Internet Censorship in Italy: an Analysis of 3G/4G Networks

Giuseppe Aceto[*,◇], Antonio Montieri[◇], Antonio Pescapé[*,◇]
[*]University of Napoli "Federico II" (Italy) and [◇]NM2 srl (Italy)
{giuseppe.aceto, pescape}@unina.it, montieri@nm-2.com

***Abstract*—Users trying to access censored content may experience different results, depending on the technique adopted to enforce Internet Censorship, that in turn depends on different factors. Administrative control of the network (i.e. the entity managing network devices) is one of such factors. To the best of our knowledge, we are the first to focus on censorship detection on 3G/4G (hereafter *mobile*) network operators, investigating the extent of differences in applying censorship inside a single country. To do so we performed an experimental campaign in Italy using the five major mobile operators. We introduce the censorship detection platform and tests we adopted, and aggregate the results according to the outcome of the tests in classes, related with censoring techniques and circumvention capabilities. Overall 15 different aggregated behaviors have been found in the experimental campaign. The analysis of measurement results reveals wide dis-homogeneity of treatment for a given censored resource across different mobile operators, with 99.5% of resources showing at least two different behaviors when probed. The discussion of reported results informs about the unexpected variability on transparency and precision of censorship, and also on effective detection and circumvention strategies, as measured from mobile networks in a single country.**

***Index Terms*—Active measurements, censorship detection, Internet censorship, Italy, mobile networks.**

I. Introduction

In recent years censorship has raised the attention of the ICT scientific community that has started to focus on the regulatory actions enforced by governments to limit or prevent the access to online information. Internet Censorship is a practice usually employed neither by the user, nor by the service operator, but by a third party (e.g., government judicial authorities or intentionally established agencies) with the purpose of impairing a client application in its ability to reach a requested resource or service. Technical means required by censorship application usually interfere with the common behavior of standard network protocols and consequently of network applications. This centralized control is then in conflict with the distributed and fair nature of the Internet and impacts with different degrees both the end users, network operators, and content providers [1]. Most of the times, users are not even able to recognize such techniques, since these give the false impression that an outage has occurred, causing the inaccessibility of the requested resource. Indeed, their tangible effect is almost always a communication error, but depending on the specific censoring technique applied, the effectiveness, the side effects, and the means for circumventing the censorship differ significantly.

Despite a growing corpus of research on Internet Censorship, to the best of our knowledge none has focused on censorship detection from mobile networks, before this work[1]. In the following we present the platform utilized for our experimentation, together with the defined and adopted methodology, and several results of the measurement campaigns aimed at the analysis of censorship as enacted by the five major 3G/4G *Mobile Network Operators (MNOs)* in Italy. Presented results provide both a big picture of the overall user experience when varying the MNO and the requested resource, and insights about the specific censorship techniques enforced by considered operators.

The paper is organized as follows: Sec. II surveys the most-related literature and positions the paper accordingly; Sec. III presents *UBICA*, the platform used to conduct the study of censorship; Sec. IV describes the methodology we propose and adopt for the analysis; Sec. V discusses the main results arisen from the experimentation for the main Italian MNOs taken into account; finally, Sec. VI ends the paper with the concluding remarks and possible future directions.

II. Related Work

Several works have analyzed the application of Internet Censorship in specific countries or scenarios: the Chinese network is likely the first and most deeply investigated, ranging from side-effects of censoring techniques (like the case of *YouTube* video service unreachable also from *outside* the country in 2010 [3]), to active probing of its complex censoring system, dubbed *Great Firewall of China* [4], [5]; other countries have been studied several times from different points of view, like Iran, in [6], [7]; Pakistan has been considered in [8], [9] and in our previous works [10], [11]. The adopted analysis methods varied greatly, including indirect inference from usage trends of circumvention applications [12], [13], analysis of logs of censoring machines [6], active probing [8], [14], [15], and monitoring of censored nodes [16]. A framework for censorship circumvention based on a combination of Mobile IPv6 and Moving Target Defense strategy is proposed in [17], [18]. For an in-depth discussion of censorship detection techniques we refer to [1]. In this work we have adopted an active measurement approach and subsequent analysis along the lines of our previous works [10], [11] (that included also Italy among other countries), with the addition of specialized HTTP testing

[1] Preliminary results of this measurement campaign have been presented as a poster in [2].

(varying the HTTP request using mobile User-Agent strings) and further analysis steps (statistics on aggregated behaviors) specifically focused on Mobile Network Operators.

III. UBICA

UBICA (*User-Based Internet Censorship Analysis*) is a platform endowing users with a *censorship monitoring* system. A central *Management Server* orchestrates a number of globally distributed *probes* of different kinds (i.e. router-based, headless client, GUI-client). In detail, UBICA provides: (i) dynamically updated censorship tests; (ii) dynamically updated targets to be tested; (iii) support for different types of probing clients; (iv) automatic censorship detection and censorship technique identification. UBICA clients are composed of a core measurement-related part, and leverage standard UNIX utilities and well-known network diagnostic tools; clients have been explicitly designed to be highly portable. The principal components of UBICA architecture are shown in Fig. 1.

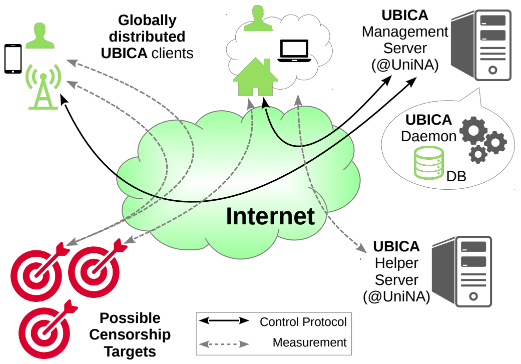
Globally distributed UBICA clients
UBICA Management Server (@UniNA)
UBICA Daemon
DB
Internet
UBICA Helper Server (@UniNA)
Possible Censorship Targets
Control Protocol
Measurement

Figure 1. UBICA architecture diagram. Home and mobile clients are shown.

The evidences of censorship are collected through *active measurements* performed by the probes that periodically retrieve the list of targets and the code (i.e. the test requirements) kept on the Management Server. Each probe then packs the results of these measurements in a report that is uploaded back to such server. A SQL database maintains the most significant information asynchronously parsed, extracted, and stored by the Management Server from the report files. Finally, the *Analysis Engine* performs a series of analyses for the censorship detection and censorship technique identification, periodically processing the results of the evidence collection stored in the database. In the following we briefly describe the different measurements performed by the UBICA probes, and the detection algorithm applied by the Analysis Engine to infer the presence of censorship.

DNS resolution. This test elicits a name resolution in order to collect evidences about this phase: given a fully qualified domain name, a DNS request of `type A` is issued from the probe towards its default resolver. The probes adopted in the experimental campaign described in this paper leverage the tool `nslookup` to send the request. In addition, the same request is issued also towards several open resolvers to distinguish between different DNS tampering techniques (either *hijacking* or *injection*) [1]. Following tests are applied to both the IP returned by the default resolver and the one returned by the control (open) resolver.

TCP reachability. To verify a possible filtering triggered by the *IP:port* pair, this test tries to set up a TCP connection with a specific target starting a three-way handshake with a given timeout.

HTTP reachability. This test issues an `HTTP GET` request: the response—or lack of it—and additional application level values are collected from the server. The *User-Agent (UA)* field of the HTTP request header is conveniently set selecting it from a list of predefined UA strings. Fitting examples of such strings are reported in Tab. I. The tool `curl` is used to send the request and collect application level information. The report from this test includes several values, such as content type, HTTP response code, redirects, etc., not reported for the sake of brevity.

Analysis Engine. The goal of this module is to tell if a resource is censored, with what kind of technique, and the associated degree of confidence. Considering a target resource and a probing client, the possible outcome provided by the Analysis Engine is one of the following: (i) *insufficient data* if there is not enough data to run the algorithm; (ii) *not censored*; (iii) *censored*. In detail, the censorship detection algorithm employed takes as input a viewpoint, target, and time interval; the viewpoint represents the network hosting the probe at a certain aggregation level (e.g., address range, service provider, country). Returned output is one of the aforementioned verdicts, and comprises also the list of censorship techniques identified and a confidence index that summarizes the number and the detection frequency of these techniques. The algorithm is made up of three steps performed in sequence; each step leverages the outcomes of the preceding one. Specifically, the first step is the identification of possible *DNS tampering*; the second one is the evaluation of *packet filtering* and *TCP connection disruption*; finally, the last step concerns the detection of *HTTP tampering* techniques. These tests are repeated periodically, and the relative frequency of a detection verdict contributes to the associated confidence level. Continuous monitoring over time allows also the detection of changes in censoring policy.

We refer to [19] for further details on the platform and the analysis algorithm it adopts.

IV. Methodology and Dataset

In this section we detail the experimental factors taken into consideration, and formalize accounted parameters and outcomes that contribute to define the methodology applied in our experimentation. Then, information on elaboration and validation of the collected dataset is provided.

A. Factors of Interest

The factors potentially impacting the enforcement and detection of censorship are summarized in Tab. I. Focusing our analysis on the Internet Censorship in Italy, we have selected *the four major MNOs* (i.e. H3G, TIM, Vodafone, Wind) accounting for the 96.6% of the Italian market, and

Table I
SUMMARY OF FACTORS AND CONSIDERED VALUES.

Factor	Values
MNO	H3G, PosteMobile, TIM, Vodafone, Wind
Target	200 censored targets
DNS	Default (MNO-provided), Open (Public DNS resolvers)
User-Agent (UA)	Safari 5.1 (iPhone - iOS 5.0), IEMobile 7.11 (HTC Touch 3G - Windows Mobile 6.1), Google Chrome 41.0 (Desktop - Windows 7)

Table II
AGGREGATED BEHAVIORS.

Parameter	Values
DNS comparison	Same (S), Different (D)
Dependence on UA	UA-dependent (U), UA-independent (I)
Outcome from Default DNS resolution	Legitimate (L), Failing IP (F), No IP (N), Block Page (B), Missing URL (M)
Outcome from Open DNS resolution	Content (C), `NXDOMAIN` (X), Connection Timeout (T), TCP Reset (R), HTTP Error (E)

the leader in the field of virtual MNOs (i.e. PosteMobile), that holds the 52.1% of the market share of the Mobile Virtual Network Operators (MVNOs) [20]. In terms of UBICA, each MNO represents a viewpoint. The list of *targets* to be tested contains URLs that have been labeled as censored by different sources, including an observatory of documents of the Italian judicial authority [21] and a crowdsourcing online report [22]. Preliminary experimentation conducted with UBICA by probes connected to our institution network (Rete Italiana dell'Università e della Ricerca [23]) allowed us to choose those targets that were accessible by these probes and therefore evaluated as *not censored* by UBICA when accessed from the outside of selected MNOs' networks. This process led to a smaller list of 200 *targets*, that have been used for the experimentation described in this work. Name resolutions are elicited through both the *default* (MNO-provided) resolver and a list of *open* (Public DNS) resolvers [8] in order to detect possible DNS-based censorship techniques enforced (*DNS tampering*). Lastly, the list of *UA strings* usable in HTTP reachability measurements has been conveniently set to test both mobile and desktop agents of different kinds (see Tab. I).

When a client—representing a generic mobile user in our scenario—tries to retrieve a content from a target, it experiences a number of distinct behaviors depending on the specific combination of factors shown in Tab. I. In detail, given a MNO and a target, a user can experience: (i) same or different default and open DNS resolutions; (ii) possible redirections of the request dependent or not on the UA; various outcomes obtained through the (iii) default and (iv) open DNS resolutions. Tab. II reports the list of parameters taken into account and their possible outcomes. To ease the dissertation of feasible behaviors—note that not all combinations of values in Tab. II can happen—we refer to each quartet of outcomes as an *aggregated behavior*. Hereafter we describe the outcomes from querying the default and open DNS resolvers.

When a user leverages the default resolver, the response returned by the DNS server might not correspond to the **legitimate (L)** DNS database entry. Indeed, a censoring server replies with a *forged response*, which is a Resource Record of `type A` containing an IP address different from the one that would be obtained from the legitimate resolution of the requested resource (i.e. *DNS hijacking*) [1]; specifically:

- **Failing IP (F)**: a non-Internet-routable IP address (e.g., private or shared address space);
- **No IP (N)**: an empty IP address (DNS returns `NODATA` and `NOERROR`);
- **Block Page (B)**: an IP address of a web server returning a page stating that the target hostname has been explicitly blocked and possibly informing about the cause of the block;
- **Missing URL (M)**: an IP address of a web server returning a page stating that the requested hostname is non existent or misspelled.

Another DNS-based censoring technique, *DNS injection*, is enforced by a middlebox reading the DNS request and responding with a spoofed fake DNS reply, fooling the client. This technique is hard to circumvent, while DNS hijacking can be easily circumvented changing the default DNS resolver (acting as the *censoring device*), with an open resolver. This way, a user can attain open access to the Internet, but might still experience several erroneous outcomes that prevent the correct retrieval of the requested **content (C)**. Indeed, the open resolver could reply with a `NXDOMAIN` response (i.e. "no such domain") if the domain name referenced in the DNS query does not exist (possibly has expired before the censorship is revoked). In addition, the TCP connection might be disrupted by incurring a **connection timeout (T)**, or forcing its prompt termination through a **TCP reset (R)**. Finally, the client could receive an **HTTP error (E)** response message (i.e. 4xx or 5xx status codes).

B. Experimental Dataset

To collect the dataset, the UBICA reference architecture shown in Fig. 1 has been set-up with headless clients (running Kubuntu 14.04), and connected to the Internet through mobile terminals acting as gateways via tethering USB. The dataset reports a verdict for each combination of the factors presented in Tab. I. Specifically, the collection algorithm performed by means of UBICA has returned a total of 6000 measurements in each run. This raw data was then conveniently post-processed in order to extract the aggregated behaviors, that summarize the single verdicts into outcomes effectively influencing the user experience. In details, an aggregated behavior has been associated to each of the 1000 possible $(target, viewpoint)$ pairs. This process has been conducted through automated data-cleansing and detection procedure accomplished by means of UBICA, followed by a manual validation of the results.

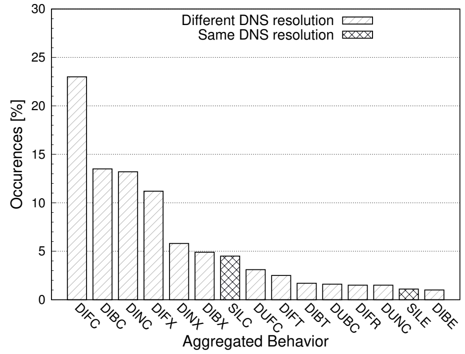

Figure 2. Aggregated behaviors (topmost 90% of occurrences).

V. Results

The aim of this section is to discuss the most interesting results stemming out from the experimental campaign conducted between February and March 2016 adopting the methodology introduced in Sec. IV.

A. General Overview of the User Experience

In the following we provide an overview of the factors that mostly affect the browsing experience of a mobile user requesting a censored target.

Our experimentation shows that choosing between default or open DNS servers heavily impacts the user experience, for all MNOs. Considering the aggregated behaviors we found that censorship is enforced through DNS tampering (more specifically, *hijacking*) for 92.2% of the $(target, viewpoint)$ pairs. When the returned Resource Record is legitimate (no DNS tampering), users correctly retrieve the requested page in 6% of the cases (corresponding to targets that are not censored by some of the MNOs). Non DNS-based censorship techniques (discussed in Sec. V-E) account for 14.4% of occurrences.

To further detail the analysis, we refer to Fig. 2, that shows the topmost 90% of the aggregated behaviors observed, ordered by decreasing occurrence percentage (unreported behaviors account for less than 1% each of the occurrences). Interestingly, for the three most frequent behaviors the default and open DNS resolutions differ (D); moreover, the target website is always censored (F/B/N) when the default DNS server is chosen, whereas it is correctly reached (C) leveraging an open resolution. These results highlight the prevalence of the DNS hijacking variant. The behaviors corresponding to DNS injection, i.e. same resolution (S) but censored content, are not present in the dataset, as only legitimate (L) addresses are retrieved (the two cases highlighted in Fig. 2). Censoring techniques operating just at the DNS stage are not affected by HTTP User-Agent, thus all these behaviors are independent (I) of the UA specified in the requests following the resolution.

B. Impact of the MNO on Aggregated Behaviors

One of the goals of our experimentation is to evaluate how and how much the MNO adopted to access to the Internet influences the aggregated behaviors observed. With reference to Fig. 3, it can be seen that only 1 target (0.5%) presents the same aggregated behavior for all the MNOs, whereas half of them (100 out of 200 targets) exhibits 3 different behaviors. These results confirm that, by varying the MNO that offers connectivity, a user is highly likely to experience different outcomes, even requesting the same content. This will not necessarily happen for all the targets and all the MNOs, as 95% of the targets exhibit at most 4 aggregated behaviors, having at least 1 behavior in common between 2 MNOs.

Table III
Pair-wise variation in censorship application.

MNO	**H3G**	**PosteMobile**	**TIM**	**Vodafone**	**Wind**
H3G	0%	92.5%	32.5%	94%	75%
PosteMobile		0%	99%	60%	95%
TIM			0%	99.5%	65.5%
Vodafone				0%	65%
Wind					0%

The variations shown in the aggregated behaviors between accounted MNOs are summarized in Tab. III, where for each pair of MNOs the percentage of targets that show different aggregated behaviors in the pair (0% means all targets behaved the same, 100% all targets behaved differently) is reported. On the one hand, the operators that disclose the lowest pair-wise variation (i.e. most frequently behaving in the same way) are H3G and TIM that present different behaviors for 32.5% of the targets. On the other hand, TIM and Vodafone have almost always different aggregated behaviors (99.5%). Notably, PosteMobile and Wind exhibit the same aggregated behaviors for only 10 out of 200 targets, although the former is a *MVNO* and offers its services leasing the network infrastructures and radio spectrum from the latter.

C. Dependence on (mobile) User-Agent

Another aspect we analyzed is the dependence of the observed behaviors on the class (mobile or desktop) of users that makes a request, as characterized by one of the User-Agent (UA) strings reported in Tab. I. With this aim, we have discriminated the requests subject to a redirection dependent on the UA, from those which are not subjected to redirections. We have found that at least 10% of the requests exhibit a redirection that depends on the UA for all the MNOs. The greatest impact is attained by Wind (14% of the requests), whilst the lowest is obtained with Vodafone (9.5%). H3G, PosteMobile, and Tim show a dependence from the UA for 11% of the requests. Thus, the UA turns out to be a factor

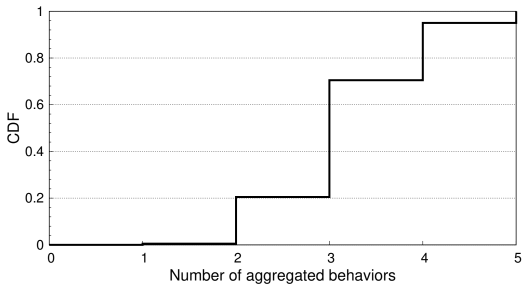

Figure 3. Distribution of targets exhibiting a different number of aggregated behaviors after a change in the MNO used to access to the Internet.

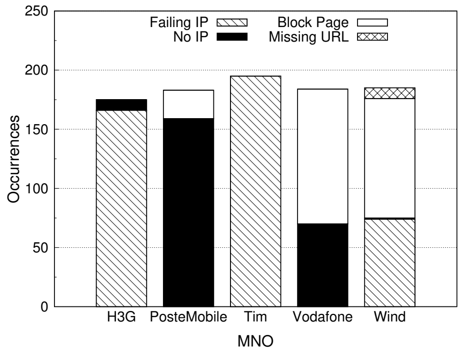

Figure 4. DNS hijacking variants enforced by each MNO. Only Wind-provided resolver returns all four types of forged responses.

affecting the final outcome of censorship, as the MNOs force a specific user class to be redirected towards the content or a block page, but this occurrence is secondary, regarding 11.3% of overall requests.

D. DNS-based Censorship Techniques

As we have seen in previous sections, the censorship techniques enforced by the Italian MNOs are mainly based on DNS hijacking. Interestingly, we have found that forged responses obtained strongly depend on the specific DNS resolutions made by the default resolver of each MNO. The most frequent of such responses contains a non-routable *Failing IP* address, whereas only for one MNO (Wind) the returned IP address points to a web server that hosts a Missing URL page.

Going into the details of this analysis, we have found that each MNO prefers to enforce different DNS-based censorship techniques through specific forged responses. It is interesting to note that only Wind uses all four types of forged responses described in Sec. V-A.

Fig. 4 highlights DNS hijacking techniques employed by every MNO. The return of a Failing IP is the most popular censorship technique for 3 out of 5 MNOs. Only PosteMobile- and Vodafone-provided resolvers do not return such a type of response at all, preferring to provide an empty IP address (i.e. No IP), or redirect the requests toward a block page, respectively. Note that despite PosteMobile is a Virtual MNO serviced by Wind's network infrastructure, in 87% of the cases it returns a No IP response, which has been observed for only one target with Wind. Generally speaking, a No IP response is returned for at least a target by 4 MNO-provided resolvers. Indeed, TIM enforces one single DNS-based censorship variant based on the reply of a Failing IP address for all censored targets. Moreover, a user relying on TIM or HG3 can't really know if the unreachability of the requested resource is due to censorship or caused by a communication error. Conversely, when the default resolver returns an explicit block page (in the case of PosteMobile, Vodafone, and Wind), users are able to realize that a given target is actually censored and they are possibly apprised about the causes of the block. Finally, a Missing URL page is obtained for 9 targets just with the Wind-provided resolver.

As shown in Fig. 4, at least 87.5% of the targets—for H3G—are censored through a DNS-based technique. The maximum value is reached by TIM, that enforces DNS hijacking for 195 out of 200 targets. In order to circumvent DNS hijacking, we have elicited a name resolution for every target leveraging also a number of Public DNS resolvers. Thus, accordingly to these results, it is interesting to investigate the faculty of obtaining the requested page through this open DNS resolution. Considering all the requests toward each target from every MNO, changing the default resolver to an open one allows to correctly obtain the content for 56% of the requests.

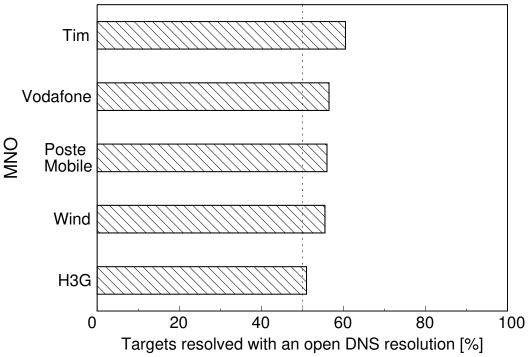

Figure 5. Percentage of targets resolved for each MNO leveraging an open DNS resolution.

Fig. 5 provides details about the different outcomes observed between MNOs when an open DNS resolution is leveraged. As expected, for every MNO at least 50% of censored resources can be correctly retrieved changing the default resolver. The greatest impact is obtained for TIM, for which 121 out of 200 targets (60.5%) can be reached through an open resolution, whereas the lowest one is observed for H3G (51%). We can note that even leveraging a Public DNS resolver, the requested resource might not be correctly retrieved. In these cases, averaging the occurrences of the observed outcomes over the 5 MNOs, we obtain (i) a NXDOMAIN error response for 23.6% of the targets, (ii) a connection timeout for 6.4%, (iii) one HTTP error code between 403, 404, 500, 503, 504 for 4.9%, and (iv) the termination of the connection by TCP reset for 3.1%. These outcomes reveal the presence of non DNS-based techniques, that can be also enforced in combination with DNS hijacking. We analyze in more detail these cases in the following section.

E. Other Censorship Techniques

Our results reveal that some targets are unreachable despite a correct default (or open) DNS resolution. As a consequence, just changing the default DNS resolver to an open one would not be an effective circumvention technique. We have investigated the root causes of this unreachability to distinguish temporary and permanent failures from actual cases of censorship enforced by the Italian MNOs. With this aim, a secondary experimental campaign has been performed to compare the outcomes of the targets that were still unreachable even in this supplementary campaign, despite a correct DNS resolution.

We have found that, also for these, targets are treated differently by the MNOs. In all the cases taken into account we have found that Vodafone blocks the access to the resource intercepting the request through a *proxy server* (whose IP address falls in the range of Vodafone networks). This is a fine-grain censorship enforcement, as it can inspect both the request (in the HTTP `host` header) and the response (looking for specific content), therefore this technique can in principle reduce overblocking (i.e. unintended blocking of resources due to poor discrimination capabilities of the censoring technique) to a minimum. The proxy returns an HTTP response with a status code "5xx server error". Conversely, the other MNOs always block the requests disrupting the connection during the *TCP handshake*. TCP sessions are disrupted either sending a *TCP RST* (reset) packet or filtering the TCP connection altogether (i.e. TCP SYN packets don't receive any response). This is the most coarse-grained censorship technique, even more than DNS-based ones, as in case the blocked $(IP, port)$ pair corresponds to a virtual hosting or CDN service, it blocks many online resources potentially completely unrelated.

For only one target the requested web page is correctly retrieved from some operators (H3G and Wind), possibly the beginning or the end of a censorship enforcement on the target.

VI. Conclusion

In this work we have investigated Internet Censorship in Italy, for the first time in literature focusing on Mobile Network Operators (MNOs). The main goals of the analysis were to infer (i) what censoring techniques are employed by MNOs, (ii) if and how much the usage of a specific operator affects the experience of the user, (iii) if mobile terminals are treated differently from desktop ones (while connecting through MNOs), and (iv) if Mobile Virtual Network Operators (MVNOs) behave similarly to the MNO they lease the infrastructure from. We adopted the UBICA platform to perform experiments on the five major Italian MNOs (including one MVNO and the related infrastructure MNO), modeling the resulting measurement outcomes in classes (aggregated behaviors) characterizing the detected censoring techniques.

From the analysis of UBICA responses, we found that (i) most employed techniques are DNS-based (hijacking), but a small number of cases (14.4%) shows evidence of TCP tampering and HTTP tampering; (ii) the specific operator used to access the Internet seriously affects the experience of the user, that for a given target can obtain an explicit block page, an obscure failure, or the original content, depending on the operator; (iii) mobile terminals, featuring characteristic HTTP header strings, are treated differently from desktop ones in a non-negligible fraction of cases (up to 14%, depending on the MNO); (iv) the considered MVNO presents a behavior wildly different from the underlying MNO leasing the infrastructure (for 95% of the targets, the censoring technique is different).

These findings shed light on the transparency, censorship detection, and circumvention strategies in the main 3G/4G networks in Italy. Notably, experiments considering only one MNO will provide no information about other operators and also on the served MVNO. As a consequence, censoring behaviors are not generalizable and a common technique for censorship circumvention is hardly definable. We plan to perform additional measurements (over time, on more countries) and analyses, to further extend the coverage of our findings.

References

[1] G. Aceto and A. Pescapé, "Internet censorship detection: A survey," *Computer Networks*, vol. 83, pp. 381 – 421, 2015.

[2] G. Aceto, A. Montieri, and A. Pescapè, "Internet Censorship in Italy: a first look at 3G/4G networks," in *International Conference on Cryptology and Network Security*. Springer, 2016, pp. 737–742.

[3] M. L. Mueller, "China and global Internet governance," *Access contested: security, identity, and resistance in Asian cyberspace, ed. Ronald J. Deibert, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain*, pp. 177–194, 2012.

[4] X. Xu, Z. M. Mao, and J. A. Halderman, "Internet censorship in China: Where does the filtering occur?" in *Passive and Active Measurement*. Springer, 2011, pp. 133–142.

[5] G. C. Feng and S. Z. Guo, "Tracing the route of China's Internet censorship: An empirical study," *Telematics and Informatics*, 2012.

[6] S. Aryan, H. Aryan, and J. A. Halderman, "Internet Censorship in Iran: A first look," in *3rd Workshop on Free and Open Communications on the Internet*. USENIX, 2013.

[7] M. Wander, C. Boelmann, L. Schwittmann, and T. Weis, "Measurement of globally visible dns injection," *IEEE Access*, vol. 2, pp. 526–536, 2014.

[8] Z. Nabi, "The anatomy of web censorship in pakistan," in *Presented as part of the 3rd USENIX Workshop on Free and Open Communications on the Internet*. Berkeley, CA: USENIX, 2013. [Online]. Available: https://www.usenix.org/anatomy-web-censorship-pakistan

[9] S. Khattak, M. Javed, S. A. Khayam, Z. A. Uzmi, and V. Paxson, "A look at the consequences of internet censorship through an isp lens," in *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 2014, pp. 271–284.

[10] G. Aceto, A. Botta, A. Pescapè, N. Feamster, T. Ahmad, and S. Qaisar, "Monitoring Internet Censorship with UBICA," in *Seventh International Workshop on Traffic Monitoring and Analysis (TMA'15) Barcelona, Spain*, April 2015.

[11] G. Aceto, A. Botta, A. Pescapé, M. F. Awan, T. Ahmad, and S. B. Qaisar, "Analyzing Internet censorship in Pakistan," in *IEEE 2nd International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (IEEE RTSI)*, Bologna, Italy, Sep. 2016.

[12] A. Di Florio, N. V. Verde, A. Villani, D. Vitali, and L. V. Mancini, "Bypassing censorship: a proven tool against the recent internet censorship in turkey," in *Software Reliability Engineering Workshops (ISSREW), 2014 IEEE International Symposium on*. IEEE, 2014, pp. 389–394.

[13] L. Dixon, T. Ristenpart, and T. Shrimpton, "Network traffic obfuscation and automated internet censorship," *IEEE Security Privacy*, vol. 14, no. 6, pp. 43–53, Nov 2016.

[14] A. Filastò and J. Appelbaum, "OONI: Open Observatory of Network Interference," in *Proc. 2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI 2012)*, August 2012.

[15] S. Burnett and N. Feamster, "Encore: Lightweight measurement of web censorship with cross-origin requests," *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 4, pp. 653–667, Aug. 2015. [Online]. Available: http://doi.acm.org/10.1145/2829988.2787485

[16] P. Winter and S. Lindskog, "How the Great Firewall of China is blocking Tor," in *Proc. 2nd USENIX Workshop on Free and Open Communications on the Internet (FOCI 2012)*, August 2012.

[17] V. Heydari, S.-i. Kim, and S.-M. Yoo, "Anti-censorship framework using mobile ipv6 based moving target defense," in *Proceedings of the 11th Annual Cyber and Information Security Research Conference*, ser. CISRC '16. New York, NY, USA: ACM, 2016, pp. 7:1–7:8. [Online]. Available: http://doi.acm.org/10.1145/2897795.2897815

[18] V. Heydari, S. i. Kim, and S. M. Yoo, "Scalable anti-censorship framework using moving target defense for web servers," *IEEE Transactions on Information Forensics and Security*, vol. PP, no. 99, pp. 1–1, 2017.

[19] G. Aceto, "Monitoring Internet censorship: the case of UBICA," March 2014, (PhD Thesis). [Online]. Available: http://www.fedoa.unina.it/9786/

[20] "MVNO News - Osservatorio MVNO," http://www.mvnonews.com/osservatorio-mvno/, Feb. 2017.

[21] M. d'Itri, "Osservatorio sulla censura di internet in italia," https://censura.bofh.it, Feb. 2017.

[22] Herdict, "Web platform for reporting inaccessibility of web sites," http://www.herdict.org, Feb. 2017.

[23] "GARR - La Rete Italiana dell'Università e della Ricerca," http://www.garr.it/, Feb. 2017.