across regions. For instance, our results show that well connected regions such as Europe expose fewer caching headers than regions such as Africa, where transit is costly. While revealing, it should also be noted that exposing this information is a potential security threat, due to the frequent use of old and vulnerable middlebox software [7].

Our future work will focus on exploring how these trends evolve. Interestingly, many HTTP/2.0 browser implementations are following an encrypt everything model, which will undermine some middlebox functions. This is perhaps concerning as our work indicates a widespread dependence on their functionality. Hence, ISPs may endeavour to find ways around this [27]. This is particularly the case for security and performance oriented middleboxes, which may be considered critical to business operations. Hence, we believe that the continued monitoring of this process could offer fascinating insight into how network operators react and optimise to changes in Web protocols.

9. REFERENCES

[1] The CAIDA UCSD AS classification dataset. `http://www.caida.org/data/as_classification`.

[2] Hola. `https://hola.org/`.

[3] Luminati FAQ. `https://luminati.io/faq`.

[4] Verizon injecting perma-cookies to track mobile customers, bypassing privacy controls. `https://www.eff.org/deeplinks/2014/11/verizon-x-uidh`.

[5] AT&T stops using invasive perma-cookies, but it may turn them back on. `http://www.wired.com/2014/11/att-hits-pause-privacy-busting-perma-cookie-test`, 2014.

[6] The relative cost of bandwidth around the world. `https://blog.cloudflare.com/the-relative-cost-of-bandwidth-around-the-world`, 2014.

[7] Squid: Security vulnerabilities. `https://www.cvedetails.com/vulnerability-list/vendor_id-823/Squid.html`, 2016.

[8] Carlucci, G., De Cicco, L., and Mascolo, S. Http over udp: an experimental investigation of quic. In *Proc. ACM SAC* (2015).

[9] Chung, T., Choffnes, D., and Mislove, A. Tunneling for transparency: A large-scale analysis of end-to-end violations in the internet. In *Proc. ACM IMC* (2016).

[10] Detal, G., Hesmans, B., Bonaventure, O., Vanaubel, Y., and Donnet, B. Revealing middlebox interference with tracebox. In *Proc. ACM IMC* (2013).

[11] Dhamdhere, A., and Dovrolis, C. Twelve years in the evolution of the internet ecosystem. *IEEE/ACM Transactions on Networking (ToN)* (2011).

[12] Dimitropoulos, X., Krioukov, D., Riley, G., et al. Revealing the autonomous system taxonomy: The machine learning approach. In *Proc. PAM* (2006).

[13] Elkhatib, Y., Tyson, G., and Welzl, M. Can spdy really make the web faster? In *IFIP Networking* (2014).

[14] Fanou, R., Francois, P., and Aben, E. On the diversity of interdomain routing in africa. In *Proc. PAM* (2015).

[15] Fanou, R., Tyson, G., Francois, P., and Sathiaseelan, A. Pushing the frontier: Exploring the african web ecosystem. In *Proc. WWW* (2016).

[16] Huang, L.-S., Chen, E. Y., Barth, A., Rescorla, E., and Jackson, C. Talking to yourself for fun and profit. In *Proc. Workshop on Web Security and Privacy* (2011).

[17] Klyne, G. Message headers. `http://www.iana.org/assignments/message-headers/message-headers.xhtml`, 2015.

[18] Kreibich, C., Weaver, N., Nechaev, B., and Paxson, V. Netalyzr: Illuminating the edge network. In *Proc. ACM IMC* (2010).

[19] Naylor, D., Finamore, A., Leontiadis, I., Grunenberger, Y., Mellia, M., Munafò, M., Papagiannaki, K., and Steenkiste, P. The cost of the S in HTTPS. In *Proc. ACM CoNEXT* (2014).

[20] Naylor, D., Schomp, K., Varvello, M., Leontiadis, I., Blackburn, J., López, D. R., Papagiannaki, K., Rodriguez Rodriguez, P., and Steenkiste, P. multi-context TLS (mctls): Enabling secure in-network functionality in TLS. In *Proc. ACM SIGCOMM* (2015).

[21] Poese, I., Uhlig, S., Kaafar, M. A., Donnet, B., and Gueye, B. Ip geolocation databases: Unreliable? *SIGCOMM CCR* (2011).

[22] Sanchez, M. A., Otto, J. S., Bischof, Z. S., Choffnes, D. R., Bustamante, F. E., Krishnamurthy, B., and Willinger, W. A measurement experimentation platform at the internet's edge. *IEEE Transactions on Networking (ToN)* (2014).

[23] Scott, W., Bhoraskar, R., and Krishnamurthy, A. Understanding open proxies in the wild. Technical Report. `http://netlab.cs.washington.edu/squid/paper.pdf`, 2015.

[24] Shavitt, Y., and Zilberman, N. A geolocation databases study. *IEEE Journal on Selected Areas in Communications* (2011).

[25] Sherry, J., Hasan, S., Scott, C., Krishnamurthy, A., Ratnasamy, S., and Sekar, V. Making middleboxes someone else's problem: network processing as a cloud service. *SIGCOMM CCR* (2012).

[26] Tyson, G. Dataset. `http://bit.ly/1qg7PT4`.

[27] Vallina-Rodriguez, N., Amann, J., Kreibich, C., Weaver, N., and Paxson, V. A tangled mass: The android root certificate stores. In *Proc. ACM CoNEXT* (2014).

[28] Vallina-Rodriguez, N., Sundaresan, S., Kreibich, C., and Paxson, V. Header enrichment or ISP enrichment? Emerging privacy threats in mobile networks. In *Proc. HotMiddlebox* (2015).

[29] Vallina-Rodriguez, N., Sundaresan, S., Kreibich, C., Weaver, N., and Paxson, V. Beyond the radio: Illuminating the higher layers of mobile networks. In *Proc. MobiSys* (2015).

[30] Weaver, N., Kreibich, C., Dam, M., and Paxson, V. Here be web proxies. In *Proc. PAM* (2014).

[31] Web index. `http://thewebindex.org`, 2016.

[32] World bank data repository. `http://data.worldbank.org`, 2016.

[33] Zhou, Z., and Benson, T. Towards a safe playground for HTTPS and middle boxes with QoS2. In *Proc. HotMiddlebox* (2015).