

# Automated Discovery of Internet Censorship by Web Crawling

Alexander Darer  
Oliver Farnan  
Dept. Computer Science  
University of Oxford  
Oxford, UK

Joss Wright  
Oxford Internet Institute  
University of Oxford  
Oxford, UK

## ABSTRACT

Censorship of the Internet is widespread around the world. As access to the web becomes increasingly ubiquitous, filtering of this resource becomes more pervasive. Transparency about specific content and information that citizens are denied access to is atypical. To counter this, numerous techniques for maintaining URL filter lists have been proposed by various individuals, organisations and researchers. These aim to improve empirical data on censorship for benefit of the public and wider censorship research community, while also increasing the transparency of filtering activity by oppressive regimes.

We present a new approach for discovering filtered domains in different target countries. This method is fully automated and requires no human interaction. The system uses web crawling techniques to traverse between filtered sites and implements a robust method for determining if a domain is filtered. We demonstrate the effectiveness of the approach by running experiments to search for filtered content in four different censorship regimes. Our results show that we perform better than the current state of the art and have built domain filter lists an order of magnitude larger than the most widely available public lists as of April 2018. Further, we build a dataset mapping the interlinking nature of blocked content between domains and exhibit the tightly networked nature of censored web resources.

## KEYWORDS

censorship; DNS; filtering; transparency; monitoring

## ACKNOWLEDGMENTS

This work was supported by EPSRC through the Centre for Doctoral Training in Cyber Security, University of Oxford and The Alan Turing Institute under the EPSRC grant EP/N51012.

Alexander Darer & Oliver Farnan are funded by the Centre for Doctoral Training in Cyber Security.

Joss Wright is partially funded by the Alan Turing Institute as a Turing Fellow under Turing Award Number TU/B/000044.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*WebSci '18, May 27–30, 2018, Amsterdam, Netherlands*

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5563-6/18/05...\$15.00

<https://doi.org/10.1145/3201064.3201091>

## 1 INTRODUCTION

The effort expended by censorship regimes around the world attempting to filter Internet resources they deem to be too sensitive or against the morality of their own interests is on-going. As Internet access has become more ubiquitous, the scale of deployed filtering systems is increasing. A recent study has shown that blocking the reachability of popular sites at national levels is widespread and disruptive [24][29]. Advocates for free-speech and a free Internet push for transparency and openness, while censors attempt to repress the flow of certain information within their networks. Key to this is the blocking of specific webpages and the URLs that point to them.

In response to large scale filtering of web resources, there have been numerous studies over recent years aimed at determining the type of content being blocked in different countries. Of particular interest are periods of time when blocking has occurred and the development of techniques to monitor filtered URLs and keywords [12][13][18][23][33][41].

We introduce an approach for discovering filtered domains in different countries at scale and reasonable cost. The system is fully automated and does not require per-country expertise or cooperation - meaning that the safety of individuals within censored regimes won't be compromised. Our method applies web crawling techniques to find blocked content and uses a seed list of known filtered URLs to initiate the search. We make use of DNS servers within a target country as measurable devices. This allows us to monitor the filter status of individual domains and sub-domains without human intervention. The system is recursive so newly discovered filtered URLs are fed back into the search to allow on-going measurement. Results from our experiments using four different test countries have shown that our approach can be used to find filtered URLs that are not present in the original seed lists. Furthermore, we collect data about the linked nature of various filtered domains to gain further insight into how different pieces of filtered content are associated.

### 1.1 Related Work

Over the last decade there have been many approaches for detecting censorship of the Internet around the world. Of these, many are country specific and have focused on China [10][18][22][25][40], Indonesia [21][36], Iran [7][8], Pakistan [2][27] and Thailand [19] among others.

The most widely adopted and current URL filter lists are maintained by the *CitizenLab* [9]. They are constructed using local knowledge and reports of filtering in different countries and collate data from different sources such as *OONI* [16].

Developing new techniques for discovering filtered URLs is a challenging problem. Yet, this is a rich research field with numerous techniques published over recent years [4]. The use of DNS as a means of testing censorship of web content is not new, but can be advantageous due to its scalability and remote nature [35]. These attributes make DNS a common tool for other censorship monitoring architectures such as *UBICA* [3], *FilteredWeb* [13] and *CensMon* [33].

Building in-depth and accurate URL filter lists is an important aspect for censorship research. These collections are in widespread use among the research community for various different measurements and tests for internet reachability, web content blocking and circumvention techniques [30][38]. Furthermore, the subsequent and on-going maintenance of these lists provides opportunities for insight into the condition of internet filtering around the world. The data collected by the aforementioned monitoring architectures is vital if we are to construct a model of censorship as it develops.

## 1.2 Contributions

This paper introduces a new approach for discovering filtered domains within target censorship regimes. We have created an implementation of the technique and, through experimentation, shown it to be an effective tool for building URL filter lists. Furthermore, our results reveal that the approach has found significantly more filtered URLs for the test countries than were currently available in the largest public filter lists. Our formalised contributions are:

- A new approach for discovering previously unknown filtered domains
- Experimental analysis of the technique through measurement of filtering activity within four known censorship regimes
- Category breakdown of the types of content being blocked within these regimes
- Analysis of forward filtered links and filtered backlinks of webpages on filtered domains

A substantial research output from this body of work is a test list containing a large number of currently filtered domains within China, Indonesia, Iran and Turkey that have previously been unpublished. We aim to make this list available as soon as possible to the wider censorship research community.

## 2 TRAVERSAL OF FILTERED WEBPAGES

Traversal between webpages using embedded hyperlinks is the most widely used method for content discovery on the Web. Our aim is to exploit the connections between different filtered webpages to efficiently crawl sites in search for more blocked content. An important assumption for this approach is that different filtered webpages do indeed link to others. We demonstrate this through experimental analysis of four different countries that are known to filter websites via DNS manipulation. The method we describe is not dissimilar to conventional web crawling techniques widely used by large search engines. Given this, we aim to build a new dataset that contains information pertaining to backlinks<sup>1</sup> of filtered webpages.

This technique works on a simple premise - filtered webpages contain links to other filtered webpages. We begin the discovery




Figure 1: High-level overview of filtered webpage traversal

by seeding the system with a number of known filtered URLs with the presupposition that these will contain hyperlinks to further blocked content. A high-level overview of the technique is shown in Figure 1 and works as follows:

- (1) Start with a list of known filtered URLs for country  $c$
- (2) Retrieve webpages for all known filtered URLs in our list
- (3) Extract any URLs from the downloaded webpages
- (4) Isolate the URLs that are filtered in country  $c$  from the extracted URLs
- (5) Add the newly identified filtered URLs to the list, then goto step 2

A number of hyperlinks in any webpage will point to resources that do not provide utility for our discovery. We ignore any URLs that point to static HTML assets - such as javascript, css or image files and also remove any self-referencing URLs - hyperlinks to the same domain for the webpage. We aim to reduce the possibility of having the crawler becoming stuck in cliques such as affiliate or adult site networks this way. For purposes of analysis of the approach we visit each unique URL only once.

### 2.1 Methods for Checking Webpage Availability

A key part of this system is the capability to determine if a domain is filtered in a certain country or not. This is important as crawling unnecessarily large portions of the Web will make the discovery inefficient. Furthermore, an inaccurate checking procedure could produce large numbers of false positives - detracting from the usefulness of the results. We must also be prudent in regards to ethical issues when taking internet measurements using such a system. Since, we will be probing for blocked content, we want to ensure that the safety of individuals is not compromised as described by [14].

Given these requirements, we employ a checking system that uses DNS infrastructure to determine if domains are being filtered within a given country. The exploitation of DNS as a means for blocking access to certain web resources is in widespread use

<sup>1</sup>Backlinks are hyperlinks that point to a certain page from other pages.

around the globe [15][24][25][30]; and because resolvers operated by large Internet Service Providers (ISPs) are often open, we can use these as a basis for our measurements.

To determine the filter status of a domain, we require a globally non-censored DNS server as a control and a measurement DNS server located within a target country. Using these, we run through a process to comprehensively check if a domain is filtered by the measurement server. We examine responses to DNS queries made for the test domain to determine if the server is poisoned [15] or acting rogue. This procedure is described in Algorithm 1.

The following six checks are used to ascertain the filter status of domains:

- (1) A DNS query is intercepted in the target country when sent to a non-existent DNS server
- (2) The measurement server times out but the control server does not
- (3) The measurement server responds with a private IP address but the control does not
- (4) The measurement server resolves an IP that times out on an HTTP GET request, whereas the resolved IP from the control server does not
- (5) The control server resolves an IP that times out on an HTTP GET request, whereas the resolved IP from the measurement server does not
- (6) The content length of the webpages from each resolved IP differs by more than a defined percentage amount

If the results from any of the above are found to be positive, we consider the domain to be filtered by the measurement DNS server.

This procedure has a number of useful features. Firstly, it does not require cooperation of any person or individual within a censored country, it solely makes measurements on infrastructure. Second, it is an efficient mechanism that is scalable and yields fast results. Third, we can perform measurements from outside target countries giving us the capability to analyse filtering within a wide array of censorship regimes.

The limitation with this approach is that we lose fine grain information about individual URLs that may be blocked. This is due to the sole use of DNS servers as a checking mechanism - since one can only query about entire domains or sub-domains. Further, we require that measurement DNS servers respond to queries from remote countries in the same manner they do for queries made domestically. Yet, as we need to make a trade-off between effectiveness, efficiency and ethical considerations, this method is sufficient for use to generate good results to show the usefulness of this technique for discovering filtered domains as a whole.

## 2.2 Ethical Considerations

Implications of censorship measurements open us up to a number of ethical issues that we must give thought to. First and foremost, it is imperative that we do not cause harm to any persons or organisations that are unaware of our actions, intent or motivation. This can be a causal problem in numerous ways, not least because testing of internet filtering often requires sending network traffic to and from—in and out—of countries controlled by censorship regimes [11]. Certain studies within the field have required the use of aware

---

### Algorithm 1 Pseudocode for Domain Filtering Check

---

```

MAXDIFF  $\leftarrow p$  {content length % difference that indicates filtered
domain}
dom {domain to check}
mDNS {measurement DNS server in target country}
mDNSFake {fake DNS server in target country}
cDNS {control DNS server}

{the following variables are NULL on timeout}
1: mFakeIP  $\leftarrow \text{resolveDNS}(dom, mDNSFake)$ 
2: mIP  $\leftarrow \text{resolveDNS}(dom, mDNS)$ 
3: cIP  $\leftarrow \text{resolveDNS}(dom, cDNS)$ 
4: mIPContent  $\leftarrow \text{httpGETRequest}(mIP)$ 
5: cIPContent  $\leftarrow \text{httpGETRequest}(cIP)$ 

6: if mFakeIP  $\neq \text{NULL}$  then
7:   return TRUE {DNS query was intercepted in target country}
8: end if

9: if mIP == NULL and cIP  $\neq \text{NULL}$  then
10:  return TRUE {mDNS is rogue server}
11: end if

12: if mIP  $\neq cIP$  then
13:  if isPrivateIP(mIP) == TRUE and isPrivateIP(cIP) ==
FALSE then
14:    return TRUE {mIP is private address}
15:  end if
16:  if mIPContent == NULL and cIPContent  $\neq \text{NULL}$  then
17:    return TRUE {mDNS is rogue server}
18:  end if
19:  if mIPContent  $\neq \text{NULL}$  and cIPContent == NULL then
20:    return TRUE {mDNS is rogue server}
21:  end if
22:  if length(mIPContent)/length(cIPContent) > MAXDIFF
then
23:    return TRUE {mIP points to incorrect content}
24:  end if
25: end if
26: return FALSE {dom not filtered}

```

---

volunteers who are located within countries of interest. While these individuals are generally knowledgeable of the motivation of the study and potential ramifications of their actions if implicated, this is not something we as researchers should take lightly. In many cases, it is simply not appropriate to use human participants for this type of work. Furthermore, there are a number of legal issues with measurements of censorship based on the techniques used - especially if inference is made using direct observations within a target country [42].

We must also consider the use and deployment of these kinds of discovery techniques by antagonists. Since we aim to build a system that can automatically find alternative content that is blocked based on *known* blocked content, such a framework could be utilised in an adverse way to filter further web resources. Unfortunately, we

cannot guarantee that this use-case will never occur given the fact that censors generally do not publish technological details about their infrastructure and systems.

These concerns should not however reduce our willingness to practice this kind of research. If considerable effort is made to ensure our measurements will not affect individuals, we are able to provide empirical data concerning censorship around the world. This can give us as researchers a substantial insight into complex socio-political issues that are of benefit to the community and are of wider public interest given the state fragile of international relations. Moreover, our proposed approach does not pose a risk to individuals or rely human volunteers and vulnerable subjects. We take measurements directly from infrastructure in a manner that the services were originally designed for.

### 3 EXPERIMENTAL ANALYSIS

We conduct experiments on four different countries with an aim to build domain filter lists that are longer and more in-depth than are currently available. This was achieved using an implementation of the approach written in *Python* with the following parameters:

- Control DNS server: 8.8.8.8
- *MAXDIFF* (content-length difference that indicates filtering): 50%
- Filter check timeout: 10 seconds
- Maximum recursion depth<sup>2</sup>: 100
- Seed URLs obtained from the CitizenLab filter lists [9]

The *MAXDIFF* value is used based on a study that found the content-length of censorship block pages are 95% likely to differ by more than 50% compared to the genuine page [1]. Further, we ensure that the system does not follow links that self-reference the parent site - this is to say we attempt to stop looping behaviour with pages that link to others on the same domain. Also, we never revisit a URL that has previously been seen - it will be counted in the statistics we gather, but not checked again.

The target countries tested were: China, Indonesia, Iran and Turkey. Each experiment ran for seven days, or until no more filtered domains were found. The DNS servers used for each test country are shown in Table 1. The real DNS servers were selected from large ISPs in the target countries and the fake from the pool of unallocated IP addresses also owned by the same ISPs. We do this because we aim to take measurements on mass infrastructure within the target countries rather than smaller organisations or individuals. It is also extremely important to keep testing within the bounds of how the network was designed to operate. We *only* make DNS queries and do not send any other crafted (or benign) packets to the target servers. While we acknowledge that our requests will likely contain queries for censored sites, the infrastructure we measure has been developed to process these kinds of transactions.

#### 3.1 Results

Table 2 depicts the number of unique URLs extracted over the course of each experiment and how many of those were filtered in the given country. We also perform a count on the number of unique filtered domains within the list of filtered URLs. As a measure for the breadth of each run, the Alexa Top 1000 domains were removed

so we can analyse how deep the system is able to penetrate to lesser known sites with lower numbers of visitors and backlinks.

In total, we extracted over 80 million URLs from filtered web pages, of which 5.7 million were themselves from a filtered domain. The number of blocked domains identified for Turkey and Indonesia are an order of magnitude larger than those found for China and Iran. This is due to the widespread censorship of adult related sites within these particular censorship regimes. Turkey passed a law in 2007 prompting the explicit blocking of over 80,000 sites, of which many contained adult content [6], and Indonesia, a similar ban in 2010 [20] & 2017 [31].

We perform a comparison with the most widely available public URL filter lists, maintained by the CitizenLab. To ensure a fair comparison, we run these lists through our filtering check and report those numbers. The figures are shown in Table 3. From this we can show that we have performed efficiently and identified more filtered domains than were present in the original seed lists. To gain further insight into the types of content filtered in Turkey and Indonesia, we remove the adult domains to create separate counts for better comparison.

Our results demonstrate that this approach is effective at finding previously unknown filtered domains. A major advantage of this technique is that *only* URLs from filtered domains are visited, meaning that we can achieve efficient web crawling.

### 4 FURTHER ANALYSIS

The experiments we performed have yielded an interesting dataset that lends itself to further investigation. We are able to track the paths that lead to filtered content by analysing routes taken by the crawler. This gives us a useful base for examining how deeply connected collections of blocked sites are. Further, we can identify the backlinks of filtered pages and the outbound (forward) links to other filtered sites, we can discover networks of filtered sites.

We find that the results found in Turkey and Indonesia contain large numbers of adult sites - which as explained, are known to be banned. This observed behaviour of our tool may be due to the way that adult websites and businesses associate their domains together with the use of vast networks of traffic brokers, domain redirectors and link collections [39]. Based on this networking effect the web crawler may traverse content within this subject matter given the tightly linking nature of the sites - site A references site B and site B references site A, etc. However, this is important behaviour for this approach because different pages within each site may contain distinct filtered URLs. The limitation is that the crawler may get stuck in a loop within a closed network. Even so, our results contain over 1292 filtered non-adult domains for Indonesia and 528 filtered non-adult domains for Turkey.

The results for China and Iran show significant improvement over the original seed lists of filtered domains, with our number for China over 10 times greater than the input to the system and Iran over 60% higher.

#### 4.1 Top-Level-Domain Enumeration

We perform an enumeration of all publicly available top-level-domains (TLDs) that can be attributed to different domains - and

<sup>2</sup>This is the maximum depth of recursion from the seed URLs

**Table 1: DNS servers used for experiments**

	Real Servers	Fake Servers	ISP
<b>China</b>	202.46.32.29 180.76.76.76	220.181.57.217 223.96.100.100	Shenzhen Sunrise Technology Co. Ltd.
<b>Indonesia</b>	202.134.0.155 202.134.1.10	202.134.2.10 180.131.144.44	PT Telkom Divisi Multimedia
<b>Iran</b>	94.183.43.170 2.179.167.100	94.183.92.90 5.161.128.10	Aria Shatel Company Ltd
<b>Turkey</b>	195.175.39.39 195.175.39.40	195.175.30.39 195.175.30.100	Turk Telekomunikasyon Anonim Sirketi

**Table 2: Results from experimental analysis**

	Extracted URLs (HTML assets and self-linking URLs removed)	Filtered URLs	Filtered Domains (Alexa Top 1000 removed)	Filtered Domains (Alexa Top 1000 removed)
<b>China</b>	33,082,217	2,098,264	1576	1454
<b>Indonesia</b>	12,580,357	835,395	47,143	47,065
<b>Iran</b>	15,381,873	1,868,852	651	576
<b>Turkey</b>	19,250,931	913,213	39,725	39,614
<b>Totals:</b>	80,295,378	5,715,724	89,095	88,709

**Table 3: Comparison of results to CitizenLab filter lists**  
*CitizenLab figures accurate as of 1st Sept 2017*

	Filtered Domains (Alexa Top 1000 removed)	
	<i>CitizenLab</i>	Darar et al.
<b>China</b>	127	1454
<b>Indonesia</b> ( <i>Adult domains removed</i> )	124	1280
<b>Iran</b>	351	576
<b>Turkey</b> ( <i>Adult domains removed</i> )	131	513

therefore different DNS records. We use the Public Suffix List maintained by the Mozilla Foundation [17]. This list of TLDs contains all known public suffixes, common examples such as *.com* and *.org*, and less well-known instances such as *pvt.k12.ma.us*. For each filtered domain discovered in a target country, we remove the TLD and check the domain, along with any subdomains, with all suffixes in the list for filtering in that country. For Indonesia and Turkey, we run the test on the non-adult domains only for better comparison. Results of the enumeration are shown in Table 4.

Having completed this process, we find a large number of alternative TLDs for the filtered domains discovered through the traversal are also themselves filtered. During this process, we find that many of the enumerated domains found to be blocked by DNS in the target countries do not have records associated with them held by

**Table 4: Filtered domain counts after TLD enumeration**

	Filtered Domains	Of which, hosts exist
<b>China</b>	97,167	5408
<b>Indonesia</b>	1479	1543
<b>Iran</b>	5970	4527
<b>Turkey</b>	789	584

the control server. In particular, 94% of the enumerated domains found to be filtered in China received NXDOMAIN responses from the control which could therefore not resolve them. A reason for this could be that censored websites may be "retired" or move onto new domains and hosting infrastructure to evade the block. While this is a case for completely removing them from the set of results presented here, they are still explicitly filtered within the country - showing that the authorities continue to block access to them. This could be due to the stance of the censorship regime or the fact that once a site is filtered, the process for removing them from blacklists is less than trivial.

## 4.2 Categories of Filtered Domains

To gain insight into the types of content being blocked, we run a category analysis on our list of filtered domains using the WebShrinker Categories API [37]. This returns a list of categories attributed to each domain and allows us to isolate from a high level different genres of websites that are being blocked in each target country.

Figure 2 shows the breakdown of categories for each country. From this we can see that certain types of site are overwhelmingly being blocked over others. Of particular interest is the filtering of news and media, search engines and translators by China, personals and shopping by Indonesia and games and streaming media by Turkey. We also note that the proportion of proxy and filter avoidance sites blocked by China and Iran to be comparatively high too. This is in line with recent statements from the Chinese government concerning mandatory blocking of VPNs by network providers in the country [28] and a similar circumstance around the Iranian presidential election in 2013 [34].

Figure 3 shows a comparison of categories of the filtered domains between the four test countries. This is the proportion of filtered domains per category per country. From this we can infer the different types of content that are under attention by the different regimes. For example, filtering of content within the topic of weapons is even between China, Indonesia and Iran, however censorship of religious sites is more prevalent in Iran.

### 4.3 Geographical location of blocked hosts

In addition to inferring the types of content being blocked, we identify the locations of the servers hosting filtered domains in each test country. This is achieved by making a DNS query for each domain to the control server and using MaxMind GeoIP2 country database [26] to locate the resulting IP addresses by country of origin. The breakdown of the origin of hosts of filtered domains to test country is shown in Figure 4.

Unsurprisingly, we find that the largest number of servers are hosted within the United States. This is expected due to the way many content-delivery-networks maintain peers in North America and the fact that over 50% of all Internet hosts are located on this continent [5].

During the course of this investigation we observe that a disproportionate percentage of blocked domains for Turkey were hosts in the Republic of Ireland. On further analysis of the domains and IP address records we find that the country appears to block any sub-domain of *evennode.com* which is a hosting provider for NodeJS and Python web applications. The IP addresses of the blocked domains are owned by Amazon Technologies Inc. as part of their datacentres supporting Amazon Web Services. Further examination of this peculiarity was not performed, but it opens the questions as to whether certain censorship regimes will filter entire blocks of IP addresses and domains based on their hosted locations.

Other cases of interest are the irregular blocking of Dutch sites by Indonesia and Russian sites by China.

### 4.4 Backlink Analysis

For a more in-depth look into the networking effect between blocked websites, we find the number of filtered backlinks to and filtered forward links from each blocked webpage<sup>3</sup>. This allows us to see how deeply integrated each censored site is within the network of filtered content. We can look at the number of sites referencing *a given* blocked domain and also which filtered sites reference the most *other* blocked domains.

<sup>3</sup>Note that the backlinks and forward links are also themselves filtered in the given target country

To calculate these, we log every backlink we find to a filtered domain along with the filtered domains found to be linked *from* each filtered domain (forward filtered links). This results in a large graph of interconnected nodes (where each node is a filtered domain) and edges representing hyperlinks between them. From this, we can gain an insight into which domains are highly referenced within the network and which domains contain the most references to other filtered domains. Figures 5, 6, 7 and 8 show the backlinks of filtered domains for each target country.

Notable observations in Figures 6a and 8a are that the top sites that link to other filtered domains appear to be adult link collections which supports the findings in [39]. We can also see in Figures 5b and 7b that many of the linkers to filtered content are freedom of expression and independent news sites, both of which often contain political criticism.

### 4.5 Limitations

As mentioned previously, the filter check is limited by the sole use of DNS. While this reduces cause for ethical concerns, it does mean that content filtered by other means - such as IP filtering, keyword filtering or Deep Packet Inspection - will not be marked as blocked. Improvements to this check could increase the performance of the tool. Despite this, we still achieve good results.

A second limitation of this approach is the way that localised loops can form between networks of filtered content. This is a key issue with any web crawling system and often requires human interaction to break the loops - large search engines offer the ability for webmasters to provide links to new sites to improve reach. The looping behaviour we encounter can reduce the effectiveness of the system since the crawler does not have a means to connect other networks of filtered sites. Currently, this can only be altered by manipulating the seed URLs, but is not a fundamental issue with the approach. For purposes of testing and evaluation, limits were not imposed on the traversal between different domains and webpages, but a future implementation could handle looping behaviours in a similar way that search engines avoid spider traps [32].

## 5 CONCLUSIONS

This work has presented a new approach for building domain filter lists. We demonstrate the method is effective and capable at discovering censored web content in multiple different countries. Given the recursive nature of this method, we envisage that it will be a useful tool for organisations who maintain lists of blocked URLs. Furthermore, the system does not require large amounts of infrastructure to operate or special access to third-party systems or APIs. The use of DNS as a means of checking for filtering has scope to be improved, however, it allows us to test the effectiveness of these kinds of techniques, without incurring ethical issues in regards to the safety of individuals.

Through experimentation on four censorship regimes, we have discovered a large number of filtered domains that have not been previously published. This information will be of significant benefit to future studies concerning research within this field and for organisations that build circumvention tools. We aim to release this data as soon as possible.




Figure 2: Category breakdown of filtered domains for each target country




Figure 3: Category comparison of filtered domains between target countries




Figure 4: Location breakdown of hosts serving filtered domains for each target country

Our analysis of the collected data shows the relationship between backlinks of filtered webpages and hyperlinks to other filtered pages. This shows there is indeed a networking effect between different pieces of filtered content and provides a basis for future investigation. Furthermore, our analysis of the types and locations

of content being blocked gives insight into the current state of Internet censorship within these regimes.




Figure 5: Backlinks of discovered filtered domains - China




Figure 6: Backlinks of discovered filtered domains - Indonesia

## 6 FUTURE WORK

The approach described in this paper lends itself to refinement and extension. Firstly, the method of checking the filter status of URLs could be improved so it takes into account more factors than only DNS, although care will need to be taken to limit potential harm to people inside censored regions of the world. This could improve the

accuracy of the system and potentially increase the scope within which it can operate.

Secondly, the technique could be integrated with others to form a hybrid system. This may improve performance and reduce the reliance on individual networks of filtered URLs. For example, the search engine based method used by [13] would integrate well with this approach. A combined system of this type could improve both






Figure 7: Backlinks of discovered filtered domains - Iran

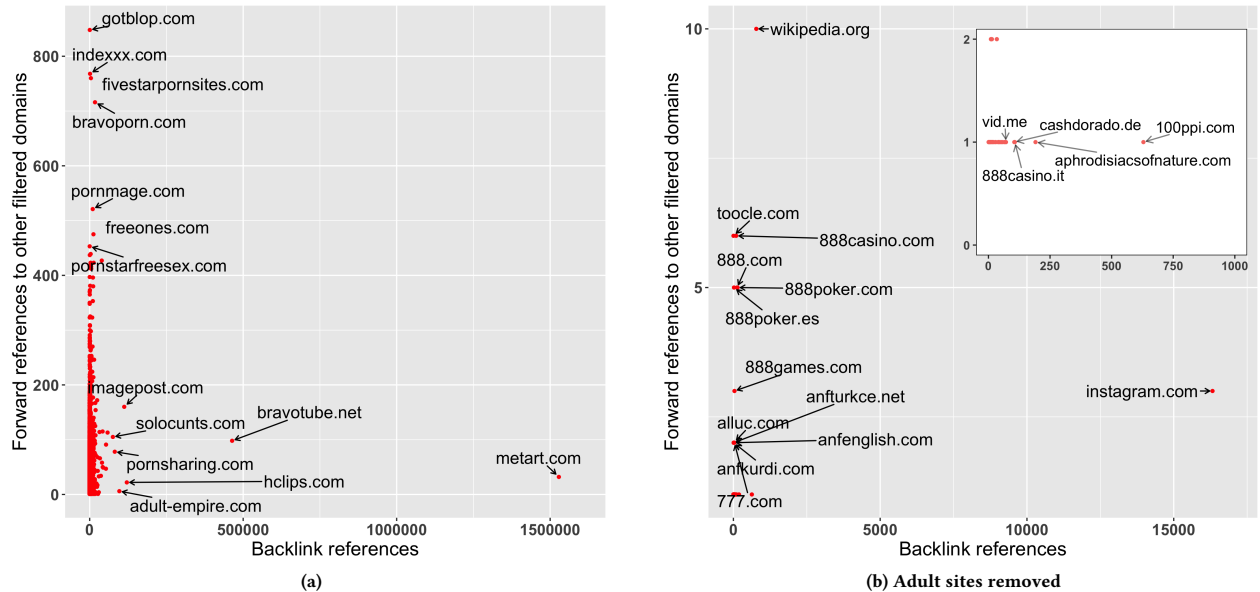


Figure 8: Backlinks of discovered filtered domains - Turkey

the breadth and depth of discovery for filtered URLs by traversing hyperlinks as well as making web searches. Furthermore, this may reduce the closed looping behaviour of solely web crawling.

Thirdly, the data collected by traversing between filtered URLs has potential for further analysis and experimentation. We have touched upon the connectivity between filtered URLs, but there is opportunity for deeper investigation into this concept.

## REFERENCES

- [1] Giuseppe Aceto. 2014. Monitoring Internet censorship: the case of UBICA. (2014).
- [2] Giuseppe Aceto, Alessio Botta, Antonio Pescapè, M Faheem Awan, Tahir Ahmad, and Saad Qaisar. 2016. Analyzing internet censorship in Pakistan. In *Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI), 2016 IEEE 2nd International Forum on*. IEEE, 1–6.
- [3] Giuseppe Aceto, Alessio Botta, Antonio Pescapè, Nick Feamster, M. Faheem Awan, Tahir Ahmad, and Saad Qaisar. 2015. Monitoring Internet Censorship with UBICA. In *Traffic Monitoring and Analysis*. Springer.

- [4] Giuseppe Aceto and Antonio Pescapè. 2015. Internet Censorship detection: A survey. *Computer Networks* (2015).
- [5] U.S. Central Intelligence Agency. Accessed Oct 2017. *Internet hosts, CIA World Factbook*. <https://www.cia.gov/library/publications/the-world-factbook/rankorder/2184rank.html>.
- [6] Mustafa Akgul and Melih Kirlidog. 2015. Internet censorship in Turkey. *Internet Policy Review* 4, 2 (2015), 1–22.
- [7] Collin Anderson. 2013. Dimming the Internet: Detecting throttling as a mechanism of censorship in Iran. *arXiv preprint arXiv:1306.4361* (2013).
- [8] Simurgh Aryan, Homa Aryan, and J Alex Halderman. 2013. Internet Censorship in Iran: A First Look.. In *FOCI*.
- [9] citizenlab.org. 2017 (accessed May, 2017). *citizenlab/test-lists*. <https://github.com/citizenlab/test-lists>.
- [10] Richard Clayton, Steven J Murdoch, and Robert NM Watson. 2006. Ignoring the great firewall of china. In *Privacy Enhancing Technologies*. Springer, 20–35.
- [11] Jedidiah R. Crandall, Masashi Crete-Nishihata, and Jeffrey Knockel. 2015. Forgive Us our SYNs: Technical and Ethical Considerations for Measuring Internet Filtering. In *Ethics in Networked Systems Research*. ACM.
- [12] Jedidiah R Crandall, Daniel Zinn, Michael Byrd, Earl T Barr, and Rich East. 2007. ConceptDoppler: a weather tracker for internet censorship.. In *ACM Conference on Computer and Communications Security*. 352–365.
- [13] Alexander Darer, Oliver Farnan, and Joss Wright. 2017. FilteredWeb: A Framework for the Automated Search-Based Discovery of Blocked URLs. In *Network Traffic Measurement and Analysis*. IFIP. [http://tma.ifip.org/wordpress/wp-content/uploads/2017/06/tma2017\\_paper32.pdf](http://tma.ifip.org/wordpress/wp-content/uploads/2017/06/tma2017_paper32.pdf)
- [14] David Dittrich, Erin Kenneally, et al. 2011. The Menlo Report: Ethical principles guiding information and communication technology research. *US Department of Homeland Security* (2011).
- [15] Oliver Farnan, Alexander Darer, and Joss Wright. 2016. Poisoning the Well: Exploring the Great Firewall’s Poisoned DNS Responses. In *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*. ACM, 95–98.
- [16] Arturo Filasto and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference.. In *FOCI*.
- [17] Mozilla Foundation. 2017 (accessed Sept, 2017). *Public Suffix List*. <https://publicsuffix.org/>.
- [18] King-wa Fu, Chung-hong Chan, and Marie Chau. 2013. Assessing censorship on microblogs in China: Discriminatory keyword analysis and the real-name registration policy. *Internet Computing, IEEE* 17, 3 (2013), 42–50.
- [19] Genevieve Gebhart, Anonymous Author, and Tadayoshi Kohno. 2017. Internet Censorship in Thailand: User Practices and Potential Threats. In *European Symposium on Security & Privacy*. IEEE. <http://homes.cs.washington.edu/~yoshi/papers/GebhartEtAl-IEEEEuroSP.pdf>
- [20] Open Net Initiative. 2011 (accessed Jan, 2018). *ONI Country Profile - Indonesia*. <http://access.opennet.net/wp-content/uploads/2011/12/accesscontested-indonesia.pdf>.
- [21] Arturo FilastÅs Khairil Yusof Tan Sze Ming Kay Yen Wong, Maria Xynou. 2017 (accessed May, 2017). *The State of Internet Censorship in Indonesia*. <https://ooni.torproject.org/post/indonesia-internet-censorship/>.
- [22] Gary King, Jennifer Pan, and Margaret E Roberts. 2013. How censorship in China allows government criticism but silences collective expression. *American Political Science Review* 107, 02 (2013), 326–343.
- [23] Jeffrey Knockel, Jedidiah R Crandall, and Jared Saia. 2011. Three Researchers, Five Conjectures: An Empirical Analysis of TOM-Skype Censorship and Surveillance.. In *FOCI*.
- [24] Marc Kührer, Thomas Hupperich, Jonas Bushart, Christian Rossow, and Thorsten Holz. 2015. Going wild: Large-scale classification of open DNS resolvers. In *Proceedings of the 2015 Internet Measurement Conference*. ACM, 355–368.
- [25] Graham Lowe, Patrick Winters, and Michael L Marcus. 2007. The great DNS wall of China. *MS, New York University* 21 (2007).
- [26] MaxMind. Accessed Oct 2017. *GeoIP2 Databases*. <https://www.maxmind.com/en/geoip2-databases>.
- [27] Zubair Nabi. 2013. The anatomy of web censorship in Pakistan. *arXiv preprint arXiv:1307.1144* (2013).
- [28] Bloomberg News. Accessed Oct 2017. *China Tells Carriers to Block Access to Personal VPNs by February*. <https://www.bloomberg.com/news/articles/2017-07-10/china-is-said-to-order-carriers-to-bar-personal-vpns-by-february>.
- [29] Paul Pearce, Roya Ensafi, Frank Li, Nick Feamster, and Vern Paxson. 2017. Augur: Internet-Wide Detection of Connectivity Disruptions. In *Symposium on Security & Privacy*. IEEE. <http://www.ieee-security.org/TC/SP2017/papers/586.pdf>
- [30] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. 2017. Global Measurement of DNS Manipulation. In *USENIX Security Symposium*. USENIX. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-pearce.pdf>
- [31] The Jakarta Post. 2017 (accessed May, 2017). *Indonesia blocks 800,000 web-sites*. <http://www.thejakartapost.com/news/2017/01/07/indonesia-blocks-800000-websites.html>.
- [32] Maria Praetzellis. Accessed Oct 2017. *Identify and avoid crawler traps*. <https://support.archive-it.org/hc/en-us/articles/208332943-Identify-and-avoid-crawler-traps->.
- [33] Andreas Sfakianakis, Elias Athanasopoulos, and Sotiris Ioannidis. 2011. CensMon: A Web Censorship Monitor. In *Free and Open Communications on the Internet*. USENIX.
- [34] Yeganeh Torbati. Accessed Oct 2017. *Iran blocks use of tool to get around Internet filter*. <https://www.reuters.com/article/us-iran-internet/iran-blocks-use-of-tool-to-get-around-internet-filter-idUSBRE9290CV20130310>.
- [35] Matthäus Wander, Christopher Boelmann, Lorenz Schwittmann, and Torben Weis. 2014. Measurement of globally visible DNS injection. *IEEE Access* 2 (2014), 526–536.
- [36] Barney Warf. 2011. Geographies of global Internet censorship. *GeoJournal* 76, 1 (2011), 1–23.
- [37] WebShrinker. Accessed 2017. *WebShrinker Categories API*. <https://www.webshrinker.com/>.
- [38] Zachary Weinberg, Mahmood Sharif, Janos Szurdi, and Nicolas Christin. 2017. Topics of Controversy: An Empirical Analysis of Web Censorship Lists. *Privacy Enhancing Technologies* 2017, 1 (2017), 42–61. <https://petsymposium.org/2017/papers/issue1/paper06-2017-1-source.pdf>
- [39] Gilbert Wondracek, Thorsten Holz, Christian Platzer, Engin Kirda, and Christopher Kruegel. 2010. Is the Internet for Porn? An Insight Into the Online Adult Industry.. In *WEIS*.
- [40] Joss Wright. 2014. Regional variation in Chinese internet filtering. *Information, Communication & Society* 17, 1 (2014), 121–141.
- [41] Joss Wright, Alexander Darer, and Oliver Farnan. 2015. Filterprints: Identifying Localised Usage Anomalies in Censorship Circumvention Tools. *arXiv preprint arXiv:1507.05819* (2015).
- [42] Joss Wright, Tulio Souza, and Ian Brown. 2011. Fine-Grained Censorship Mapping: Information Sources, Legality and Ethics. In *Free and Open Communications on the Internet*. USENIX.