Informing Protocol Design Through Crowdsourcing: the Case of Pervasive Encryption

Anna Maria Mandalari
University Carlos III of Madrid

Marcelo Bagnulo
University Carlos III of Madrid

Andra Lutu
Simula Research Laboratory

Abstract— **Middleboxes play an important role in the modern Internet ecosystem. They perform advanced functions, but they can turn net into a hostile ecosystem for innovation. It is therefore essential, when designing a new protocol, to first understand its interaction with the elements of the path. We show how to make informed protocol design choices by using a crowdsourcing platform. We consider a specific use case, namely the case of pervasive encryption in the modern Internet. We perform large-scale TLS measurements to advance our understanding on whether wide adoption of encryption is possible in today's Internet.**

I. Introduction

The indisputable success of the Internet and, consequently, the increasing demand from end-users for more secure and faster access to the online services drives the need for continuous innovation. Modern networks often rely on dedicated hardware components generically dubbed *middleboxes* to perform advanced processing functions like, for example, enhancing application performance, traffic shaping, optimizing the usage of IPv4 address space or security.

One major criticism of middleboxes is that they might filter traffic that does not conform to expected behaviors, thus ossifying the Internet and rendering it as a hostile environment for innovation [3].

This does not mean that it is impossible to deploy new protocols, but that in order to ensure success it is imperative to first understand the interaction of the proposed solutions with the middleboxes active along the path. Recent studies [5],[4] on middleboxes behavior attempt to provide such information. However, the existing measurements use only a very small number of vantage points, e.g., in [5] only 142 measurement points are used.

Until recently, large-scale Internet measurement infrastructures necessary to perform this type of analysis were available only to large Internet players, such as Google, Akamai and large ISPs. Consequently, the lack of public access to such resources makes it hard to repeat or verify their results.

The emerging sea of crowdsourcing (such as the Amazon Mechanical Turk, Microworkers and others) can provide an accessible alternative to perform large scale Internet measurements. By expanding the traditional crowdsourcing focus from the human element to use a diverse and numerous group of end-user devices as measurement vantage points [2] we can leverage on crowdsourcing platforms to run Internet wide measurements.

In this paper, we show how to make informed protocol design choices by using a novel methodology for performing large scale Internet measurements, using a crowdsourcing solution approach. We exemplify next the efficiency of our methodology in the case of evaluating the feasibility of pervasive encryption in the modern Internet ecosystem.

In other words, we need to establish at this point whether using encryption in traditionally unsecured ports is even possible in today's Internet. In this paper, we attempt to initiate TLS connections in 68 different ports that normally do not use any form of encryption and analyze the success of the connection. This is a first necessary step towards a full comprehension of the behavior of middleboxes relative to pervasive encryption.

II. Methodology

In this section, we describe the measurements methodology we employ to assess the potential success of deploying secure protocols in the Internet using crowdsourcing. We try to establish TLS connections from a large number of vantage points (from now on, *measurement agents (MAs)*) to a large number of ports, which traditionally do not use TLS in a target server (from now on, *measurement server (MS)*), using crowdsourcing based measurements. Crowdsourcing platforms connect *employers* and *workers* from around the world. The employer is the one who creates the task (or the *"micro-job"*) for workers and specifies the parameters of its campaign, e.g., the size of the set of users performing the task or their geographical location at country level.

We argue that this approach can become an important tool for evaluating innovation solutions, primarily due to the large number of accessible and diverse measurement vantage points. Additionally, we can benefit from the freedom of deploying our own custom-designed measurement tests.

We recruit users through the Microworkers crowdsourcing platform to complete measurements on the feasibility of pervasive encryption in the current Internet ecosystem. To capture how effective would pervasive encryption actually be if deployed in today's Internet, we collect and analyze the results from more than 2,000 MAs that try to establish TLS connections in a large number of ports which normally do not use TLS.

The MAs attempt to establish both HTTP and TLS connections to 68 different ports, namely 10 well-known ports, 56 registered ports and 2 ephemeral ports. We use the success rate of the HTTP connections as the benchmark against which we compare the number of TLS successful connections. We establish then the success rate of the TLS connection by contrasting the result against the status of an unencrypted HTTP connection established in the same port. We also store and analyze in detail the

TABLE 1: Results

table

Analysis	Port 80 (%)	Average other ports (%)
Aggregated	16,5	5,8
Fixed	9	6,95
Mobile	25	4,54
Proxy	70	4,23
Non-proxy	10	6,8

TABLE 2: Packets analysis

table

Analysis	Fixed Line		Mobile	
	SYN(%)	NO SYN(%)	SYN(%)	NO SYN(%)
All	96,8	3,2	36	64
Port 80	88,3	11,7	27,7	72,3
Proxy			22,2	77,8
Non-proxy			12,7	87,3
Proxy (80)			9,6	90,4
Non-proxy (80)			36,4	63,6

server side packet exchanges.

The procedure is as follows. First, we start by asking the user to connect using a HTTP connection in port 80 to a webpage we provide. Meanwhile, in the background, HTTP and HTTPS connections are performed from the measurement devices to our servers in all the other 67 ports. In this case, data about the performance are collected in the MS.

Second, the webpage we provide contains a short form asking for additional input about the type of Internet access they are using. Finally, on the server side, we also collect and store metadata on each of the MAs that connect to our servers, such as the IP address, the user agent type, the language, and any other information included in the HTTP header.

III. Results

In the campaigns for fixed lines, we recruit 1,165 *workers* from 53 different countries. The MAs are hosted in 286 ASes overall.

For the mobile case study, we recruit 956 *workers*, from 45 different countries and 183 *ASes*.

Considering that each MA performs 68 connections to our MS, we build a complex dataset for a total of 114,228 connections [1].

Table 1 refers to the results obtained from aggregated results, for both fixed line and mobile network (label *Aggregated*), the results from users that use a fixed line and from users connected to a cellular network (labels *Fixed*, *Mobile*). We also compare the rate of successful TLS connections for users we detect using a proxy and for users that do not (labels *Proxy*, *Non-proxy*) in mobile network scenario. To better understand how proxies or other middleboxes behavior impacts the performance of the TLS protocol in unconventional ports, we focus on the packet analysis, splitting the analysis for fixed line, mobile and for users that use or not a proxy.

Table 2 refers to the percentage of SYN we receive when users try to establish a TLS connection to our MS from fixed line use case and from mobile network, considering all port and particularizing the analysis for port 80 (labels *All*, *Port 80*). Moreover, in the case of mobile network we particularize the analysis for proxy/non-proxy case (labels *Proxy*, *Non-proxy*, *Proxy (80)*, *Non-proxy (80)*).

We observe that in the case of proxies, 90% of the SYN packets are missing. While this may seem non causal at first (as the SYN packet is forwarded before the middle box actually knows whether this is a regular HTTP connection or a TLS connection), proxies usually wait until they receive the GET from the client to establish the connection to the server in order to apply their policies. This explains why in the case of TLS, we miss a high number of SYN packets.

We also try to understand if the filtering of TLS is consistent across the different ports for a given MA. In other words, if the TLS connection fails in a given port, how likely is that it will fail in other ports. In order to quantify this, we estimate the conditional probability of failure in a given port X given that the TLS connection in port 80 has failed. We choose the port 80 as it is in general a port with high failure rate. We estimate the aforementioned conditional probability for the case of fixed line and for the case of mobile network without proxies. We can see that the estimated conditional probability is around 90% in both cases (slightly higher in the fixed line case), implying that when the TLS connection fails in port 80, it is very likely that it will fail in the other ports.

IV. Conclusion

In this paper we describe an experimental model for using crowdsourcing platforms to perform large-scale Internet measurements. Our research efforts expand the traditional crowdsourcing focus from the human element to use a diverse and numerous group of end-user devices as measurement vantage points. We demonstrate the described approach while assessing the feasibility of deploying encryption by default in the Internet. We focus our crowdsourcing campaigns on building a representative dataset to show the potential success of widespread adoption of TLS encryption for existing protocols in their native ports.

We find that in average the failure rate of TLS over different ports is near the 6%. We also find that in the case of mobile networks where proxies are used, the failure rate can be as high as 70%. We conclude that it is probably feasible to roll out TLS protection for most ports except for port 80, assuming a low failure rate (6%).

We believe that our results can serve as a lower bound for the failure rate for using protocols other than expected in different ports.

[1] The data set is freely available on *http://it.uc3m.es/amandala/dataset.php*

References

[1] A. Bittau et al. The case for ubiquitous transport-level encryption. In *USENIX*, 2010.

[2] A. Doan et al. Crowdsourcing systems on the world-wide web. *ACM*, 2011.

[3] M. Handley. Why the internet only just works. *BT Technology Journal*, 2006.

[4] B. Hesmans et al. Are TCP extensions middlebox-proof? In *Workshop on Hot topics in middleboxes and network function virtualization*. ACM, 2013.

[5] M. Honda et al. Is it still possible to extend TCP? In *ACM IMC*, 2011.