

AIG Paper Changes

Review/CITE papers:

B. Gueye, A. Ziviani, S. Fdida, J. F. de Rezende, O. C. M. B. Duarte. « Two-Tier Geographic Location of Internet Hosts ». In Proc. 7th IEEE International Conference on High Speed Networks and Multimedia Communications (HSNMC). July 2004

Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, C. Huang. « Towards Street-Level Client-Independent IP Geolocation ». In Proc. USENIX NSDI 2011

Ingmar Poesse, Steve Uhlig, Mohamed Ali Kâafar, Benoit Donnet, Bamba Gueye: IP geolocation databases: unreliable? Computer Communication Review 41(2): 53-56, 2011.

S. Siwipersad, B. Gueye, and S. Uhlig. Assessing the geographic resolution of exhaustive tabulation for geolocating Internet hosts. Proc. of Passive and Active Measurement conference (PAM), Cleveland, Ohio, USA, April 2008.

B. Gueye, S. Uhlig, and S. Fdida. Investigating the Imprecision of IP Block-Based Geolocation. Proc. of Passive and Active Measurement conference (PAM), Louvain-la-neuve, Belgium, April 2007.

Bamba Gueye, Steve Uhlig, Artur Ziviani, Serge Fdida: Leveraging Buffering Delay Estimation for Geolocation of Internet Hosts. Proc. of IFIP Networking, 2006.

Provide Additional Details/Clarifications:

~~The introduction lacks of references. For instance, the paragraph starting with « One possible approach towards IP geolocation... » Or the one starting with « Database methods are often exclusive and updated manually ». Otherwise, those paragraphs sound like free statements. → Valid point, I need to find references done~~

- I found the way authors « divide » the world in regions pretty strange. Why limiting a region to Pakistan? Why not considering Middle East and Extreme East (for instance)? Anyway, this 2-tiers architecture looks like being heavy to deploy. -> **I think it is obvious that we divide regions based on available landmarks (need to think about it).**

~~Based on the color code, it seems to me that, on north America (see Fig. 6), only 4 nodes are active for ping. And only 5 for Europe (Fig. 8) → I should improve the figures by using more landmarks using all landmarks for alpha analysis now.~~

~~Looking @ Fig. 7 and 9, it seems that North America and Europe have been less extensively measured than Pakistan. To me, this may introduce some bias in the measures and, by extension, in the conclusions drawn by authors. Please discuss that point in the paper (BTW, why so much landmarks in Pakistan and so little in Europe and North America?). More generally, what is the exact impact of the number of Pinger nodes on the alpha parameter evaluation? → We have much more PerfSONAR landmarks now and we should add that.~~

~~From the paper, it is not clear how AIG builds its constraints per landmark for each measurement of RTT (one for the upper bound, one for the lower bound). In addition, why considering 7 alpha values for the upper bound and only one for the lower bound? What is the exact impact of the number of alpha considered on the measurement results? → The latest AIG used in TULIP now has 3 lower alphas. I will mention this and also add suitable figures to illustrate.~~

~~Why the minimum alpha value is set to 0.15 for North America. That's strongly different from the minimum value observed (see Table I). → We no longer use 0.15. I will update this~~

~~When performing your measurement, did you use all landmarks (in each region) for CBG? Or do you restrict yourselves to a subset (following Gueye et al. suggestions on the number of landmarks to use — see ToN 2006) → Need to look into this~~

~~Why results are (nearly) similar in Pakistan? And so different in other regions? Is it because Pakistan only consider PingEr landmarks? Please, clarify. → Too high landmark density~~

~~Does CBG appear on Fig. 22? If so, please, clarify the plot legend. -> There is no CBG here just AIG~~

~~Comparing with databases is a nice idea. But, please, consider well known databases, like IP2Geolocation, Maxmind, ... Further, it seems that your results with databases are pretty good. This is in contradiction with what has already been observed (see, for instance, Poesse et al. CCR 2011). Is it because landmarks are mostly PlanetLab nodes, thus within universities. And universities are pretty well localized by databases? → The results are either very good or very bad just as one would expect~~