection of this information relatively simple, and so IP geolocation remains the standard technique, and it has proven useful at the scale of millions of hosts [2].

The Gueye et al. paper, "Constraint-Based Geolocation of Internet Hosts", from SIGCOMM IMC 2004 (CBG) [1], was written in a different context. IP geolocation as a concept was still in its infancy; the first large-scale attempt hand only just begun in 2002 by a then unknown company named MaxMind, and the technique really only became effective many years later.

Researchers seemed more interested in what they termed 'measurement-based' approaches. These involve measuring some live values in the network (usually delay or throughput) and then inferring location from these. The most naive approach simply had a target ping several landmark hosts with know locations, and then estimated the target's location as the position of the landmark host with the shortest delay [3].

2 CONTEXT OF THE ORIGINAL PAPER

CBG was an improvement on this approach, because it gave a 'continuous' location rather than just finding the nearest landmark. Inherently, CBG's technique assumes that the network is dense enough to be roughly approximated by 'as-the-crow-flies' communication channels between hosts that operate at a delay tightly related to distance traveled.

The paper was valuable because it both provided a technique for solving an important problem (locating internet hosts) and gave insight into the structure and function of the internet (the relationship between geographic location and delay).

Today, there are many non-measurement-based techniques for locating an internet host which are highly accurate (and perhaps more accurate than any measurement-based system can be) [5]. Despite this, it's useful to examine the extent to which the assumptions that CBG makes about the structure of the internet are true today, and the most salient way to do that is to try to replicate the results on new data.

3 REPLICATION DESIGN

We aimed to replicate Figure 1 and Figure 3 of the CBG paper using more current data.

We used data from the RIPE Atlas of 36 hosts across the United States, and 54 hosts across Western Europe [4]. This is a small subset of 333 worldwide 'anchor' hosts that all ping each other every 2 minutes and record that data for research purposes. We chose these subsets by taking every RIPE location in the specified region that had RTT ping data available for the date and time we tested. Notably, the RIPE Atlas project is the spiritual successor of the now decommissioned RIPE Test Traffic Measurement (TTM) project that was used by the original paper.

We had to make several small additions to the basic algorithm outlined in the paper to make it flexible enough to handle real-world data. Most of the changes were minor, and involved gracefully handling edge cases where no ping data was available for some reason. However, one case warrants explanation: the paper never specifies what the algorithm should do when no polygon of circle intersection points exists. That is, a series of landmarks and distance estimates can be thought as a set of circles on Cartesian plane that denote possible locations of the target. Then if any one landmark underestimates the distance to the target, then there will be no polygon that represents an area that all landmarks agree the target will be in. While the bestline approach attempts to solve this by always erring on the side of overestimating distance, it is not perfect, since it relies on a limited set of collected data that cannot represent all possible RTTs.

We encountered this rarely but persistently in our analysis. Across timespans and locations, a small proportion (roughly 1%) of our estimation routine calls would fail, because there was no polygon. The important thing to note here is that it is incorrect to simply throw out these cases; the algorithm cannot pick and choose what cases it has to predict. The most basic thing to do would be to just make an arbitrary guess (that is, in the center of the region tested) and continue. However, we developed a simple technique that does better.

Instead of making an arbitrary guess, we repeat the circle intersection procedure, but with larger radii at all landmarks. By increasing the radii, we increase the chances that a previously underestimating landmark will now overestimate, and we will be able to construct a polygon. We repeat this procedure with successively larger multiplicative factors (1.05, 1.1, 1.25, 1.5, 2.0, ...) until we get an intersection. In our tests, the largest factor ever used was 1.5, suggesting that the underestimate is usually not grossly less than the actual distance.

We also took the minimum RTT over a very short amount of time (4 minutes) rather than the 2.5th percentile of RTT over many weeks. While in theory, the latter approach should be much better, we achieved reasonable results without needing to collect more data. We discuss this more in section 6.

4 RESULTS

Just as the original paper did, we first tried to verify the basis of constraint-based geolocation. To summarize, we know that every delay measurement between two points can be thought of as having two components: the propagation delay, which is usually on the order of the speed of light times the path length, and the network-based delay, a broad category of possible queueing, holding, processing and other delays added by intermediaries on the path.

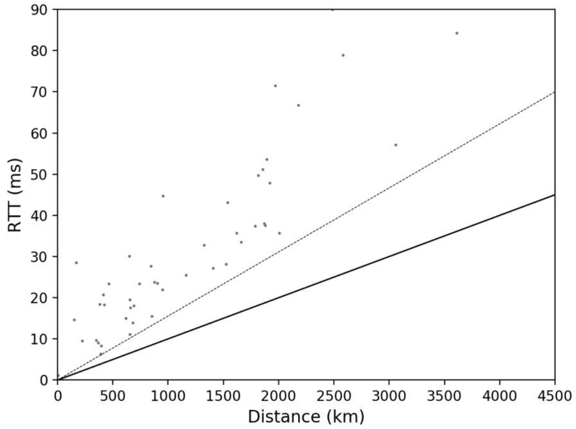

Figure 2: The observed measurements from a sample host, and its 'bestline' and 'baseline'.

A graph of all RTTs and observed distances is in Fig. 1. You can see that there is clearly some relationship, with the vast majority of measurements happening in a relatively small band of RTT values.

The key observation of CBG is that this second class of delay is only additive; it is generally not multiplicative nor subtractive. Then we can think of the minimum RTT over some period of time as approaching an estimate of the best (ie least distorted) path between two hosts. Then taking many of these minimum RTTs to locations various distances away, we can determine a relationship between distance away and minimum RTT for each source landmark. Then, using estimated distances (from measured RTTs) from many known locations, we can perform standard triangulation and guess a target's location.

The specific relationship that we determine is the 'bestline' - the line with positive y-intercept that best estimates all the (distance, RTT) points without overestimating any RTT. This is so that given some RTT, our distance estimate is always an overestimate, so that when we intersect all of the circles, we will get a proper region of intersection rather than a series of disjoint regions. An example of what such a calculation looks like is in Fig. 2.

Finally, we test the central claim: can we use CBG to accurately predict the locations of target hosts. Fig. 3 shows the cumulative probability density of the error distance of both our replication of CBG and an ablated version of the model described below. We observe mean errors in the U.S. of 178km and in Western Europe of 174km. We also observe median errors in the U.S. of 67km and in Western Europe of 69km. The ablated version has much higher median errors of 384km and 131km in the U.S. and Western Europe, respectively.

5 DISCUSSION

Our implementation appears to replicate the results of the original paper accurately. We appear to achieve similar results in the United States and Western Europe, though a precise comparison to the original is difficult because only the graphs of the original paper are available, and the data source (RIPE TTM) is no longer available. Our median error is higher than the original for Western Europe, but lower for the United States.

As an ablation test, we also attempted to run the model with a baseline that is based just on the speed of light in copper (which we have assumed is 100 km / ms, or very close to 2/3 the speed of light in a vacuum). This is to test the paper's claim that the 'bestline' approach better captures the dynamics of RTT vs. distance. With this restriction, the triangulation scheme does much worse, showing that the prior information about RTT vs. distance captured by the bestlines is actually useful. Moreover, we can see that just triangulating speed-of-light information from many geographically diverse hosts is not enough to determine location; we appear to need individual information from each host about how distance from that host and RTT are related.

We also confirm a smaller intuition of the authors. The authors note that in their analysis, the gap between CBG and the naive approach is much larger for Western Europe than the U.S. They explain this as a result of the fact that the Western Europe test had fewer landmarks, and so therefore the naive approach was much worse. Intuitively, the CBG model outperforms baseline approaches the most when there is more limited information. In our results, we note that the trend is flipped - the gap is much bigger in the US. However, we actually used only 36 hosts in the US, but 54 in Western Europe, making our result match the intuition of the authors; since the US had fewer landmarks and therefore more limited information upon which to base an estimate, CBG outperformed the naive approach by a larger margin.

Firstly, this data suggests that the internet regularly transfers information at a relatively large fraction of the universal limit. This is a reminder that, at least from a latency perspective, the internet pretty surprisingly close to optimal.

The fact that we are able to replicate results from 2004 on an entirely different set of hosts from 2017 suggests that the CBG technique is robust. We suspect this is due to the simplicity of the technique and the small number of assumptions that it makes about the structure of the internet.

More broadly, this indicates a key fact about the internet that was true in 2004 and is still true today: at a birds-eye view, the distance and the speed at which internet moves most information between two points are tightly related. In fact, the minor variation in the relationship can be modeled extremely simply (using a line) and then combined to produce more accurate estimates.

The fact that we took minimum RTT over only 4 minutes (rather than the 2.5th percentile over weeks) and still got comparable results might indicate that the amount of data the previous authors collected was not necessary. It also means that the system can be more readily deployed and 'live' - we only need to collect data for 4 minutes, and then we can locate hosts for at least that period of time. In other words, a large-scale implementation of this procedure would just rerun the analysis every 4 minutes to keep the distance-to-RTT information highly accurate.

6 LIMITATIONS & WEAKNESSES

As a comparison, our ablation test is not quite the same as the original paper's point-of-comparison, GeoPing. We chose to replace GeoPing for two reasons: firstly, re-implementing it seemed to challenging, since the paper that introduced it lacked details on how it worked and no existing implementation was available; secondly, we believe our ablation test better captures the idea of testing the original authors intuition about bestlines. In other words, while the original authors may have wanted to show that their system was better than the state-of-the-art at the time, we are less interested in that goal, and would instead like to show that the delay information captured in the 'learning' phase is actually put to use.

Most immediately, the map projection actually introduces non-trivial error in our solution. Briefly, a map projection is what allows the spherical (or more accurately, ellipsoidal) surface of the Earth to be represented as an x-y Cartesian plane. While the original paper doesn't actually specify a projection, all of its equations and description imply a 2D representation. It is certainly possible to instead perform the above using intersecting spheres and ellipsoids, but it is significantly more challenging to implement. All map projections incur some degree of error, whether in direction or distance. We chose the Web Mercator projection, a popular standardized projection used by most commercial and open-source modern online mapping systems. While this is a popular projection and there are many existing implementations, it tends to incur a large amount of error in distance measurement. During testing, we observed an up to 30km

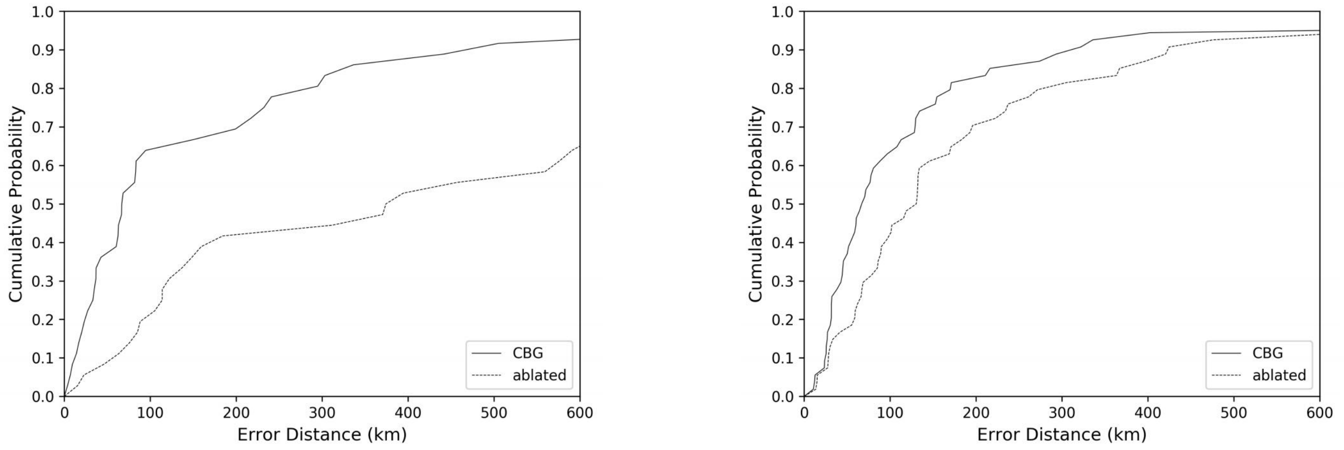

(a) U.S. Dataset

(b) Western Europe Dataset

Figure 3: Error distance for CBG and the ablation test.

error across the US dataset purely due to the projection. It may pay dividends in terms of accuracy to choose a different projection; additionally, for comparability, it would be nice to know what projection the original authors used.

While we carried out a broad range of analysis, we only sanity checked our 4 minute window - that is, to achieve more robust results, we would need to run this over all 4 minute windows of some number of weeks or months to be absolutely sure that the technique can work at scale.

7 FUTURE WORK

One weakness of our analysis, and the original CBG paper, is the fact that we only ever try to locate the anchors themselves. That is, if the anchors tend to have some kind of pattern in how they are situated on the network (for example, if they all happen to sit on large trunk lines). The next step is to see if this technique can work on hosts that are not RIPE anchors; that is, if it can work for normal end-user hosts. This would require a access to a geographically diverse of set of hosts that issue pings to the RIPE Atlas anchor hosts regularly. While such an experiment is possible within the Atlas system, we'd need more time and could not just conduct the study on historical data.

Another possible avenue for exploration would be target type; that is, could a version of this technique work on devices that connect over something like LTE? We suspect that a naive implementation of CBG will probably not work, since latency is far more variable and less related to distance. Still, a proper evaluation might yield insight into how the approximations used in CBG may or may not apply more generally.

Finally, one way to further measurement-based techniques for geolocation is to expand the notion of measurement. Currently, most techniques just use some ensemble of latency information; the next step would be to analyze other information, like throughput behavior, hop count, path, and more. One can imagine a system that takes a large ensemble of this data and uses standard machine learning techniques to predict location. On the other hand, the beauty of simpler models like the one presented in this paper is that they are explainable from first principles.

8 CONCLUSION

We have successfully replicated the result from Gueye et al. using data from 2017. Our replication shows that the internet can be reasonably approximated as relating distance and RTT linearly. We achieved median geolocation errors in the United States and Western Europe of less than 70km, and showed through ablation that the 'bestline' technique is, in fact, essential to the success of CBG.

Anyone who would like to replicate our results can do so easily. Our code and instructions for running it are available at https://github.com/menonsamir/CBG-replication.

REFERENCES

[1] Bamba Gueye, Artur Ziviani, Mark Crovella, and Serge Fdida. 2006. Constraint-based geolocation of internet hosts. *IEEE/ACM Transactions On Networking* 14, 6 (2006), 1219–1232.

[2] Zi Hu, John Heidemann, and Yuri Pradkin. 2012. Towards geolocation of millions of IP addresses. In *Proceedings of the 2012 Internet Measurement Conference*. ACM, 123–130.

[3] Venkata N Padmanabhan and Lakshminarayanan Subramanian. 2001. An investigation of geographic mapping techniques for Internet hosts. In *ACM SIGCOMM Computer Communication Review*, Vol. 31. ACM,

173–185.
[4] NCC Ripe. 2010. RIPE atlas. (2010).
[5] Yuval Shavitt and Noa Zilberman. 2011. A geolocation databases study. *IEEE Journal on Selected Areas in Communications* 29, 10 (2011), 2044–2056.