Poster: Verifying Geographic Location Presence of Internet Clients

AbdelRahman Abdou
School of Computer Science
Carleton University
Ottawa, ON, Canada

Ashraf Matrawy
School of Information Technology
Carleton University
Ottawa, ON, Canada

Paul C. van Oorschot
School of Computer Science
Carleton University
Ottawa, ON, Canada

***Abstract*—Verifying the geographic locations of Internet clients is a challenging problem. It has various security applications such as fraud protection, location-based access control and location-based authentication. Geolocation techniques that are based on the IP address do not go beyond the network-layer of the TCP/IP protocol stack, thus fail whenever the client employs IP-hiding tactics like proxies and VPNs. This poster explains Client Presence Verification (CPV), an approach designed to verify the geographic locations of clients on the web. CPV relies on network delay measurements to verify locations, and is designed to mitigate previously studied security flaws of measurement-based geolocation systems. Real-world evaluation of CPV's false reject/accept rates shows its promising potential for practical adoption.**

I. Background

Over the Internet, location-sensitive service/content providers are those that customize their services based on the geographic locations of their web-clients. This is commonly done for contractual restrictions related to licensing or broadcasting content. Many Audio/Video-on-Demand (AVOD) and media content providers, like Hulu and Netflix, employ geographic-restriction policies, *e.g.,* to adhere to license agreements. A *geographic-restriction policy* is a policy that dictates the type of services provided to clients at different geographic regions. Furthermore, verification of clients' locations could be used to reduce credit card fraud, or as an additional authentication factor.

One class of approaches to geolocating (*i.e.,* determining the geographic location of) web-clients is inference based on unverified assertion obtained directly from the client such as GPS coordinates [3], or indirectly such as the client's IP address [4] or domain name location hints [5]. WiFi Positioning System (WPS) are also commonly used by browsers through the W3C Geolocation API [6]; WPS is based on multilateration of surrounding WiFi access points with known geographic locations, where the distance between the client and each access point is estimated based on the signal strength.

Another approach involves delay measurements from a set of machines with known locations, and using multilateration to determine the client's location [7]. This class relies fundamentally on the strong correlation between Internet delays and geographic distances [8], and has proven high location accuracy in recent years [9].

In the presence of an adversary interested to forge its geographic location, *e.g.,* to gain location-dependent benefits or avoid accountability, Internet geolocation methods are subject to evasion [5], [10]. These methods are typically designed for non-adversarial environments, and therefore the locations they return are best treated (from a security perspective) as unverified assertions which require verification of some form to provide stronger assurance or confidence in adversarial presence. To that end, a location verification technique is required to provide assurance of the results returned by an Internet geolocation technique.

II. Delay-based Location Verification

We explain Client Presence Verification (CPV), a measurement-based technique which, in contrast to previous geolocation techniques, is specifically designed for adversarial environments. In CPV, network delays between a client and a set of three verifiers with known locations are used to verify the client's geographic location. The delays are measured over the application layer (*e.g.,* using Websockets [11]) to mitigate standard IP-masking tactics like proxy servers and Virtual Private Networks (VPNs).

Findings in previous literature establish that Internet delays and geographic distances are strongly correlated [8], [9]; we leverage that, and propose heuristics to reduce the effect of delay irregularities, such as delay spikes, on the verification result. These heuristics are as follows.

First, CPV relies on one-way delay (OWD) estimates rather than round-trip times (RTTs). RTTs are commonly used in delay-based geolocation mainly because they are easier to measure, requiring less cooperation from the client than that in estimating OWDs. The standard protocol for OWD measurement [12] requires the two parties to synchronize their clocks, and exchange timestamp messages. However, CPV does not rely on this kind of client cooperation because the client is assumed the ability to manipulate any parameters under its control in order to forge its true geographic location. As such, a OWD-estimation technique was devised that requires substantially less client cooperation [13].

This poster discusses results in a recently accepted journal paper [1], extending work on Internet location-verification principles introduced earlier in a conference version [2]. It is part of the poster session of the 2016 IEEE Symposium on Security and Privacy.

Fig. 1. Example of a client asserting to be in Albuquerque, NM; the three chosen verifiers are at Tucson, Arizona; Oklahoma city, Oklahoma; and West Valley city, Utah. Map data: Google, INEGI.

Second, in CPV delays are measured actively between the verifiers and with the client to reflect the most recent network status [14]. Accordingly, network congestion, altered routes, or other factors that affect the estimated delays are spontaneously taken into consideration in the verification decision.

Third, CPV does not use a universal delay-to-distance mapping function. Instead, delays between the client and the verifiers are compared to those between the verifiers themselves, and the client's location is verified accordingly.

Fourth, CPV estimates delays iteratively to better exclude delay outliers, and reflect current network conditions.

Finally the three verifiers are chosen, as part of the CPV mechanism from a pool of available verifiers, to be in geographic proximity and encompassing the client's asserted location. This is useful in reducing the number of spanned Autonomous Systems along the route, reducing route circuitousness, and thus leading to stronger positive correlation between delays and distances [15].

How CPV works

First, the client's geographic location is asserted using any standard geolocation technique, such as the client's submitted GPS coordinates. The CPV algorithm then chooses three verifiers encompassing the client's asserted location, and sends their IP addresses to the client (*e.g.,* web-browser). The client connects to these verifiers, and the location verification processes commences. The verifiers' objective is to provide greater assurance regarding the asserted location by checking if the client is truly present inside the triangle geographically determined by their locations (see Fig. 1). The area of that triangle is thus the verification granularity; as the triangle gets larger, granularity becomes coarser.

The verifiers iteratively estimate the smaller of the forward and reverse OWDs between themselves and the client. After n iterations of delay estimation ($10 < n \leq 100$), the delays are fed to an algorithm that geometrically checks if the point representing the client's asserted geographic location is consistent with a point that would fall within the triangle determined by the three verifiers in a 2-D Cartesian plane. If that is true for $n\tau$ of the n iterations for some tunable threshold τ ($0 < \tau \leq 1$), the asserted location is considered positively verified.

III. Evaluation

We evaluated CPV using PlanetLab [16], conducting near 2,500 experiments using 80 PlanetLab nodes in the US and Canada. We used the nodes to represent verifiers, legitimate clients (truly inside triangles), and adversaries (outside triangles). Knowing the ground truth of legitimate clients and adversaries, we quantified CPV's false reject (FR) and false accept (FA) rates at various values of n and τ. Assuming appropriate triangle choices, *i.e.,* legitimate clients not being on the triangle's borders or too close thereof, FR and FA rates were 2% and 1% respectively. Our analysis also shows that any $\tau > 0$ was sufficient to reject 90% of all adversaries.

References

[1] A. M. Abdou, A. Matrawy, and P. C. van Oorschot, "CPV: Delay-based Location Verification for the Internet," *IEEE Trans. Dependable and Secure Computing, TDSC (to appear; accepted June 14)*, 2015.

[2] A. Abdou, A. Matrawy, and P. C. van Oorschot, "Location Verification on the Internet: Towards Enforcing Location-aware Access Policies Over Internet Clients," in *CNS*. IEEE, 2014.

[3] D. Hu and C.-L. Wang, "GPS-Based Location Extraction and Presence Management for Mobile Instant Messenger," *LNCS Embedded and Ubiquitous Computing*, vol. 4808, 2007.

[4] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye, "IP geolocation databases: unreliable?" *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 2, 2011.

[5] J. A. Muir and P. C. van Oorschot, "Internet geolocation: Evasion and counterevasion," *ACM Comput. Surv.*, vol. 42, no. 1, 2009.

[6] A. Popescu, "Geolocation API Specification," http://www.w3.org/TR/geolocation-API/, Oct 2013.

[7] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, "Constraint-based geolocation of Internet hosts," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, 2006.

[8] Z. Dong, R. D. Perera, R. Chandramouli, and K. Subbalakshmi, "Network measurement based modeling and optimization for IP geolocation," *Elsevier Computer Networks*, vol. 56, no. 1, 2012.

[9] Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang, "Towards street-level client-independent IP geolocation," in *NSDI*. USENIX, 2011.

[10] P. Gill, Y. Ganjali, B. Wong, and D. Lie, "Dude, where's that IP? Circumventing measurement-based IP geolocation," in *USENIX Security*. USENIX, 2010.

[11] I. Fette and A. Melnikov, "The WebSocket Protocol," RFC 6455 (Proposed Standard), IETF, Dec. 2011.

[12] S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)," RFC 4656 (Proposed Standard), IETF, Sep. 2006.

[13] A. M. Abdou, A. Matrawy, and P. C. van Oorschot, "Accurate One-Way Delay Estimation with Reduced Client-Trustworthiness," *IEEE Commun. Lett.*, vol. 19, no. 5, pp. 735–738, 2015.

[14] Y. Zhang and N. Duffield, "On the Constancy of Internet Path Properties," in *the 1st ACM SIGCOMM Workshop*, ser. IMW. ACM, 2001.

[15] R. Landa, R. G. Clegg, J. T. Araújo, E. Mykoniati, D. Griffin, and M. Rio, "Measuring the Relationships between Internet Geography and RTT," in *ICCCN*. IEEE, 2013.

[16] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "PlanetLab: An Overlay Testbed for Broad-coverage Services," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 3, 2003.