

Monitoring Internet Censorship; Linguistic Connectivity within the Webgraph



Alexander Paul Darer

Linacre College

Centre for Doctoral Training in Cyber Security

—
Department of Computer Science
University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

December 2020

*For Mum & Dad
Harry, Max & Joe*

A few words of Thanks

It's an odd feeling writing this page since it is really quite impossible to thank every person who has contributed in some form to my completing this piece of work. From parents to old school teachers, friends, family, lecturers, associates & partners—to one and all who have been there, for always or but for a moment, thank you.

A huge slice of gratitude must of course go to my excellent supervisor Joss, thank you for all of the help, advice and guidance over the course of the past years. It's been a ride and I'm proud of our achievements from the papers we've published to being caught by the BBC in Thirsty Meeples (shout-out to Oli for taking one for the team and providing them with interview material).

A second slice of thanks to Andrew for all your hard work and dedication for bringing the CDT into existence here at Oxford and keeping the helm over the years. I hope you have enjoyed and grown the experience as much as I (and I'm sure everyone else) has. Maureen & David, third slice of thanks to you for keeping everyone sane and running at the same time. I appreciate both your hard work immensely as herding a gaggle of academics is no easy task, the place would not have functioned without the two of you; and, I wish you both the best for your future endeavours after the conclusion of the CDT.

To all those who've directly helped or impacted this thesis particularly Ivan & Michael for being my transfer and confirmation assessors, and Dad & Mira for proofreading in the final stages. Joss & Oli for all the hard work into the publications we wrote and presented over the past few years. A big word of thanks to Oli for presenting my paper at TMA in Dublin when I unfortunately couldn't be there. Also, thank you to all those at the CDT who made it a great place and a fountain of interest and thought. Finally, the support and administrative staff at Linacre—you were all fantastic!

Mum & Dad, an eternity of gratitude to the both of you, for all your decades of love and support. I literally wouldn't be here without you. Harry, Max & Joe, I literally could be here without you, but I'm glad I'm not. Thanks for being born my brothers. No really, thank you—I'm glad I spent my youth on driving holidays to Spain and France rather

than jet-setting in luxury to the Caribbean. But jokes aside, I am proud of you all and supremely glad of the times we've had together.

Many thanks to all of my great friends whom I met in and around Oxford over the years: Emily & Oli, Florian & Rachel, Rodrigo & Stefanye, Sam & Sarah, Danny & Lina, James, Charlie, Alex(s). Love to you all. I also couldn't write this without saying a thank you to Kelly; and, while you aren't with me now, you were a great support and indeed part of my life throughout the majority of my time in Oxford. Raman, Franco & Kenny, thank you for extending the length of time taken to submit my thesis by distracting me with all the great times, parties, trips and laughs.

Mira, you weren't a part of me until towards the end of this journey, but I am infinitely happy that you are now. Thank you for all your love, affection, support and kindness. As this chapter closes, I feel a great surge of excitement looking to the future and what amazing experiences we will have next.

Finally, for all those who I've argued with, learnt from and loved. You are the embodiment and essence of who I am and who I will become. Once again, thank you.

Abstract

This work offers a significant contribution to the ongoing endeavours in monitoring the effects of internet censorship. It can be freely accessed online by anyone who lives in a censorship free society where limitations on academic texts are not in effect. However, there are numerous places across the globe where this would be highly impractical—due to censorship mechanisms—and potentially illegal. A universal catalogue containing censored pieces of online content, online services and websites does not exist. This thesis discusses new approaches for monitoring internet censorship and the insights gained from experimental analysis with the results.

Key contributions of this work are: firstly, a method for determining if a website is censored in particular country from a remote vantage point; secondly, new approaches for constructing lists of censored domains via a recursive discovery strategy; thirdly, a first look into the relationships between newly discovered censored websites from the perspective of network topology and linguistics.

A number of experiments were conducted to evaluate newly designed frameworks for monitoring censored websites. Using a set of known censored websites from existing lists resulted in the discovery of an order of magnitude more censored material than was previously published. Furthermore, the discovery process yielded useful data and insight into how these censored websites exhibit a multitude of hard and soft connections between them. These results improve the perspicacity of analysis into how online material is censored and give new ways of identifying the motivations and intent of censorship regimes.

These new methods for monitoring internet censorship are significantly more effective than those previously in use, whilst maintaining a strong stance in regard to ethical issues with taking measurements for censorship research.

Publications Arising from this Thesis

- [1] Alexander Darer, Oliver Farnan, and Joss Wright. “*FilteredWeb: A framework for the automated search-based discovery of blocked URLs.*” Network Traffic Measurement and Analysis Conference (TMA), 2017. IEEE, 2017.
- [2] Alexander Darer, Oliver Farnan, and Joss Wright. “*Automated Discovery of Internet Censorship by Web Crawling.*” WebSci: 10th ACM Conference on Web Science, 2018. ACM, 2018.

Co-authored Publications

- [3] Oliver Farnan, Alexander Darer, and Joss Wright. “*Poisoning the well: Exploring the great firewall’s poisoned dns responses.*” Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society. ACM, 2016.
- [4] Joss Wright, Alexander Darer and Oliver Farnan. “*On Identifying Anomalies in Tor Usage with Applications in Detecting Internet Censorship.*” WebSci: 10th ACM Conference on Web Science, 2018. ACM, 2018.
- [5] Oliver Farnan, Alexander Darer, and Joss Wright. “*Analysing Censorship Circumvention with VPNs via DNS Cache Snooping.*” Proceedings of the 2019 IEEE Workshop on Traffic Measurements for Cybersecurity. IEEE, 2019.

External Media Publications

[6] Chris Baraniuk. *“Millions of censored web pages discovered in massive study”*
<https://www.newscientist.com/article/2166746-millions-of-censored-web-pages-discovered-in-massive-study/>. New Scientist, 18th April 2018.

Referencing: [2] *“Automated Discovery of Internet Censorship by Web Crawling.”*

[7] Matt Reynolds. *“As Turkey went to the polls, a global network was mapping online censorship in real-time”* <https://www.wired.co.uk/article/netblocks-internet-censorship-shutdowns-turkey-elections>. Wired, 24th June 2018.

Referencing: [4] *“On Identifying Anomalies in Tor Usage with Applications in Detecting Internet Censorship.”*

Table of contents

1	Introduction & Motivation	1
1.1	Introduction	2
1.2	Scenario	3
1.3	Research Questions	3
1.4	List of Contributions	4
1.5	Summary of Chapters	5
1.6	Language: Censorship, Filtering & Blocking	7
2	The State Of Cyber And Censorship	9
2.1	What is Cyber?	11
2.2	What is Censorship?	11
2.2.1	Types of Censorship	12
2.3	Basic Concepts	15
2.3.1	Computer Networks	15
2.3.2	Unified Resource Locators (URLs)	20
2.3.3	Encrypted Communication	22
2.3.4	[D]omain [N]ame [S]ystem Infrastructure	23
2.4	Cyber [and] Censorship	25
2.5	Methods for Internet Censorship	30
2.5.1	IP Address Filtering	30
2.5.2	Censorship of Internet Domains	31
2.5.3	Keyword Filtering	31
2.5.4	Filtering of Webpages	32

2.5.5	Blocking of Social Media	33
2.6	Methods for Circumventing Internet Censorship	34
2.7	Existing Approaches for Detecting Internet Censorship	37
2.7.1	Existing data sources	37
2.7.2	Collecting new data sources	39
2.7.3	Approaches using existing infrastructure	41
2.7.4	DNS based approaches	43
2.7.5	Combination methods	44
2.7.6	News and social media monitoring	45
2.7.7	Summarising the opportunities for further research	46
2.7.8	Catalogue of Existing Approaches	48
2.8	Ethical Issues with Research into Censorship	51
2.8.1	Risks of measuring censorship	52
2.8.2	Misuse by adversaries	53
3	A Method For Determining The Censorship Status Of Webpages And Domains 55	
3.1	Detecting Blocked Domains	56
3.1.1	Route 53: How DNS Works	57
3.1.2	Different DNS Response Types	58
3.2	DNS Hijacking	59
3.3	DNS Cache Poisoning	61
3.4	An Instrument for Filtered Domain Measurements	63
3.4.1	Finding Evidence of Censorship in DNS Responses	65
3.4.2	Comparison with existing DNS measurement approaches	67
3.4.3	Advantages	67
3.4.4	Limitations	68
3.4.5	Summary	69
4	A Framework For Modelling Characteristic Linguistic Connections Between Filtered URLs	71
4.1	Text Mining in Webpages	72
4.1.1	Descriptive Tags for Web content	73
4.1.2	Information Extraction	74
4.2	Recursive Discovery for Filtered Webpages	75
4.2.1	System Overview	76

4.2.2	Methodology	77
4.2.3	Process for Extracting Descriptive Tags	79
4.2.4	Use of Existing Search Engine Infrastructure	80
4.3	Chinese DNS Blocking Experiment	80
4.3.1	Implementation	81
4.3.2	Parameters	82
4.3.3	Results	82
4.4	Evaluation of Approach	83
4.4.1	Efficiency of Discovery	85
4.5	Further Analysis of Results	87
4.5.1	Enumeration for additional filtered suffixes	87
4.5.2	Locations of filtered hosts	89
4.5.3	Seed domain connectivity	91
4.5.4	Limitations & Issues	93
4.6	Summary	93
4.6.1	Future Improvements	94

5	Automated Filtered Resource Discovery Using Hyperlink Traversal within the Webgraph	95
5.1	The Connectedness of the Web	97
5.1.1	Links between Online Communities	98
5.2	Iterative Traversal between Filtered URLs	99
5.2.1	URL Extraction from Webpages	102
5.3	Experimental Analysis	103
5.3.1	Further Ethical Consideration	103
5.3.2	Country Specific Testing	104
5.3.3	Results	105
5.4	Evaluation of Approach	107
5.4.1	Top-Level-Domain Enumeration	107
5.4.2	Geographical location of blocked hosts	108
5.4.3	Limitations & Drawbacks	110
5.5	Summary	110
5.5.1	Future Improvements	111

6	Using Filtered URLs To Gain A Deeper Understanding Of Internet Censorship in the Webgraph	113
6.1	Filtered Webpage Category Breakdown	114
6.1.1	Co-occurring Categories	117
6.2	Descriptive Tag Distribution in URLs	122
6.2.1	Searching for Descriptive Tags in Filtered URLs	122
6.2.2	Descriptive tag word embeddings	125
6.2.3	Connection through descriptive tags	127
6.3	Content Citations	127
6.3.1	Routes of discovery for Chinese censorship	129
6.3.2	Backlinking of filtered webpages	131
6.3.3	Filtered Domain Discovery Power	136
6.4	Summary	136
6.4.1	Future Research	137
7	Conclusion	139
7.1	Resolution of Research Questions	141
7.2	Future Work	142
7.2.1	Summary of future research threads	143
7.3	Final thoughts	145
	List of References	147
	Appendices	161
A	Lists of Discovered Filtered Domains	163
A.1	China	163
A.2	Indonesia	167
A.3	Iran	171
A.4	Turkey	173

List of Figures

2.1	Global undersea cable map as of March 2019. <i>Source:</i> <i>https://www.submarinemap.com/</i>	20
2.2	Global internet connectivity by network. <i>Source:</i> <i>https://commons.wikimedia.org/</i>	21
2.3	Map of root DNS authoritative servers. <i>Source:</i> <i>https://root-servers.org/</i>	24
2.4	Censored version of The Birth of Venus by William-Adolphe Bouguereau	25
2.5	How Tor works. <i>Source:</i> <i>Electronic Frontier Foundation</i>	35
2.6	Example of bypassing censorship via email	37
3.1	Direct DNS query interference	60
3.2	Wireshark traffic capture when making queries to a DNS server in China	61
3.3	Poisoning of DNS server cache	62
3.4	Process for determining filter status for a domain in specific target countries	64
4.1	Tags derived from text in Example 4.1.1	73
4.2	High-level overview of the recursive filtered URL search	78
4.3	Filtered domain distributions	84
4.4	Performance of filtered domain discovery	86
4.5	Breakdown of filtered suffixes	88
4.6	Locations of filtered hosts	90
4.7	Seed domain connections	92
5.1	Original diagram for the Web in a proposal to CERN’s management—Information Management: A Proposal [18]	96

5.2	Original diagram for hyperlinks in a proposal to CERN's management—WorldWideWeb: Proposal for a HyperText Project [19]	97
5.3	Network connections within Wikipedia (source: Wikimedia Commons)	99
5.4	Graph of RFCs (source: Björn Höhrmann) The colour of the nodes indicate the subject cluster where the size and location depend on the number of requests and inbound-links for each document respectively.	100
5.5	High-level overview of filtered webpage traversal	101
5.6	Filtered URL traversal process: High-level methodology	101
5.7	Location breakdown of hosts serving filtered domains for each target country	109
5.8	Location breakdown of hosts serving filtered domains for each target country (normalised)	109
5.9	Location comparison of hosts serving filtered domains between each target country	109
6.1	Category counts over discovered filtered domains across all of the measurement (target) countries <i>Pornographic domains removed</i>	115
6.2	Filtered website category comparisons	116
6.3	Category correlation heatmaps per country	118
6.3	Category correlation heatmaps per country (<i>continued</i>)	119
6.3	Category correlation heatmaps per country (<i>continued</i>)	120
6.3	Category correlation heatmaps per country (<i>continued</i>)	121
6.4	Top 75 descriptive tags found in filtered webpages	123
6.4	Top 75 descriptive tags found in filtered webpages (<i>continued</i>)	124
6.5	Descriptive tag word embeddings per country	126
6.6	Graph of links between descriptive tags and filtered domains for China <i>Blue nodes: tags Green nodes: parent filtered domains Red nodes: child filtered domains</i>	128
6.7	Graph of links between filtered domains in China <i>Alexa Top 1000 removed Turquoise nodes: parent filtered domains Red nodes: child filtered domains</i>	130
6.8	Backlinks of discovered filtered domains China	132
6.9	Backlinks of discovered filtered domains (Top 1000 sites removed) China	132
6.10	Backlinks of discovered filtered domains Indonesia	133
6.11	Backlinks of discovered filtered domains (adult sites removed) Indonesia	133
6.12	Backlinks of discovered filtered domains Iran	134
6.13	Backlinks of discovered filtered domains (Top 1000 sites removed) Iran	134
6.14	Backlinks of discovered filtered domains Turkey	135
6.15	Backlinks of discovered filtered domains (adult sites removed) Turkey	135

List of Tables

2.1	Catalogue of approaches for detecting Internet censorship	50
3.1	DNS response types when resolving an IP address for a domain	58
4.1	Discovered filtered URLs	83
4.2	Comparison with alternative filtered URL lists for China	85
4.3	Filtered domain counts after suffix enumeration	89
5.1	DNS servers used for experiments	105
5.2	Results from experimental analysis	106
5.3	Comparison of results to CitizenLab filter lists <i>CitizenLab figures accurate as of 1st Sept 2017</i>	106
5.4	Filtered domain counts after TLD enumeration	108

List of Algorithms

3.1	Pseudocode for Domain Filtering Check	66
4.1	Pseudocode sample for the framework	78
4.2	Descriptive tag extraction process	79

"A talent following the ways of yesterday is not sufficient to improve the ways of today"

King Wu-ling 307 BC

CHAPTER 1

Introduction & Motivation

Filtering, blocking, suppression and restriction. Censorship of literature, information and intellectual resources—books, news, communication and invention—is common and widespread across the globe. The dissemination of works that discuss, critique, propose or display ideas and thought is both attacked and applauded. In a survey conducted by Freedom House, 45%—or 88—of all assessed countries are rated as “free” [139]. This means that there are 107+ countries in the world that are judged as “partly free” or “not free”. Over 2.7 billion—37% of the world’s population—live in nations described as “not free”. For many in western countries, this is an exorbitant figure, something that feels alien and far removed from day to day life. Yet, for many people, a life with censorship control does not instil fear or anger. On the contrary: it is often recognised, desired and applauded. Reasons that quickly spring to mind are the blocking of illicit and immoral content—such as child pornography—or the restriction of online gambling institutions within Islamic countries. These are both examples of censorship, yet many people in the West would view the latter as an imposition on free-speech, and the former an important control that benefits wider society. Censorship, it seems, is widely utilised, widely condoned and widely accepted. This thesis aims to study and analyse this activity under the *context of Cyber*, whilst also discussing the wider socio-political circumstances that greatly affect the use and state of censorship on the Internet.

1.1 Introduction

Censorship of the Internet is rife across the globe [198]; while some countries publish open doctrines on their policy, many do not. Filtering of Internet resources and content remains a covert activity with little description of technical implementation or the specific data that is blocked. This can often be for good reason where knowledge of a particular censorship system, the resources it attempts to filter, and methods for updating would compromise its effectiveness. This is a similar problem that anti-cheat software creators face when developing techniques for halting bad actors within online multiplayer games [162] such as the immensely popular League of Legends. Knowledge of the anti-cheat software and methods would allow cheaters to gain an edge in constructing exploits to gain unfair advantage—it's a constant arms race between the game and cheat developers. This is congruous to the issues that censorship regimes face and members of their population who attempt to bypass the filtering imposed by said regimes. That being said, and contrary to the field of anti-cheat, it is often the case that a regime will also own the *physical infrastructure* within their jurisdiction upon which networked systems rely. This allows for far greater and finer grain control over the way that people communicate, which services can be accessed, which protocols are allowed, permission of encryption and, ultimately, the content that can flow through the network. This makes bypassing censorship a more difficult and complex task, even more so in a clandestine manner, since accessing censored content is often an illicit activity. Nevertheless, analysis of censorship infrastructure remains a popular research topic, with members of certain countries potentially risking their livelihoods for the sake of free access to information their respective governments have deemed sensitive, harmful, opposing or otherwise immoral. However, while the blocking & filtering of content on the Internet has become more widespread over time, it remains a tough challenge for censors given the evolving nature of commerce, business, education and communications within society. Many day-to-day activities are conducted almost exclusively online or with some form of nationwide or globally networked action. In order to compete on global markets, a country *must* integrate with others in cyber space. The People's Republic of China, a nation entrenched in censorship policy and activity, has the world's second-largest economy by nominal GDP [177] and first by purchasing-power parity [143]. China is heavily integrated with Internet commerce and global markets, yet the government continues to seek control over *political* news and critique as a means of populace control. As a subject of academic investigation, this nation's censorship infrastructure has been assessed, analysed, appraised and denounced more than any other. To this day, it remains a major theme and motif in censorship research.

1.2 Scenario

Research into censorship is an important subject within the wider fields of cyber security and socio-political studies. The censorship policy and intent of national governments can influence numerous factors within the political, business and innovation strategies upon which countries develop. Decisions made towards the freedom of speech, communication and information can have far-reaching effects beyond that which may have been originally intended. Throughout history, censorship has formed major parts of authoritative strategies in order to directly control a population. However, there have also been “sincere” attempts at reducing access to content that is deemed unethical or immoral by a controlling government authority. A good example of this is the suppression of piracy-related sites that disseminate copyrighted works for free and without the rights holder’s agreement. This is unquestionably censorship that is frequently enforced through a mixture of legislative and technical means, often forced upon Internet Service Providers (ISPs). Whether or not this is itself an *“ethical” use of forced filtering* is currently being debated by different stakeholders and advocates for free speech and authors rights alike. Moreover, there is also contention between different nation states on the same issue, especially with regards to piracy—for instance, the legality of downloading pirated materials in Switzerland [169] compared to that of the UK. The use, or non-use, of censorship to control access to information is largely different among countries across the world. This thesis presents new approaches for monitoring these activities, while respecting the numerous ethical issues regarding censorship research.

1.3 Research Questions

This thesis deals with a number of topics across censorship and computer science. The broad research themes cross several fields of study within linguistics, natural language processing, censorship research, network research and the political/social sciences. Censorship is a *“human”* derived phenomenon; it is abstract and exists only within the political & social agendas of controlling actors and their respective subordinates. Where this subject lends itself to an appreciable unit of academic thought and work is where the process and decision-making aspects of censorship intersect these different fields. Technology has always been at the forefront of censorship, whether used by the censors or the agitators. The field is incredibly diverse since it touches on several different themes, making it well worth exploring with numerous research opportunities.

Below is the list of specific research questions that are dealt with in this thesis:

- To what end can we determine if Internet resources are censored/blocked whilst keeping ethical concerns minimal?
- How can we use automated censorship discovery methods to monitor Internet censorship?
- Can patterns and structure in one set of filtered content be used to find further filtered content?
- Are there homogeneous links between different filtered content that are apparent through the analysis of linguistic patterns?

1.4 List of Contributions

The projects undertaken to form this thesis have contributed significantly to the field of censorship studies and the wider cyber security research community. Two peer reviewed publications that form the main base of the work have been accepted into important conference venues. Furthermore, there have been follow-up pieces of academic research that were directly influenced by this work.

The main contributions are summarised below:

- A novel, analytical strategy for determining if a domain is censored in a remote target country or Internet Service Provider (ISP). The technique is fully automated and avoids the use of volunteer participants in regions where censorship measurements are to be taken.
- Descriptions of novel methodologies for automating the detection of censorship of specific Internet resources & content that outperform current methods by an order of magnitude. The approaches are scalable and do not require manual effort to operate after the initial set-up stage.
- A novel technique that leverages existing and commercially available search engine infrastructure as devices to find newly filtered domain names. This method incorporates the scale and sophistication of commercial search engines without the need to operate such a service.
- An analysis and breakdown of the types of censored content across a set of experimental target countries as discovered during the evaluation of the censorship detection methods. This is framed in the context of website categories and keyword extractions where we can see what topics and themes are more widely blocked in specific regions.
- A comparative analysis between different geographic regions of the nature of “*connectedness*” within Internet censorship, with analysis of how censored websites link, share and disseminate their content. We investigate the graph of censorship between different websites and how certain domains act as “hubs” of censorship activity.

- Block lists of the censored domains discovered during the study for China, Indonesia, Iran and Turkey. These lists are over 10 times longer in length than the most widely used current lists—shown in Appendix A

1.5 Summary of Chapters

Chapter 2: The State Of Cyber And Censorship

Introduction to censorship and how cyber interfaces the problem on numerous fronts. The main aim is to create useful definitions for these topics for future discussions in the thesis. There is no simple or singular way to define cyber and/or censorship, so descriptions will be taken and amalgamated from several sources on the subjects. The chapter will introduce key concepts that are used later on as well as giving a quick overview of the different types and methods of Internet censorship, and a summary of literature in this research area. Furthermore, Section 2.4 introduces a discussion on cyber censorship, how this medium alters the way censorship is administered and how the wider media reports on such activities. This provides justification for the research in this field and why we, as academics, should be interested in studying this subject.

Chapter 3: A Method For Determining The Censorship Status Of Webpages And Domains

Key to this work is the ability to find if a given webpage or internet domain is, in fact, censored in a target country. This chapter introduces existing approaches for tackling this problem, as well as the newly developed method used for the main body of research in this thesis. Reducing ethical concerns was of utmost importance during this work, therefore there is thorough discussion on appropriate methods for taking censorship measurements.

Chapter 4: A Framework For Modelling Characteristic Linguistic Connections Between Filtered URLs

The first new approach discussed for automated censorship detection is a novel method that leverages existing Internet services to perform web searches for potentially filtered content. This chapter introduces the notion that the language used across filtered sites may be common or even shared. This allows us to model the patterns in order to build search queries that are executed on large-scale search engines. The resulting URLs are then checked for censorship. Using this method, and an initial seed list of known filtered URLs, it was possible to find an order of magnitude more blocked websites in China than were published in the largest and most widely used filter list. Furthermore, this approach also gives insight into the types of content being filtered and potential keywords that are blocked in the media of the target country.

Chapter 5: Automated Filtered Resource Discovery Using Hyperlink Traversal within the Webgraph

Building on the method in Chapter 4, this chapter introduces a second novel approach that uses web crawling as a means to discover further censored sites. This technique exploits the “connectedness” of Internet sites by extracting hyperlinks from known censored web pages and then determining if they are themselves censored. Positive matches are then used to re-seed the system through a further process of scraping and checking for more filtered hyperlinks. Experiments were conducted on four target countries that engage in large-scale Internet censorship—China, Indonesia, Iran & Turkey. The results show that this method can be very effective for the task of building more in-depth URL filter lists.

Chapter 6: Using Filtered URLs To Gain A Deeper Understanding Of Internet Censorship in the Webgraph

The two approaches for discovering previously unknown or unmonitored censored web sites yield large amounts of useful and interesting information. This chapter explores this data by identifying key trends in the types of content that is censored and analysing how it was discovered. Furthermore, we can visualise the “links” between filtered sites as dependencies which show how different pieces blocked media references and is referenced by others. There is an introduction of a new potential metric of censorship that is derived from how well-cited a particular piece of filtered content is and how many other independent pieces are cited from it—*Discovery Power*.

Chapter 7: Concluding Thoughts

The thesis is concluded with a summary of the work, contributions and results. Future avenues of research, new ideas and improvements are discussed as well as the limitations and shortfalls of the approaches presented. Finally, there is a closing statement on the published papers that arose from and supported this work and the impact that they have had within the wider research community.

1.6 Language: Censorship, Filtering & Blocking

Words for describing censorship activities vary depending on research community, publication and writing style. There are occasionally misunderstandings or ambiguities surrounding the language used within this research area. Where some individuals prefer the use of particular words and others use alternatives, this can sometimes cause contentions and misinterpretations of ideas and discussion. For the purposes of this thesis, and unless otherwise stated, I will be using the above words interchangeably with regards to censorship activities.

During this thesis, I will make numerous references to those that censor—the individuals, governments, corporations or other organisations that impose censorship onto others. The language used around these entities will vary depending on the context of the discussion, for example, the phrase *“controlling government/regime”* references the political organisation that holds enough power to enforce censorship onto citizens under its jurisdiction. Similarly, *“controlling corporation”* will pertain to a business that holds power over its employees, a marketplace or other significant resource.

"Cyber = Society"

Andrew Martin, University of Oxford

CHAPTER 2

The State Of Cyber And Censorship

Censorship is control. Censorship is regulation. Censorship is power. Throughout a significant portion of human history, and with the ability to freely share expression and identity through speech, writing or action, censorship has cast a shadow. As I write this, sitting on a train in a typically “Western” country, the idea that a higher authority could effectively silence my voice or manipulate my work to their “vision” of morality is uneasy and concerning. Tailoring my thoughts, my research or the ideas I wish to convey within this thesis, consciously or unconsciously, to play in-line with the standing and opinion of a controlling power, such as a government, is not a position that I, or many others, wish to be in. Of course, throughout my life journey, from curious child to curious academic, my thoughts, mentality, opinions and ideas have been shaped by the culture I grew up in and the societal order and conventions that were apparent. These things have been manipulated in many degrees, through malicious or benign intent by numerous controlling authorities—a government, my parents, my teachers, etc. I can only become the product of my experiences and genetic make-up, there are no truths except those born in discovery or my own internal consciousness: *“Cogito, ergo sum”*—“I think, therefore I am”. To this extent, the position in which I stand principally and intellectually will have been derived almost entirely from my lived experience, and the teachings from hundreds of people, social and academic. These learnings have shaped the way I conduct myself in almost all aspects of life, and many could have been influenced for the benefit of other individuals, or society as a whole. Censorship can play an integral role in this kind of manipulation.

“Good governance: legitimate, accountable, and effective ways of obtaining and using public power and resources in the pursuit of widely-accepted social goals”

Michael Jonston—United Nations Public Administration Network [78]

The restriction of information, freedom of speech and expression, the burning of printed literature or the dismantling of printing machines, the forced, consolidated filters placed on internet resources: are these legitimate, accountable and effective methods for achieving a societal goal? Can censorship be used morally, to shape a “good character” within the people? What is “good character”? Who decides? To what extent do *the people* determine social goals? Through democratic elections? Through religious beliefs? Through revolution? These are some of the questions that many studies of censorship try to extrapolate, or even answer. Yet it seems that opinions regarding the ethical, or unethical, use of censorship as a tool of power and control depend largely on an individual’s or organisation’s fundamental principles. Freedom of expression free from interference is affirmed within Article 19 of the Universal Declaration of Human Rights [14].

“Everyone has the right to freedom of opinion and expression; this right includes freedom to hold opinions without interference and to seek, receive and impart information and ideas through any media and regardless of frontiers.”

Article 19—The Universal Declaration of Human Rights

This document has been adopted by, and influenced, many national constitutions. However, it is not legally binding. It is open to interpretation. It places no emphasis on the *origin* of opinion, thoughts or ideas. If media outlets within a nation only report a case from a single, biased position, opinions will form around this. If multiple positions are discussed, then it stands to reason that consumers will also form multiple opinions. Censorship of published media and content is inextricably linked to the manipulation of public opinion, to achieve societal goals—for control or social morality. The question is: who decides the standard of social morality? Morality is deeply linked to, and derived from, principles of culture, religion and philosophy. If a particular religion deems the consumption of pornographic literature immoral, then is it appropriate to block access to this content in a society that builds its principles of morality around that religion? There are no easy answers to these questions, and opinions surrounding them appear to highly influenced by cultural norms. What is certain however, is that humanity has arrived into the age of information. As the Cyber era of society blooms, the way humans communicate is changing with the speed of information dissemination. The scale of shared knowledge, ideas and expression has been

increasing exponentially. The overlap among different cultures and societies is growing; multi-national businesses and organisations span the globe spreading ideologies through commerce and charity. Electronic dispersal of information and finance is a common grounding among nations. The side effects of this: the electronic dispersal of culture, skills, opinion and expression. Society has *become* cyber and *cyber is* society.

2.1 What is Cyber?

cyber: adjective. *relating to or characteristic of the culture of computers, information technology, and virtual reality.*

cyber: combining form. *relating to electronic communication networks and virtual reality.*

Cyber is a term used across multiple subjects. It relates to and characterises the culture of computing and information technology. Commonly, cyber is combined with other words as a prefix, for example cyberpunk or cybersecurity. These combinatorial forms appear to have simple meanings; however, they are often understated. Cyber as a concept is wide and far-reaching. The *culture of computing* is as much ingrained into modern society as language and long-standing traditions and rituals. Many societies are now completely dependent on computers, and more importantly, computer networking. The virtual world is a trove of information and consumption. It is estimated that 235 billion emails are sent and received each day [65]; this is an extraordinary figure and a huge amount of transferred information. However, this barely touches the surface of the cyber-ecosystem. From financial markets to personal communications, weather forecasting to entertainment, privacy and surveillance—cyber touches them all. The world is interconnected by mass computer networks, linking continents with the almost instant transmission of large scale data. Paper-based organisation systems are being phased out and replaced with efficient cyber-based alternatives. Cyber is integral to modern society.

2.2 What is Censorship?

censorship: noun. *the suppression or prohibition of any parts of books, films, news, etc. that are considered obscene, politically unacceptable, or a threat to security.*

According to the Beacon for Freedom of Expression [127], the term “censor” originated in 443 BC in Rome. The Censor, a position held by an elected officer, was charged with maintaining the census of the population as well as upholding public morality, which involved enforcing the moral standards decided by the Centuriate Assembly (a group of elected individuals who held all political

power in Rome). Since then, censorship has changed drastically in how it is approached, applied and perceived. However, one aspect which certainly hasn't changed is the way censorship is used to retain power over individuals within a population. Defining censorship is tricky because of the way it integrates with society; most nations around the world exhibit a form of censorship for one reason or another. Written below are a number of different definitions for modern censorship from various sources.

- *“The control of the information and ideas circulated within a society.”*
—Global Internet Liberty Campaign
- *“The suppression of speech, public communication or other information which may be considered objectionable, harmful, sensitive, politically incorrect or inconvenient as determined by governments, media outlets, authorities or other groups or institutions.”*
—Merriam-Webster Dictionary
- *“Censorship, the suppression of words, images, or ideas that are "offensive", happens whenever some people succeed in imposing their personal political or moral values on others.”*
—American Civil Liberties Union

2.2.1 Types of Censorship

Censorship manifests in several different forms depending on the regime or organisation, the aims and the intentions. Motivations behind censoring speech and expression will often depend on what ideology or system the censor wishes to influence or dominate. The bottom line is that censorship is a form of control and power. It is not always narcissistic in interest—some individuals believe that censorship can be moral and right, but usually only towards directives that *they* believe in. A good example of this is the censorship of certain types of content born of illicit activities, such as child pornography or racist hate speech. The majority of people—maybe even the perpetrators of these crimes—will deem this content immoral and, thus, the censorship of it is moral. It is still blocking, filtering, suppression and restriction; the reduction of freedom of expression and ideas, yet also widely accepted and tolerated and, sometimes, even praised or commended. Other types of censorship are generally not thought of in this way within many Western countries. Suppression of political, social, cultural ideas and critique is mostly seen as excessive control and power. It is generally not accepted or tolerated. However, in other societies, particularly those with strongly integrated cultures of organised religion, the restriction of certain communications and information is condoned by large portions of the populations. Is this due to the imposition of indoctrination

and propaganda, the control of information and ideas at an early stage of an individual's life? Or is this something that people just "felt" was right? It was their God-given freedom that enabled them to realise their faith in a higher order, then use and impose its teachings on others. From a logical and methodical standpoint, it is difficult to rationalise this. If you subscribe to the "blank-slate" view [90]—the idea that almost all of an individual's behavioural traits, their personality, their beliefs and interests are developed from environmental factors as they grow—then a controlled state of information seems a more appropriate proposition. And, even if you stand elsewhere within the *nature vs nurture* debate, it remains difficult to find a rational explanation for the immorality of heresy or critical depictions of religious teachings without prior indoctrination. Furthermore, when organised religion is ingrained within a society, its laws and customs, we find increasing suppression of alternative views and opinions throughout history.

Political censorship is the suppression of objective information and critique directed against a government or political party. It is often thought of as a systematic restriction of ideas and communications that oppose those of the controlling government. This type of censorship can be applied in numerous ways:

- Exercised control over journalists or news outlets
- Concealment or falsification of information
- Litigation practices against political dissidents

These activities are commonly paired with propaganda initiatives whereby incorrect or fake information is spread while contrary views are blocked. The use of political censorship almost always arises from a government with power over a set of armed forces that can be utilised to impose fear or enforcement of the controlling party's agenda.

Religious censorship is the suppression of freedom of speech within a clique or society that does not align to a controlling religious organisation's views or a set of religious teachings. This type of censorship has been widely practised by many religious societies throughout history and is still enacted today. Much of Western Europe was under restriction of ideas that opposed Catholic views and those of the Church throughout the Middle Ages and Renaissance, including the theory of Biological Evolution and Heliocentrism. Many Middle-Eastern countries that operate under Sharia law enforce restrictions over content on the Internet such as social media and critique of Islam. These are consequences of actions that directly arose from opposition of the respective religious doctrines and teachings.

Corporate censorship is when a business or corporation suppresses speech and expression about itself, its actions, its products/services or a competitor. This occurs when employees, associated people or other corporations are threatened with monetary sanctions, loss of business or unemployment if they engage in critique or opposition against the controlling corporation. It also includes restriction of copyrighted or protected works due to publisher/owner refusal to approve their release.

Corporate censorship is most commonplace in the art, news and entertainment industries. There are many instances where music, television and other performing arts have been restricted to keep in line with the controlling party's agendas and motivations. News outlets are particularly culpable of this activity since they often have systematic biases towards certain political objectives or sides, or relationships with third-party corporations such as advertisers and partners. This is often exercised through patterns of self-censorship where a particular journalist or media organisation will simply "not report" a situation or story because it is "impossible to document everything" [34], as remarked by Croteau and Hoynes.

Internet censorship enacted by corporations is increasing too. Much of the networking and hosting infrastructure that makes up, or provides access to, the Internet is owned and operated by private businesses. There have been numerous cases where corporations will block or restrict certain resources to consumers or producers if they deem they will incur reputational or monetary damages due to providing availability for them.

Self-censorship is cross-descriptive of the other types of censorship. It occurs when a producer or publisher decides to remove, reduce or retain a set of works or information due to self-interests, fear of reprisal or lack of interest or willingness to release a resource. Self-censorship is often practised without a conspicuous or obvious pressure from a third-party controlling organisation—though it is commonly due to the preferences or prejudice of such a third-party. Many content producers within totalitarian regimes will restrict their own works in order not to conflict with the opinions or ideologies of their governments. It may not be explicitly illegal to produce such content, however they may receive criticism or undesired repercussions from others within the community, such as customers or advertisers. This has led to many Western pieces of artistic content or news to be subjugated or altered to fit into a more *harmonised* view as dictated by a repressive government known as "rivercrabbing" in China [108].

Self-censorship can become apparent when any of the above three censorship types are enforced upon a population or organisation. It also materialises under authoritarian organisations or regimes, even if censorship isn't officially appropriated, but strict rulings with harsh

punishments are commonplace. Hence it never usually occurs without an external pressure or cultural biases and predispositions of some form.

2.3 Basic Concepts

In this section, we will explore the underlying concepts that are used within this thesis. Internet censorship is a field that crosses numerous boundaries between the physical and virtual world, where complex socio-political interactions are augmented through the use of technology. At its heart, censorship will always be a “*human*” issue—the use of content/ideology moderation and filtering is for the benefit of an individual or organisation. However, as networked communications have become universally ubiquitous for a large portion of the world’s population, the platform on which censorship is performed has shifted to the web. It is, therefore, important to have a grounding of computing and networking notions in order to properly examine the state of censorship within the online world.

2.3.1 Computer Networks

Computer networking, as we know it today, has been born from a combination of academic, military and commercial interest. It has grown organically through different mediums of research, often crossing borders between different stakeholder groups, such as governments, corporations and universities. The early days of computer networks were limited to simple data links between localised computing environments. These then grew to wider area connections among different locations across nations and eventually the globe. Many of the fundamental networking principles and protocols were developed during the growth of these interconnected networks, often to address issues that have become synonymous with distributed systems. The different procedures and conventions for exchanging binary data among competing network implementations have steadily been merged or become obsolete over time. Many of the implementations within the suite of networking protocols we use en masse today have been realised through cooperation among operators during these mergers.

2.3.1.1 Open Systems Interconnection (OSI) Layers

The OSI model is a standard for communications across telephone or computer networks. It was formed in 1983 through a merger of existing standard documents describing how telecommunication networks should operate. The basis of the model was formed from experiences learned from the construction of earlier networks, such as ARPANET (USA) and CYCLADES (France). These competing ideologies for computer networks were realised through

different academic and military led research projects. Many underlying concepts within the OSI model—such as packet-switching—were created on these networks. The original directive contained seven layers:

1. **Physical** Transmission of raw binary data over a physical system
2. **Data link** Transmission of data frames between two nodes over a physical layer
3. **Network** Multi-node network structure that allows control of traffic to routable addresses
4. **Transport** Transmission of larger segments of data between nodes on a network
5. **Session** Maintaining longer communication sessions between nodes on a network to allow back and forth data transmission
6. **Presentation** Translation from network encoding to application data
7. **Application** High level abstractions that allow for complex interaction over a network, such as file sharing

2.3.1.2 Internet Protocol (IP)

IP is a cornerstone communications protocol that underpins the majority of internet communications. The initial version of IP was created and used within ARPANET in US. IP sits within layer 3 of the OSI model and provides routing and transmission capabilities for datagrams known as IP packets, using addresses that are locatable on a network. The first major version of IP was v4—known as IPv4—which is the major networking body of the Internet. The next version, IPv6, is currently being adopted by most countries and major internet service providers. The IPv4 address range is limited to 4 bytes or 32 bits which sets the maximum number address at 4,294,967,296. The format for an IPv4 address is as follows:

<0-255> . <0-255> . <0-255> . <0-255>

127 . 0 . 0 . 1

Examples: 129 . 67 . 190 . 161

156 . 67 . 241 . 46

Four numbers between 0 and 255 separated by periods. Addresses are usually allocated in blocks to an ISP or nation by the Internet Assigned Number Authority (IANA) and the Regional Internet Registries (RIRs) using Classless Inter-Domain Routing (CIDR) which is a method for describing and allocating blocks of IP addresses. The blocks that are issued to different ISPs are publicly known. It is, therefore, trivial to identify what entity hosts a particular address. For censorship studies this is an

important feature since we can use this information to identify where an IP address is supposed to be located geographically and virtually. It is often the case that this is manipulated during censorship.

We must also be aware of private IP address blocks, these are ranges of addresses which are reserved for private, *non-internet* connected networks. They are used within internal networking systems where two machines need to communicate over a Local Area Network (LAN). Internet facing services *should never* be provisioned with a private IP on their network interfaces, in fact most computers and routers are designed to they will only ever send data to a private IP address over a LAN. The reserved blocks are:

10.0.0.0 - 10.255.255.255 CIDR: 10.0.0.0/8 (255.0.0.0)
172.16.0.0 - 172.31.255.255 CIDR: 172.16.0.0/12 (255.240.0.0)
192.168.0.0 - 192.168.255.255 CIDR: 192.168.0.0/16 (255.255.0.0)

IP Packets Data sent over an IP network is organised into small chunks called packets. These differ per version, but generally contain an IP header and payload, which depends on the higher-level protocol. The header will consist of numerous fields describing the packet and disclosing the source address and destination address. Many features of IP packets help with the transmission of data within the link layer in an attempt to improve the reliability and speed of sending and receiving traffic, such as fragmentation and reassembly. In higher-level layers, such as transport, these packets can be organised and manipulated for dependable transmission of data.

Networking Hardware Routers are specific pieces of networking hardware that will forward IP packets to external computers or other routers. They are instrumental in forming large networks and are key to directing traffic from one network to another. The term “networking device” is fairly high-level and is often used to describe many different types of network equipment such as:

- Gateways
- Routers
- Switches
- Repeaters
- Bridges
- Hubs

These devices typically operate at different OSI layers within a network. For the purposes of this, we are mainly interested in switches and routers, which usually work on layers 2 and 3 respectively.

2.3.1.3 Other Key Protocols

Transmission Control Protocol (TCP) TCP is one of the main protocols in use for data exchange over IP based wide area networks. It is designed to provide a reliable method for

transmitting data between nodes, with acknowledgment of delivery and error checking built in. Packets of data within TCP are sent with sequence numbers so the order can be corrected by the recipient. They also contain checksum values so that the integrity of the packet can be confirmed. Furthermore, the recipient will respond to the sender with acknowledgment responses for each packet received. This allows for lost packets to be resent. TCP is the most widely used protocol for many internet communications including the world-wide web (HTTP), file/resource sharing and email.

User Datagram Protocol (UDP) UDP is used for communications that don't require guarantees on data delivery, ordering or duplication. However, data integrity is offered in a similar way to TCP, using checksums. The protocol does not protect the data link from any inadequacies of the underlying network. As such, packets may be lost or malformed on delivery. However, in doing so, the overheads of TCP are not incurred. UDP is often used, therefore, where speed of data transmission is of utmost importance. It is used for applications such as DNS requests, video/audio streaming, voice over IP (VoIP), etc.

Border Gateway Protocol (BGP) BGP is part of the IP suite which is used to exchange routing information between Autonomous Systems (AS). An AS is a collection of routing prefixes for network operators that are under control of a singular entity or domain. These prefixes are assigned as AS numbers (ASNs) that denote a particular network, or set of networks, so they can be uniquely identified on the Internet. These numbers are usually allocated to large network operators such as ISPs. BGP is a protocol that allows for a common routing policy among different ASs. It is a path vector protocol that makes routing decisions to allow internet traffic to flow between different networks in order to form the wider internet. This works by ISPs announcing their ASNs with the associated IP addresses within those networks; BGP uses a set of rules to then determine where a particular IP packet should be routed to.

Hyper Text Transfer Protocol (HTTP/HTTPS) HTTP is a protocol for transferring information over distributed systems and networks. It was initially developed at CERN by Tim Berners-Lee in 1989. The key elements of HTTP are the concepts of documents and hyperlinks. A document is any accessible piece of data that can be accessed via the protocol using a hyperlink reference it. Hyperlinks are essentially addresses that allow for the location of documents on the Internet or an intranet. This allows authors to produce document that can *link* to and reference other documents—thereby allowing for the creation of a network of documents accessible via the Internet, called the World Wide Web (WWW). Hyperlinks on the Web are formatted as Uniform

Resource Locators (URLs) which are references to resources that specify a location and retrieval method. HTTP provides the necessary data transmission protocol to allow for the request of resources via URLs. HTTPS is HTTP transmitted over a secure connection using TLS or SSL to encrypt the data in transmission.

2.3.1.4 The Webgraph

The Webgraph is a description of World Wide Web in terms of a collection of directed connections—*edges*—between documents—*nodes*. Since the Webgraph is a simple network graph, computations can be applied to it like any other graph. This allows for calculations such as *degree distribution*, *PageRank* and identification of co-citation. Furthermore, community discovery and detection algorithms can be executed to find collections of linked topics and authors/publishers. The network topology of the Webgraph is simply webpages and hyperlinks, however, further relationships can be identified through the metadata or content of such webpages. An example of this is co-occurring word or pieces of language across different documents.

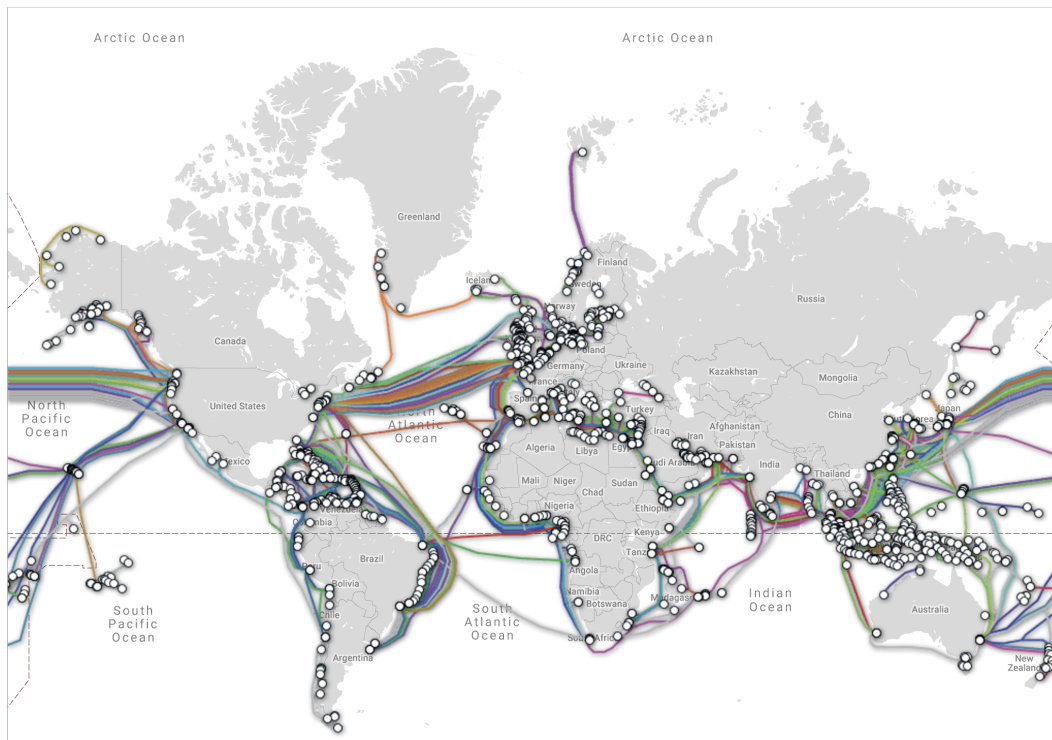
2.3.1.5 Global Connectivity Map

The majority of internet communications between large ISPs and nations are transmitted using fibre optic cables. This medium allows for huge bandwidth rates, far larger than traditional copper cabling and they are more resilient against interception and tapping. Over the course of several decades, undersea fibre cables have been laid to connect continents across oceans. The current network state provides a distributed and redundant system upon which most internet data transmission will occur. Figure 2.1 shows the state of undersea connectivity as of March 2019.

Furthermore, the inter-connectivity between major ISPs and networks is incredibly vast and complex. There are huge numbers of different possible routes from one machine over different pathways. Figure 2.2 attempts to visualise this based on data collected about different internet networks in 2015.

Figure 2.1: Global undersea cable map as of March 2019.

Source: <https://www.submarinecablemap.com/>



2.3.2 Unified Resource Locators (URLs)

A URL—or web address—is a reference to a web resource that provides a network location and method for obtaining the content. URLs can point to any kind of file, directory or identifiable thing. They are widely used for internet communications and on websites as the means to link between and share webpages. The origin of the URL was the RFC 1738 document [151], written by Tim Berners-Lee. The structure for URLs follows a strict format:

URL = scheme:[//authority]path[?query][#fragment]

Examples:

`http://www.ox.ac.uk`


`https://en.wikipedia.org/wiki/Oxford`

`https://en.wikipedia.org/wiki/Oxford#University_of_Oxford`

`ftp://mirror.ox.ac.uk/sites/ctan.org/systems/texlive/`

Figure 2.2: Global internet connectivity by network.

Source: <https://commons.wikimedia.org/>



2.3.3 Encrypted Communication

The encryption of information is the act of manipulating data so that it cannot be read or understood by unauthorised parties. It is often used to prevent eavesdropping over lines of communication. Encryption has a long history where numerous techniques for encoding messages and deciphering encrypted data have been developed. Many of these methods were born for use in military environments or where important persons of interest needed to transfer information to remote locations without allowing third-parties to access it, such as governmental communications. The goal of encryption is to convert a plaintext message into ciphertext using a scheme that then allows an authorised party to decrypt the message back to plaintext. Encryption schemes usually use keys that allow this conversion between plaintext and ciphertext and ciphertext to plaintext.

Modern encryption is the product of decades of research and development into cryptography (securing communication with encryption) and cryptanalysis (subverting encryption). Two of the most important discoveries for encrypted communications today have been in key agreement and public-key cryptography.

Key agreement: the act of securely exchanging cryptographic keys over a public channel that then allows for encrypted communication. The Diffie–Hellman key exchange was one of the first methods for securely agreeing on keys over a potentially insecure line of communication (such as the Internet). The seminal paper *New Directions in Cryptography* [44], published in 1976, describes a technique which allows two parties to establish a shared key for *symmetric encryption*—where the same key is used for encryption and decryption.

Public-key cryptography: the use of two keys in an encryption system, one which is used for encryption and one for decryption. These keys are usually referred to as the Public-key (known to all) and the Private-key (known only to the owner of the Public-key). Public-key cryptography allows for *asymmetric encryption*—the use of different keys for encryption and decryption. Furthermore, it gives the ability to *digitally sign* messages which can be used to prove a message came from a particular sender. This style of encryption was also first publicly presented within *New Directions in Cryptography*. However, research declassified by the British government in 1997 showed that public-key cryptography was known to them in the early 1970's through discoveries by James H. Ellis and Clifford Cocks, two cryptographers working for GCHQ.

The widespread use of public-key encryption was not fully realised until the world-wide web was becoming more ubiquitous. With the explosion of internet communications becoming global, the need for a scalable method for securing data transfer became increasingly important. Transport

Layer Security (TLS) and Secure Sockets Layer (SSL) are security standards that provide protocols for establishing secure, encrypted channels over computer networks. These are cornerstones for remote communications today and are widely used to allow for private data exchange between parties over the Internet. A major part of these is the use of Trusted Third Parties (TTPs) which aim to reduce the dispersal of fraudulent content online. A TTP will create identity certificates using digital signatures that can be used to determine if a message has originated from the owner of the certificate. This allows one party to effectively authenticate another on the Internet. Commonly, companies, organisations or website owners will register their respective domain names with a TTP that allows them to prove to users they are the owners of the given domain. For instance, when a user accesses their Bank's website, they need to be sure that the remote host they communicate with is indeed operated by the Bank, TTPs offer a method to do this using identity certificates.

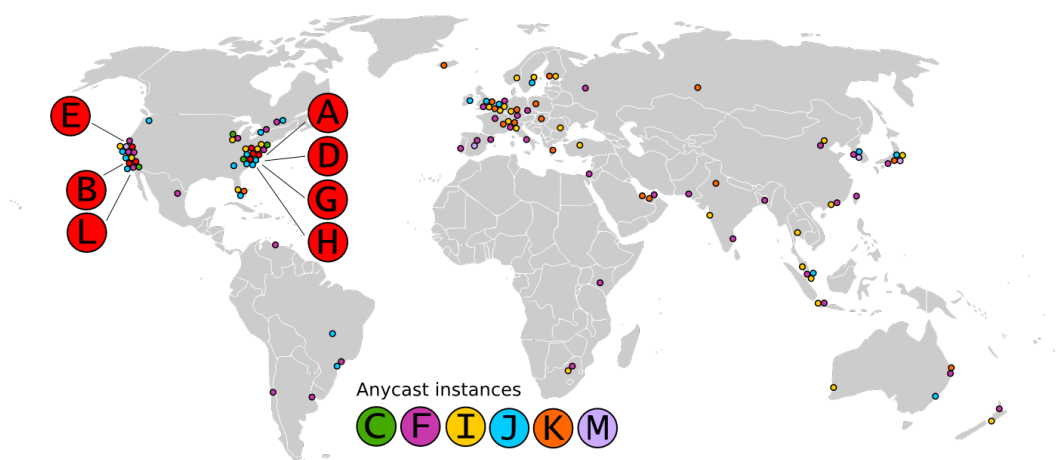
2.3.4 [D]omain [N]ame [S]ystem Infrastructure

The Domain Name System (DNS) is a major part of almost all communication over the majority of networks, academic, corporate or governmental. It provides a method for assigning names to hosts on an IP network. Most machines connected to IP networks will be allocated one or more IP addresses in order for other nodes to route data to them. Domain names are used to abstract from the specific IP addresses and instead allows nodes to perform lookups for the host they wish to communicate with. DNS is a hierarchical system which relies on authoritative naming servers that contain lookup tables translating domain names to IP address(es). When a machine wants to communicate with the host that is assigned the name *node69.corp*, it will first perform a DNS lookup by sending a query to a known DNS server for that name. The DNS server will respond with an IP address of that node or an error if the name does not exist. Because DNS is hierarchical, the naming servers can be sub-divided to provide redundancy and fault tolerance. However, it is usually the case that the top-most server is the ultimate authority. This means that if an individual, lower-level server does not have a record for a domain name, it can query a different DNS server higher in the hierarchy.

For wider Internet communications, there are several root authoritative DNS servers at different locations around the world. These act as the top-level naming servers of the root zones of the Internet. They will respond to requests with other, lower-level authoritative servers for top-level domains (such as .com, .co.uk, .org, .net). As such, the root servers are extremely critical for the Internet to operate as they will be the initial step for finding the IP address assigned to any allocated domain name. A map showing the locations of these root servers is shown in Figure 2.3.

Figure 2.3: Map of root DNS authoritative servers.

Source: <https://root-servers.org/>



DNS is a relatively old protocol with the first specifications published in 1983 [152][153]. These documents detail the original standard for DNS—naming, serving and querying conventions. Many DNS services operate over UDP on port 53. While secure implementations of DNS exist, the majority of DNS requests and responses are transmitted unencrypted. This means that most DNS queries can be intercepted, read, manipulated or dropped. A common use of such interference is where ISPs will redirect DNS requests from the intended destination to their own DNS infrastructure, potentially for monitoring or commercial interest. This can be done relatively trivially since the UDP packets will pass through their routers, which can then modify the destination address within the headers or act as a man-in-the-middle and complete the request itself by querying an alternative DNS service. DNS security and the ways in which censors can leverage DNS as part of filtering infrastructure is further discussed in Section 3.1.


2.4 Cyber [and] Censorship

Cyber has not altered the reasons why censorship is administered or applied to a population—it has purely provided a new medium for the process. The way that content is shared and consumed has drastically changed over the past couple of decades. There are enormous interconnected social and communication networks that previously did not have the capability or speed of execution that is provided when the Internet is leveraged for such activities. The speed of dissemination of data and information is outrageously fast, the idea of “viral” content is new and exciting. In 2001, a video was said to be viral if shared between one thousand or so individuals. For the same title to be applied in 2019 it requires millions. Email, file sharing, instant messaging, news sharing, video sharing; the activities that were once confined to large, non-portable desktop computers are now mostly

performed from pocket devices with access to incredibly fast wireless data connections to the wider internet; smartphones have been a key mover in the paradigm shift that has shaped how many people now interact amongst one another. Constant, always-on access to new information, false information, major events and virtual communities. The way we communicate has drastically changed. Furthermore, the *types* of data that are censored have also altered. Textual information is a common example; however, image and video data is also now widely filtered, or partly filtered too. Interestingly the way in which such data is censored is highly dependent on the desired outcome. For instance, an image may not be politically harmful but may contain nudity which, in some jurisdictions, is considered to be sensitive. This can lead to the situation where rules and legislation require *only* the offending parts to be filtered, as shown in Figure 2.4.

Censorship has also evolved over time to compensate for the underlying change of medium. A key element of any censorship infrastructure is exerting control over the information distribution channels. Today the major channels appear to be internet service providers (ISPs) and large corporations that deliver computing hardware and software. With control over one or both of these assets, a censor can enforce content moderation on a large scale. This is the same

Figure 2.4: Censored version of The Birth of Venus by William-Adolphe Bouguereau



methodology that was used in the 15th century with the invention of the printing press [94], which was wielded by the Catholic Church to spread its ideological mission. While it was used to great effect in this effort, it also provided heretics an opportunity to enhance their voice and aided Protestant reformation. Since most presses were located in educational institutions, such as universities—which the Catholic Church widely controlled—there was a ban on all printing of works without prior consent of the Church. Domination over the medium upon which information is spread allows control of the flow of ideas and communications. The idea that censorship has fundamentally changed is dubious, reasons for censoring a group are numerous and complex. It is about exerting control to gain an objective. Whether that objective is deemed ethical or moral is a subject of debate. There are many forms of censorship and it is exercised for many reasons. What has changed however, is the *speed and scale of censorship*. While one may wish to censor a postal service which handles a million letters per hour, the human effort required for such an activity is large, and the organisation of such a system would incur numerous overheads. Since the birth of the Internet, a large-scale ISP may handle network traffic for several millions of people concurrently. Much of the infrastructure of an ISP is also more centralised, where a single exchange may link to hundreds of thousands of individuals. Restricting access to web resources at an exchange is relatively trivial on modern networking equipment. There are functions *built-in* that will do this, for example, blocking access to a range of IP addresses. If such an action is performed, one could effectively cut access to an internet service widely and instantly. This represents a huge shift in how censorship can be administered, the capabilities of powerful censors, and the way they can achieve their objectives.

The governments of nations are major players within censorship activities. Purveyors for mass filtering of content they deem inappropriate, inconvenient, harmful or sensitive have numerous reasons for doing so. The goal could be for exercised control and regime survival, security and obscenity, or both. Most countries engage in some form of censorship activity, for example the blocking of file sharing websites that provide access to copyrighted materials. However, it is felt by many the more dangerous forms of censorship occur when freedom of speech and expression is suppressed. China is a commonly noted example of this kind of undertaking. News reports of censorship activity in China are common, and the reasoning wide. Xi Jinping's vision of the Internet is that nations should be able to choose how cyber-space is monitored and controlled, as set out in his speech to the World Internet Conference [136]. From politics to civil society and historical events, the term the Chinese government uses in its directives is “managed”, topics should be “managed”. A media report in Sept 2018 states that a new set of information to be controlled is its economic performance where new data may show a potential growth stagnation

[168]. The stated reason is to prevent *chaos and panic* amongst the society which could harm the government's position. China's State Administration of Press, Publication, Radio, Film and Television—the organisation that oversees all media activity in the country can enact any censorship ruling it wishes. Social media blocking in China has been widespread since its inception, whereas censorship of video streaming was relatively limited due to the complexities of identifying sensitive content in this form. However, the largest social media site in China, Weibo has been ordered to remove live streaming functionality [179], due to the increased “negative speech” and breaches of “audiovisual guidelines”. Two further sites—iFeng and ACFUN—were also affected by this order.

The Chinese censorship policy is a very interesting field of study in of itself. While China's constitution gives its citizens freedom of speech and freedom of the press [149], the Chinese regulatory bodies employ the power to crack down on published materials that it deems inappropriate or endanger national security by revealing state secrets. Definitions within Chinese media regulations are ambiguous and vague; this therefore allows the government to enforce censorship without breaching its official constitution. The very meaning of “*state secret*” is manipulated into different contexts which provides opportunity to build censorship policy. From preventing dissidents promoting alternate views on social media to questionable applications that remove content from mobile devices, reportedly targeted towards specific religious groups [158]. Of course, the Chinese government welcomes cooperation from companies and other nations, who are willing to alter business activities in order to fit in line with their aims and goals. For companies based within China this has become a very “*normal*” procedure, however, we are seeing more and more efforts from Western businesses playing in this field. Facebook's attempts to break into the Chinese market are well known, and <https://www.facebook.com> has been blocked within China for some time, yet they are attempting to alter their offering to provide a version of the service for Chinese citizens [147]—essentially a censorship tool for Facebook content. Google's efforts to provide services within China have been longstanding, where the site has been censored on and off since its inception. Google withdrew from the market in 2010 [130] amid difficulties in building a platform that conformed to the Chinese governments objectives. However, in 2018, the company is said to have prepared a censored version of the search engine, along with altered mobile applications, that will “blacklist sensitive queries” and “restrict access to content that Xi Jinping's Communist Party regime deems unfavourable” [164]. Many researchers and human rights activists will feel this move has serious implications for freedom of speech and expression, however for the company itself, the Chinese market is unquestionably lucrative, so it is not of much surprise that it wishes to increase business in this region.

The media restrictions within China and the large potential for companies to do business has given rise to interesting cultural phenomena. A number of Western corporations, particularly technology based companies, now dominate the corporation rich list by market capitalisation [125]. Many of these businesses have strong trading links with the Chinese market and manufacturers, for instance, Apple Inc. builds and assembles almost all of its products within China for the domestic and foreign sales. Apple enjoys a profitable situation in this case with the 730 million or so Chinese citizens who live within metropolitan areas with fast Internet speeds and access to e-commerce platforms. They work closely with the Chinese government to produce a product that “fits” the guidelines with software distribution channels that are restricted—for instance, most third-party VPN applications are not available for download within the Chinese App Store [137]. Apple’s CEO, Tim Cook, has been reported to have “celebrated” China’s vision of an open Internet [167], which for many human rights advocates is seen as a defence of the Chinese censorship state, maybe even lending it some credibility. These circumstances where businesses appear to create two versions of a product or service aren’t limited to Western companies either. WeChat—the most popular instant messaging application for Chinese language users—is reported to be using a censorship policy for users located in China and one for those located internationally [144]. It seems that many companies are now comfortable in building open and regulated offerings, where submitting to certain censorship regimes is becoming more commonplace.

China is not the only country that censors its citizens, far from it. Though the Chinese censorship policy and infrastructure is probably the most widely studied, many other nations engage in filtering activities as matter of course. During mass protests, rallies or period of political unrest, the blocking of content or services deemed sensitive by a controlling power becomes fairly commonplace. Countries designated as *one-party states*, such as Vietnam, often implement strict media regulation and usually build strong surveillance capabilities for monitoring public “views” on the Internet [161]. Censorship/surveillance evasion or circumvention tools are often blocked under many authoritarian regimes too. Russia is known to have blocked Telegram, an instant messaging application that provides strong end-to-end encryption. As reported widely in the media, the Russian government ordered the operators of Telegram to hand over encryption keys that would allow them to decrypt user content [124][172]. Blocking internet access to this service also had the side-effect of restricting access to numerous sites operated by Google and those hosted by Amazon Web Services due to the fact that Telegram was hosted on infrastructure owned by those two companies. Censorship of messaging applications is common across many countries; For example: Signal, blocked by Egyptian and United Arab Emirates authorities [145][165]; SMS text messaging in Cuba [159]; WhatsApp in China [160], Syria [140] and Brazil

[138]; Telegram in Iran [141]. Media reporting of the blocking of instant messaging services are regular, some contain technical details describing how it was discovered and some are anecdotal. Many of the “larger scale” blocking activities by large censorship regimes, such as China, Russia and Iran, are confirmed and studied within academia. News outlets will often publish articles on censorship activity with little to no evidence, and with the magnitude of fake news reporting on the rise, it is wise to take these accounts with a pinch of salt. Furthermore, the censorship of these kinds of application is commonly temporal, meaning that thorough investigation into them may not be practical. However, when a widely-used service, such as WhatsApp, is blocked within a country, it will usually be reported very quickly in the media or on other social media platforms.

As mentioned previously, anti-censorship, or censorship circumvention, tools or services are common targets of mass filtering. Numerous VPN providers have their services restricted within censorship regimes, such as ProtonMail ProtonVPN [156] in Turkey. Although VPN services are not necessarily designed for bypassing censorship, they are frequently used for this purpose. Tor is an anonymity network (more in Section 2.6) that was developed to provide users a powerful way to communicate with internet services without revealing their IP address or location. It is *not* designed for circumventing censorship however it remains one of the most widely used tools for doing so. This means that Tor is regularly blocked across numerous countries around the world [175][176][123].

Corporate censorship is an interesting phenomenon where the reasoning behind certain blocking activities varies in motivation and intent. For example, the Wikimedia Foundation (owners and operators of Wikipedia) enacted a deliberate restriction of their own service. As a protest against a new EU copyright directive, known as Article 13, the Italian language version of Wikipedia was blocked to all users located within Italy [126]. This activity was not malicious in intent necessarily; however, it showcases how businesses and organisations can use censorship as a powerful tool to convey a message or publicise a point of view to potentially large numbers of people. Corporate self-censorship frequently occurs within social media services, or any platform which make user-derived content accessible. Common laws, or changes to them, can influence these events regularly. Public discussions of sex work on the social media platforms Twitter and Reddit were seen to be removed through policy changes after The Stop Enabling Sex Traffickers Act (SESTA) and Allow States and Victims to Fight Online Sex Trafficking Act (FOSTA) were signed March 2018. This type of self-censorship is due to publishers attempting to protect themselves from consequences of new litigation. The most common area where this happens is within the entertainment industry, where large scale services such as YouTube are under extreme pressure to remove copyrighted and sensitive content from their platforms in different countries [170][148]. Censorship of individuals by social media platforms is a particularly concerning activity that

companies can often be pressured or forced into. This is commonly done during protests or political rallies where the controlling regime wishes to reduce the spread of information via social media. An occurrence of this was reported in Romania, where prominent activists during a mass protests against the government had their Facebook accounts suspended [134][129]. This was not confirmed, but likely due to the Romanian authorities reporting these protesters to Facebook which then proceeded to take action. For these kinds of businesses that operate world-wide, it can be difficult to remain open and free while also upholding local laws and regulations. This issue is further amplified due to investor pressure when access to a particular market could become cut-off if they oppose the controlling power's agenda. Global internet companies such as these are therefore at an intersection between their users, who likely advocate for freedom of speech and expression, and censorship regimes that attempt to control online information.

2.5 Methods for Internet Censorship

The methods and techniques for enacting censorship of the Internet, networked resource or telecommunications are numerous. *How* a specific material is blocked to users depends highly on *what* it is and *where* it is located, either physically or virtually. Content owners, distributors or operators of infrastructure can effectively censor an asset by simply not providing or refusing access to it. For instance, Timothy Jay's lyrics within the song Cop Killer were modified before release due to Time Warner (the record publisher) refusing to distribute it due to pressure from religious advocacy groups [75]. The same applies to online resources, if a publisher does not want, or is under pressure not to provide access to certain content, it effectively engages in corporate censorship for whatever reason. In the case of Internet filtering, there are several key techniques with which an authority can censor a user's access to products, services or content.

In order to enact online censorship, content that is designated to be blocked must be described. This is usually done using blacklists of words, network addresses, domains, URLs, etc. How these lists are prepared is variable, they could be manually or automatically generated.

2.5.1 IP Address Filtering

Restricting access to or communicating with hosts located at specified IP addresses is one of the simplest forms of Internet censorship. It occurs when source or destination addresses within data packets sent over an IP network are monitored by network routers. If the destination address appears on a specified blacklist, the packet will be dropped, either on ingress or egress of the packet. This will disable any transmission of data between the sender and the intended receiver.

2.5.2 Censorship of Internet Domains

DNS filtering is a very widely utilised technique for restricting access to online services. Due to the importance of DNS when communication with remote hosts and the inherent insecurities in how it is commonly operated, it is a natural target for censors to exploit. Almost all applications that require internet connectivity will do so using a known domain name. This allows for operators to switch the hosting server IP address or perform geographical load-balancing transparently using DNS as a means for delivering different information per client query. This is to say that one client connecting from Asia will usually perform a DNS lookup to a geographically closer server than one located within the USA. DNS service providers can use this characteristic to return alternate IP addresses for the end server representing the queried domain. A client will usually make queries for a given domain and then cache the result IP address for an arbitrary amount of time (sometimes given by the result as a time-to-live value in seconds). To block access to an end host via DNS filtering, a censor needs to either control a DNS server or intercept the DNS queries. Given that much of the global DNS traffic is unencrypted, a large ISP can trivially monitor the packets passing through key routing points in their network for DNS queries or responses. They can act on the detection of a blacklisted domain within a query/response in several ways:

- Drop the query/response
- Alter the result within the response
- Drop the query and respond with an arbitrary result
- Return an error
- Drop and timeout the request

DNS filtering is a very powerful means to restrict access to online content. It can often be bypassed using secure implementations, non-restricted DNS servers or encrypted tunnels (such as a VPN or SSH). However, for less-technically adept or unknowledgeable users, filtering of internet domains will usually succeed in blocking their access to online resources. It is also scalable, if a censor has control over a large DNS authority, or the underlying network, they can enforce DNS filtering over large portions of an internet user population.

2.5.3 Keyword Filtering

Searching or isolating keywords in documents and other pieces of online content allows a censor to determine subject topics within textual data. The condition of the monitored data can be variable: news articles, social media posts, blogs, instant messaging, etc. With appropriate access to any of these, a censor can identify the types of discussion being made and which discussions are focused on topics they deem sensitive. This knowledge can be used for a multitude of purposes, including

forming part of a filtering infrastructure, or for simple analytics of communication between different parties. The keywords that are used for this process are usually held on a list that may or may not be constantly updated by a censor authority.

The identification of keywords within data streams is usually achieved through a method called Deep Packet Inspection (DPI), which aims to monitor flows of IP packets for analysis [102]. Since data may be fragmented and potentially distributed via different routes in a network, accurate and reliable detection of keywords requires high-level access to internet infrastructure. Regardless, keyword filtering can still be attained with less than complete data since a packet may contain the word, or part of therein. Importantly, most DPI techniques will require unencrypted data to operate effectively. Censors will therefore attempt to disable encrypted communication over a network, or attempt to access the encryption keys through technical means (such as a backdoor) or legal requirements for ISPs, trusted-third-parties and software manufacturers to surrender access to their root certificates¹.

2.5.4 Filtering of Webpages

The identification and blocking of specific collections of webpages usually encompasses one or more of the previously explained techniques. Servers that hosts websites will usually provide access to hundreds, if not thousands of webpages or a more complex web application. A web application is a dynamic website which will produce constantly updated or changing content based on the user or their preferences, or the publisher adding new information (articles, images, video, etc.) to the site. Static websites will normally just contain simple HTML pages that don't constantly change or update. For example, the website operated by The Guardian (<https://www.theguardian.com>) can be thought of as a *web application*, it is updated with new articles many times an hour, or less. Whereas The Million Dollar Homepage (<http://milliondollarhomepage.com/>) is a static website, the content is not altered frequently, or at all. This means that static sites are simpler to analyse and catalogue to identify subjects or topics. A dynamic site, like that of The Guardian, will offer access to hundreds of different pieces of content across numerous categories and topics. Depending on what a particular censorship regime wishes to block or restrict access will determine which techniques they use to do so.

If a censor wishes to restrict access to a specific host that provides access to a certain website, they can simply block or drop connections to its IP address, if known. The issue with this (for the censor) is that with modern load-balancing and distributed systems, the IP address of the servers that actually host many websites change consistently. They also change depending on the client's

¹Most modern encryption algorithms provide perfect forward secrecy, therefore reading plaintext data sent between users and the end servers would require a man-in-the-middle attack

location¹, load-balancing by geo-location is a common method used by large internet services. This is usually achieved by routing clients to different hosts by providing different responses during a DNS lookup for a particular domain or through content delivery networks (CDNs)—where a client will connect to the nearest server via an Anycast IP address². Occasionally, people do connect to websites using the known, direct IP address of the server, this is not particularly common however, especially for non-technically advanced users.

The more common method for blocking webpages is DNS filtering. This allows a censor to cut access to entire websites by their respective domain names. This means that a client within a network that can filter or manipulate DNS requests/responses could be subject to mass surveillance or censorship of webpages. This method may, or may not be complemented with IP address filtering to further restrict access to online resources if an IP address for a web service is known. An important caveat of the use of DNS filtering in this way is that *entire* domains (or sub-domains) will be blocked. This is because a DNS query *only* contains information pertaining to the domain that the client is requesting an IP address for. Therefore, DNS filtering is coarse and can often over-block a website. For example, much of the information on Wikipedia may not be considered sensitive by a censorship regime, however if a small portion is, the decision to restrict access to <https://www.wikipedia.org> via DNS filtering may be taken. If this occurs, a connecting client within the regime's network will not be able to access *any* of the content on it.

To prevent over-blocking, or to enable a finer-grain censorship model, a censor could use keyword filtering to achieve this. Individual webpages, or requests for certain pages, can be monitored using DPI for blacklisted keywords. If any are discovered, the censor can then drop those specific IP packets to block the respective webpage. This is a far more complicated process and relies on numerous technical features within censorship infrastructure. It is also extremely resource intensive, obtaining full coverage of all access to websites over a network and will require significant computing power and network monitoring capabilities. Furthermore, and as explained previously, achieving accurate keyword detection necessitates the need for non-encrypted channels of communication or the ability to decrypt encrypted data.

2.5.5 Blocking of Social Media

Social media sites are usually incredibly complicated web applications which consolidate several different services into a single offering. For example, Facebook provides a way to share text data, share news articles, upload and distribute images and video content, instant messaging and more.

¹A client's location is usually derived from their IP address

²Anycast IP addresses allow for different routes to be set on a single IP address, using BGP. This then allows for a single IP address to correspond to multiple hosts on a network.

Social media businesses will often need to distribute their services across different data centres and often continents in order to scale to millions or billions of users. This means that restricting access to specific parts of a social media site can be very complex. It is therefore common to see domains belonging to social media business to be blocked in their entirety.

Social media services that are operated within, or by, censorship regimes will often be internally censored. Specific user derived content may be blocked or deleted through the use of keyword or topic detection. A commonly cited example is the filtering of user posts on Chinese social media sites through analysis of the content before it is published [16][58][211].

2.6 Methods for Circumventing Internet Censorship

This is a rich research landscape and numerous techniques for bypassing censorship infrastructure have been utilised over the years. Many of them were not originally developed for circumvention but repurposed to do so. This means that the users who are attempting to bypass a filter on a particular medium or piece of content can often be easily detected. For example, the use of a proxy service to hide the true recipient IP address is trivial to detect if the proxy IP is known to the censor. However, due to the constantly changing nature of internet services, their network addresses and access mechanism, a knowledgeable user can potentially use one or more circumvention techniques to access censored content.

HTTP Proxies are devices or services that will accept and redirect HTTP traffic to and from a source and destination. This is an on-path technique that hides the requester from the receiver through relaying, or proxying, the HTTP requests between them. These services have traditionally been used for anonymity on the web, so that a user can hide their IP address from the end-server¹. Since the end-user is not in direct communication with the end-server, HTTP proxies can allow for bypassing of simple censorship methods such as IP address filtering, due to the fact the end-server's IP will not be within the IP packets source header.

Virtual Private Networks (VPNs) create private networks over public networks. A computer connected to a VPN can then communicate with remote machines over a wide area network (WAN) as if they were accessible over a local area network (LAN). VPNs make this possible by creating encrypted tunnels between the VPN server and remote user machines. The users can route some, or all, of their network traffic through the tunnel via the server (similar to an HTTP proxy) as if they were connected to the internal network, their public facing IP address will then be

¹The sender's IP address could be used for geolocation of the user.

one of the outbound, internet connected nodes on this network. This is akin to how large corporate networks operate where each client machine will route its outbound internet traffic via a gateway. Furthermore, *any* network communications can be routed through the VPN, including DNS queries. Given this, if a user can join a VPN they can then potentially bypass censorship through the tunnelling of their traffic through the private network. This is due to the fact that a censor will likely *not* be able to decrypt the data being routed via the VPN. This would allow a user in China for example to route their traffic through a VPN server in the USA—which means they would bypass the filtering systems in place within the Chinese internet infrastructure.

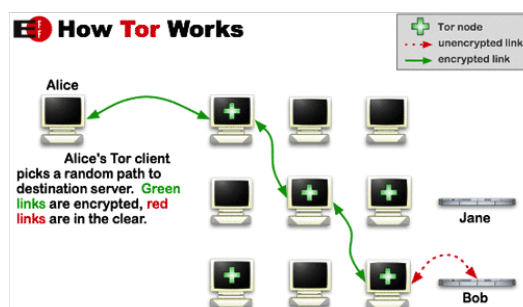
The Onion Router (Tor) is a widely known method and network for providing anonymous communications over the Internet [45]. Tor was developed in the mid-1990's through research projects conducted at the United States Naval Research Laboratory. It is designed to conceal a user's location from actors that can monitor traffic flowing through a network. It was originally developed to allow US intelligence operatives to securely communicate with their respective agencies and not for bypassing censorship systems.

However, given that it effectively creates an encrypted tunnel, similar to a VPN (see below), it can be used for accessing web resources that are blocked in a user's location. Tor consists of several thousand nodes called Relays, through which network traffic is routed from user to recipient. A user will receive a list of all publicised relay nodes and then choose a route to communicate through. The entry node (first in the route) will have knowledge of the user and the exit node (last in the route) will have knowledge of the recipient. The route will then contain an arbitrary number of relay nodes between then entry and exit. Each message sent is encoded with several layers of encryption, one per relay in the route. As the message passes through the route, each layer is "peeled" off until eventually it exits the network and reaches the recipient—hence the term *onion routing*. Figure 2.5 shows a simplified diagram of how Tor sets up anonymous connections between two parties over a (potentially) insecure network.

Tor additionally offers support for censorship resistant resources called "hidden services". These are internet services that operate using the Tor network as a buffer between the clients and end-

Figure 2.5: How Tor works.

Source: *Electronic Frontier Foundation*



servers—thereby offering the service operator anonymity. The purpose is to allow a service to be accessible via the general internet (via an effective proxy through Tor) but to prevent any client from identifying where the origin server is located, what ISP they use or who the operator is, at least via inspection of the details of the connection to the service. Hidden services can therefore be used to host and access content that would otherwise be easily blocked if available via a standard internet connection.

Psiphon is a dedicated censorship circumvention tool that uses multiple different techniques to bypass blocking attempts by censors [157]. It uses a combination of VPN & Secure Shell (SSH) tunnels and HTTP proxies to provide open access to the Internet for users within censorship regimes. As mentioned previously, many circumvention methods are relatively easy to detect, Psiphon aims to reduce this by obfuscating traffic for each protocol it uses. This means that it will be more difficult for a censor to detect if a particular user is tunnelling their traffic via a VPN for example. Psiphon, like Tor, operates numerous relay nodes in different areas of the world that a user can route communications through.

Others There are numerous other methods for circumventing censorship that have been employed over the years. Many are temporary or less scalable solutions used when a web resource suddenly becomes blocked or unavailable in a particular region.

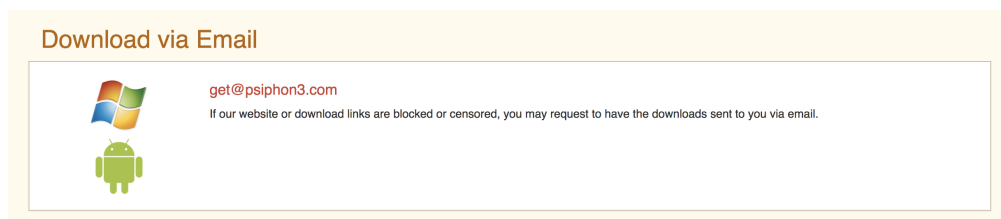
Website clones/mirrors: copies of websites that are hosted on alternative domains or IP addresses that may not be filtered

Domain fronting: where a user connects to a blocked domain via HTTPS while appearing to connect to a different (unblocked) domain. This occurs when the true domain is hidden within an encrypted HTTP header that is indiscernible by a censor.

ICMP tunnelling: where a TCP connection is setup using ICMP ping requests and replies. ICMP is commonly used as part of the IP suite to determine if hosts are reachable or not. If the ICMP message payloads are not monitored, data can be covertly transmitted within them.

Email: sending blocked content via an email system. Due to the fact email is widely used for business and personal communication, it is often less restricted than other communication channels. This means that many censorship systems can be bypassed by simply sending blocked content via email. An example of this is shown in Figure 2.6, where Psiphon will email its software to individuals if its website is censored. Email systems may be monitored however and a censor could potentially identify users who use email for sharing content that is deemed sensitive.

Figure 2.6: Example of bypassing censorship via email



2.7 Existing Approaches for Detecting Internet Censorship

The accurate detection of censorship on the Internet is a significant and challenging research problem. This is due to the dynamic nature of online content coupled with sophisticated filtering techniques employed by censorship regimes [41][67]. While detecting censorship will not itself have a direct effect on the freedom or availability of information, the right for citizens to know when their access to the Internet is restricted or tampered with is important within democratised societies; as posited by Burnett and Feamster [25].

Complex socio-political events that occur in the world have expansive effects on internet usage. Along with the vast number of different rules and regulations between countries, identifying and discriminating censorship from other effects that alter internet usage is difficult. To make matters further complicated, internet censorship and filtering activity is often hidden under veils of secrecy as matter of government policy in some nations. Consequently, to automatically determine if censorship is actually occurring in a certain country from an outsider position, third party datasets are often required.

2.7.1 Existing data sources

There are numerous different data sources that have been used for censorship detection. Previous research has not shown an optimal approach or dataset, but rather that the method used is highly dependent on what data source is selected. A number of recent studies have used data from the Open Observatory of Network Interference (OONI) [55], which was founded by the Tor Project [173]. This organisation relies on volunteers who run nodes in different areas of the world in order to run a number of tests for internet connectivity including HTTP connectivity tests, DNS lookups and traceroutes. As participants are required for OONI to collect data, there are issues concerning the privacy and safety of these people since they will be operating in regions where censorship is apparent.

While OONI maintains a network of nodes and a large amount of data and reports, it does not itself look for censorship events. There has however been external research which is based on OONI's observations. Notably, Gill et al. analysed OONI data between two five year periods: 2003-2008 [61] and 2008-2013 [60]. These studies aimed to produce an approach which could automate the analysis of OONI data as well as identifying areas where the network could be improved for future use. Moreover, there were country specific censorship reports and speculation as to how the Internet was being filtered in these regions.

Similarly to OONI, the OpenNet Initiative [154] also publishes data about blocked and filtered websites in certain countries. It does this using a different experimental network of volunteers who run software that actively checks if URLs are blocked in the county they are situated. Its breadth of data and tests are more constrained than OONI's as it simply uses a predefined list of websites and attempts to connect to them saving any responses or errors that are returned. The Initiative then summarises this data into region reports detailing the suspected filtering activity for countries that may be censoring its citizens.

Because of the close links between OONI and the Tor Project, the data collected by OONI is often compared to that collected by Tor. This allows for OONI to report whether Tor bridge nodes are blocked or not in certain countries and thus can be used as a monitor for Tor censorship detection. In addition to this, the Tor Project also runs its own censorship detection tool developed by Danezis [37]. This system is passive and so relies on the metrics collected by Tor [174]. Essentially, it works by flagging anomalous Tor usage which is achieved by comparing the ratio of connecting clients from each country over the past 7 days. The intuition is that if the ratio falls into an unexpected range for a certain country, there could be a censorship event. Danezis notes that this system can suffer as a result of relying only on the Tor Metrics dataset. This is due to the fact that if a censorship method does not influence this data, there would be no apparent change in usage. Furthermore, if a censor artificially creates Tor traffic, this too would not necessarily be detected. A potentially more suitable model would need to take more previous readings into account.

Detecting censorship by analysing longer term network traffic for anomalies could make up for some of the shortcomings of the detection system currently used by Tor. Wright et al. [203] have proposed a new anomaly detection approach which analyses the Tor Metric data—which is represented as a timeseries—for countries which deviate from global usage trends. This is achieved by using Principle Component Analysis (PCA) [79] over a 180-day rolling time window to construct an approximated model of Tor usage. The approximation is then differenced to the observed usage to produce residuals which represent localised anomalies for each country. This work extended similar approaches presented by Lakhina et al. [89] and Huang et al. [70] which

attempt to diagnose networking traffic anomalies by using data from large Internet Service Providers. In contrast to these, the system used by Wright et al. can detect anomalies on timeseries data with a large number of different variables—for the Tor Metrics, each country is represented as a different variable on the timeseries. That makes the approach efficient for large datasets and the use of PCA residuals allows the tool to ignore large global changes in patterns of usage (since only localised anomalies are identified).

This approach addresses Danezis’ problem of not having a long enough look-back over previous usage. However, it still relies on the Tor Metrics as the sole source of data, which could contain false readings, and therefore the system could produce inaccurate results. Despite this, the two methods for detecting Tor censorship discussed have the advantage of being completely passive meaning that they are non-intrusive and aren’t likely to raise ethical concerns themselves.

In contrast to this, there are other tools which exist that need to be actively installed in remote nodes or used by people around the world—similar to how OONI and the OpenNet Initiative operate. One such approach, that is also aimed at detecting Tor censorship, is presented by Winter [199]. This tool needs to be installed by volunteers that are situated in countries where Tor may be blocked and works by probing the directory authorities, relay and bridge nodes as well as the project’s official website: <https://www.torproject.org>. The results are then sent to a central server and used by Tor developers who work on censorship circumvention techniques for the Tor client. Elahi et al. present an alternative approach for achieving this based differential privacy and multiparty computation, *PrivEx* [49], a method which preserves the anonymity of users and is resistant to malicious actors, whilst still able to produce metrics on censored website visits from different client origins. This is based on statistical data collection and does not require interaction from individual users or volunteers.

Data leaks from filtering equipment or devices have offered new routes for censorship analysis, a large disclosure of log information from a *Blue Coat SG-9000* appliance in 2014 provided such an opportunity. Chaabane et al. used these data [27] to uncover the extent to which the Syrian government had imposed internet restrictions on its citizens. In particular, the use of instant messaging and social media was heavily blocked through IP, DNS and keyword filters.

2.7.2 Collecting new data sources

This kind of data that Tor collects is very important to the Tor developers, however, it is not necessarily easy, or ethical, to collect it. The issue is that you need remote machines—that are usually in use by people living in the suspect countries—to report when access to Tor is blocked. Ensafi et al. attempt to address this in their approach [50] which uses TCP connections as a side

channel for inferring a blocked connection. This method only requires a client with a global IP identifier and a server with an open port in the target region. They then make connections between the client and server and look for packet drops between them. The direction of the drop is recorded and therefore, over multiple tests from multiple clients, the technique enables them to infer which connections are censored. Their study found that only 0.63% of the measured connections to Tor from China are not blocked. This study was followed up in 2017 [111] with more comprehensive testing of 180 countries where the highest number of internet disruptions were found in *China, Iran* and *Sudan*.

Winter's tool, which requires cooperation between a number of clients and a centralised server, is an attempt to collect data which could help to infer if certain domains are filtered in nations suspected of censoring their citizens. There have been similar approaches since which are also trying to "crowdsource" this kind of information. In particular, there is Encore [26], User-Based Internet Censorship Analysis (UBICA) [2] and Web Content Monitoring Tool (WCMT) [52]. UBICA operates in the same way as Winter's approach, except it is aimed at a broader set of websites and content, not just the Tor Project. It too requires installation by volunteers in suspect countries and attempts to monitor the access to potentially restricted content over time as well as predicting what kind of censorship infrastructure is in place (i.e. is it centralised). WCMT also works in this way, and is based on how OONI works, but with less of a dependency on Tor and some extended capabilities for checking blocked services and ports.

Conversely to UBICA and WCMT, Encore relies on website owners to install a small piece of javascript on their webpages. This script will be executed by any clients that access said website and the code executes a cross-origin request to a list of URLs and reports the responses. This system has a number of advantages over other censorship detection tools since it could potentially produce a large amount of data from many different vantage points on the Internet. If a website that installs Encore has a lot of users from different countries suspected of censorship, there would be a lot of requests from different clients to potentially blocked websites. What's more is that Encore doesn't require any access to any of the users and the system is distributed in nature—so, if there were enough participating webmasters, it should be protected from network faults. However, there are again ethical concerns for the safety of the people who access these volunteer websites when essentially using them as nodes in a censorship detection tool¹. Besides this, Encore will only work if there are enough webmasters who are willing to install the necessary code, as well as a whole other host of security problems that may arise from implementing the code.

¹There were major discussions about the ethical concerns of ENCORE during a workshop held by ACM after SIGCOMM 2015. ACM made an unprecedented move to add a disclaimer before the paper explaining they did not endorse the experiments performed in the study.

Learning about the censorship infrastructure that is used by different nations is important part of this research problem. We see that some of the tools and studies aim to not only detect when and where censorship occurs, but also what filtering product/system is in place. UBICA shows that the censorship mechanism can be predicted, but only if it is within the medium you work in. Hence, UBICA can only detect censorship and the underlying infrastructure if it happens within the HTTP or DNS layer of the networking stack. Jones et al. [81] take a different approach where they are able to detect censorship block pages automatically if presented one in an HTTP response from a webserver. They use a library of known block pages maintained by the OpenNet Initiative to train a home-brew classifier that can identify a censorship block page. Their study also explores how block pages differ when produced from a single country. This information gives clues as to what kind of filtering system is in place as well as its scale. Using this approach, the researchers are able to give a much higher level picture of the censorship infrastructure in particular nations¹.

2.7.3 Approaches using existing infrastructure

All of the tools and approaches discussed so far rely on 3rd-party data sources or volunteers in suspect countries to run tests and produce the required metrics. Arguably, this only way to know for sure if certain content, in particular regions of the world, is censored or not. You need some kind of verification, which essentially boils down to requiring a number of people (or agents/nodes) in different parts of a censorship infrastructure to check if they can access a potentially blocked URL or service. Any outside tests could be biased by the conditions of the experiment, or a potentially omniscient censor who knows when a connecting client is running some sort of censorship analysis (e.g. they may execute tests from an outside country).

Yet, as mentioned previously, there are issues with this kind of censorship analysis. Aside from the obvious ethical concerns, these tools require fairly large amounts of maintenance since they need a large number of participants to operate which is itself also an issue. CensMon [184] attempts to get around this by using servers provided by PlanetLab [155]—which is a fair compromise, however it still relies on a 3rd-party service. All this being said, researchers have addressed this and developed large number of censorship detection approaches and tools which don't require any volunteer nodes to work.

An interesting example of such an approach is presented by Michel et al. [96] who compiled a corpus of over 5 million digitized books. By simply analysing the change in the number of times particular n-grams appeared over the years, they were able to detect historical censorship acts. An example of this is the suppression of published works containing the Jewish artist Marc Chagall by

¹This work uncovered a filtering device that was used by Saudi Arabia that had not been seen in the wild before the time of their study.

Nazi Germany. This is apparent by comparing the number of times his name appears in German books published in the 1910s with books published between 1936 and 1944 and then again between 1946 and 1954. There were also similar occurrences with Russia and “Trotsky” and China and “Tiananmen Square”. While this simple approach is effective, it only works in hindsight and is unlikely to provide much success in detecting censorship “in realtime”, although if it were applied to a faster moving medium (like social media), it could potentially give good results. As it happens, this approach is similar to the analysis of filtered keywords on the Internet which will be touched on later.

Michel et al. were able to analyse n-grams within their corpus because of the efforts expended by Google Inc. to digitize a large number of books and make the resulting metadata public. The use of these kinds of datasets can allow us to run analyses without the cooperation of any other party and without any internal vantage point. The same can be achieved by utilising public facing internet services such as DNS or webservers.

This kind of approach was used by Quan et al. [118] who studied Internet outages by probing IP addresses within the IPv4 Internet. While not directly aimed at detecting censorship, this method was able to track outages across the entire IPv4 address space from a single machine by targeting the probing to a small number of IP addresses in different /24 address blocks. This technique essentially reduces the number of addresses that need to be scanned to get a view of the Internet by 128 times. The authors show that their experimental results imply that around 0.3% of the Internet is unreachable at any one time—which is consistent with other studies in this area.

Furthermore, there is a collection of probes operated by RIPE Atlas [163] which measure the reachability of areas of the Internet, at the time of writing, they have over 8400 probes running tests. Yakimov [209] shows that this network can be used as a viable platform to detect routing anomalies on the Internet based on existing methodologies and therefore could also be used for censorship detection platforms. Scanning IP addresses is a semi-passive method for detecting censorship (because you still need to make connections to remote machines in suspect countries), another widely used technique is to use the public DNS servers available in different nations.

Published BGP network routes and internet connectivity scans can be used as data sources for identifying when possible nationwide censorship events have occurred [35], using this, Dainotti et al. posit that large-scale network disruptions in Libya were the result of tests conducted of firewall-based and BGP-based blocking methods. Furthermore, mapping and probing the network topology can provide information about where filtering devices are located [206].

2.7.4 DNS based approaches

DNS filtering and the black-listing of domains is widely used method for censorship [48][67][212]. It's fairly simple process and doesn't require a drastic change of any underlying network infrastructure that makes up the Internet. Some of the earliest research of internet censorship focused on DNS as a key part of detection systems [200]. DNS works using a hierarchy of multiple servers, the higher up you go, the more authority that server has. If any server in the chain doesn't know the IP address for a domain, it looks further up the hierarchy. Therefore, at any point, a server can respond with an error or an indiscriminate IP address. In most countries where DNS filtering is used, the ISPs will have a list of banned domains which, if requested for in a DNS lookup, will be responded with something other than the correct IP address for that website. Clearly this is a good avenue for censorship detection since large DNS servers must be publicly available.

Bailey et al. [15], Scott et al. [183] and Wright [202] have all used DNS as the data source for detecting the filtering of certain domains. Scott found that at least 13 countries in a 59 country experiment were using DNS filtering as a means for censoring content online. Wright's work focused on Chinese censorship and found that DNS servers within China commonly responded with fake IP addresses for black-listed domains¹. Verkamp et al. found that four of their 11 tested countries used DNS to censor the Internet [193], *Malaysia, Russia, South Korea and Turkey*. They report that, at the time of their study, these censors forwarded clients to block pages or local/private IP addresses.

One of the most in-depth looks at Chinese DNS censorship, conducted by an anonymous party [11] in 2014, found that the censorship equipment used a keyword blacklist of over 15,000 terms which, if present in a domain name, would trigger a blocking event. This group's study went as far as determining the individual IDS devices that were operating by looking at IP ID numbers during TCP connections². Similarly, a study by Weaver et al. [196] details how different filtering devices exhibit unique behaviours when sending TCP Reset packets to clients. This allowed for the accurate detection of connection interference from remote vantage points.

DNS censorship is widespread, and while many of the discussed studies focus on China, other countries have also participated in DNS censorship: India [6][63][208], Indonesia [82][194], Iran [13], Pakistan [1][103], Thailand [59] and Turkey [8][43][156]. Pearce et al. [112] found that in 2017 the top manipulators of DNS results were: *Iran, China* then *Indonesia*; and, Scott et al. report that in total, over 117 different countries were found to implement some form of DNS hijacking [182] during their measurements in 2016.

¹An interesting example being a number of different servers in China responding with an IP address for The Pet Club in Florida (<http://www.thepetclubfl.net/>) when a look up for <https://www.torproject.org> was made—<http://www.pseudonymity.net/?p=105>

²Our follow-up testing of this approach appeared to show that the IP ID counters used are now randomised, which effectively prevents the identification of devices in this manner.

2.7.5 Combination methods

It is well known that the People's Republic of China make a large and concerted effort to censor material from the Internet. The country has an Internet policy which essentially allows the government to regulate it as they wish [28] and due to the ferocity with which the state controlled Internet activity with during the 90s and 2000s, the Chinese censorship mechanism has been dubbed: “*The Great Firewall of China*”. Censorship research by Wolfgarten in 2006 shows that the firewall uses a multitude of techniques to implement content filtering [200]. Numerous efforts have been made to try and deduce how this system works, notably King et al. [84] and Clayton et al. [31]. It's generally accepted that the Firewall uses a number of different techniques to impose censorship of the Internet—including, but not limited to, DNS filtering, deep packet inspection and keyword filtering. While the history of Chinese Internet censorship shows that the government attempts to block any trace of anti-government content and other areas of interest, it has been theorised recently that the filtering effort has been aimed at collective expression rather than an individual's criticism of the government [83]. A reason for this could be the huge rise of online social media outlets over the past decade.

The Chinese government has been pushing a discourse of self-censorship for publishing institutions within the country [189], this is possible given the size of most of the media outlets that have widespread influence, but it's not possible in social media. Take a website that collects, stores and repeats the published work of millions of individuals and call it “*online social media*”, the banning of particular topics of discussion is rather difficult—especially if you cannot influence the company which operates the service. Hence, the largest social media networks in the world are heavily filtered from the Chinese internet and the larger networks within China are operated under tight scrutiny from the government.

Many published censorship detection approaches use several measurement methods to acquire data. Dalek et al. [36] shift from traditional probes and scans however, by looking to identify the machines and equipment that are operated to enact censorship, rather than the censored content itself. The authors report that *Blue Coat*, *McAfee SmartFilter*, *Netsweeper* and *Websense* devices are used in over 12 different countries outside of the original countries of manufacture (*Canada* and *USA*). Where the type/brand of filtering devices are not known, it is still possible to infer how they operate and in which locations, Park et al. [110] used probes directed at China to uncover the distributed nature of the systems in place used to scan HTML responses. Their work concluded with the observation that this “deep” level of filtering—based on the content rather than request meta data (i.e. IP address, domain, etc.)—was actually being discontinued due to the difficulty and cost of scaling such a system, along with the availability and use of more

sophisticated censorship evasion technologies (like Psiphon detailed in Section 2.6). Furthermore, methods for producing useful circumvention paths for internet censorship against specific filtering devices are available [76].

2.7.6 News and social media monitoring

A number of different approaches for detecting the censorship of social media in China have been published over the years. A widely-used method is to look at keyword filtering enforced by the Firewall, Crandall et al. [33] present ConceptDoppler, a tool which probes Chinese censorship infrastructure to build a list of blocked keywords. This approach continually checks different words which may be potentially censored to update the list over time and in doing so interestingly discovered that the largest ISP, ChinaNET, performs over 80% of the black-listing. Identifying useful keywords for censorship probes is an open research question, Espinoza et al. [53] deviate from most other methods by extracting names, entities and text from news articles which are then tested against Chinese search engines to check if they result in block pages. It was also discovered that mobile applications can be used as source for capturing blocked terminology because their developers operating in China cooperate with the authorities and regulators by providing details of their internal keyword blacklists [86].

The largest social media network in China is Sina Weibo, this is a micro-blogging service and works in a similar way to Twitter. Zhu et al. report that the network is heavily engaged in the filtering of certain content [210], however, their results suggest that these efforts aim to stamp down on the virality of sensitive topics rather than the complete censoring of them. This backs up the theory of the blocking of collective expression mentioned earlier. Zhu et al. take their study further in [211] and speculate on the mechanism of censorship used by Sina Weibo by looking at the timescale of microblog post deletions. They find that if a sensitive topic is already trending on the network, a post could be deleted on submission. Whereas, if the topic trends at a later time, the post may be deleted then. This is an interesting find and helps to identify what content is being censored in China at specific times. Fu et al. [58] also consider the censorship of microblog posts in China and attempt to find what keywords are being blocked at certain times. They also find that many Chinese microbloggers use homophones and puns of known blocked words to bypass the filtering. This illustrates the constant fight between the censors and the population, what's more, the rate of post deletion varies throughout the country with the highest number of deletions taking place in Tibet and Qinghai, as shown by Bamman et al. [16].

Social media self-censorship, administered by the operators of the network, is also widespread outside of China. Tanash et al. report that over 250k tweets were censored by Twitter themselves

after receiving requests from Turkey [188].

Many approaches and tools that look at social media censorship concentrate on the deletion of posts that discuss certain sensitive topics. This is mainly achieved by searching for keywords, however as discussed, bloggers can get around this by using alternative language. To block these kinds of content Chinese censors rely on people to report the communications as inappropriate. This activity has a deeper effect on the underlying social graph. Morrison [101] presents a method for automatically detecting when parts of a communication graph have been censored. He does this using a Support Vector Machine as a classifier which is trained on simulated censorship of social graphs. While this approach hasn't been applied to Chinese filtering, it does represent an interesting way of identifying what sensitive topics are being blocked by considering the communication patterns between users of a social network.

2.7.7 Summarising the opportunities for further research

Many censorship detection tools aim to not only identify periods of time when content is blocked, but also the underlying infrastructure used to perform the filtering. It is clear that nations who censor their citizens do so using various different methods. A possible reason for this could be because of different motives and aims that the authority wishes to achieve. In the case of China, we see that there are large efforts in keyword analysis for filtering content at a very large scale. Iran is another state that operates a very large censorship regime and because of the secrecy of the Iranian government, very little was known about Iran's internet infrastructure. Recently however, Aryan et al. [13] reported that the Iranian methods of censorship are similar to that of the Chinese. Aside from DNS and keyword filtering, there is evidence of connection throttling, especially for SSL requests. This has also been separately reported by Anderson [10] who explains that this kind of censorship is much more difficult to detect than the outright blocking of content. This is due to the way that connection throttling works as it slows a connection down to infeasible levels where the downloading of a webpage could take several hours. But the connection is still there. This method aims to shift the culture of a society so overall it doesn't want to access banned material because of the time it would take to receive it. OONI reports that all Tor connections were throttled in Iran in 2011, which if true, shows that the Iranian government also performed deep packet inspection to differentiate Tor traffic with "normal" HTTPS traffic. It is believed that this throttling has ceased now, however, there continues to be a very low number of Iranian Tor users. This indicates that even though Tor could be available there, it has lost most of its appeal with the citizens in Iran.

Having discussed a broad range of different approaches for detecting Internet censorship, it is clear that this area of research is active and complex. However, there are a number of avenues which

require further work in this subject.

Firstly, many studies into censorship have focused on well known censorship activity, namely that of China and Iran. There has not been a large amount of published work on smaller or less obvious censors in the world. This topic has room for a lot of research and can go hand in hand with other studies into censorship detection—especially those which identify nations which would not normally be considered as having Internet censorship or filtering agendas. This also means that many approaches are tailored towards taking censorship measurements in those countries, on the particular infrastructure that is used in those countries. Furthermore, some of these approaches can be difficult to repeat or reuse, partially because some use specific attributes in a particular censor's infrastructure—which may change over time—and also because some use non-public datasets—such as leaks or dumps of confidential information.

Secondly, up until this point, the majority of censorship detection approaches focus only on determining if a particular piece of content or resource is censored, rather than openly discovering other (possibly unknown) pieces of content, at least by other means than volunteer (human-led) efforts. Automating this discovery process could yield a significant improvement whereby fewer people are required—costing less and taking less time to collect and verify candidate censored materials for any particular country. Furthermore, frameworks for achieving this could be integrated with other approaches for acquiring measurements. The key factor that needs to be addressed is the method for generating candidate test URLs/content/keywords since any operational system will need to automatically identify a suitable number of potentially filtered items to test.

Thirdly, there are several pertinent ethical issues with censorship research, particularly around the collection of measurements. As the field moves forward, there is increasing pressure to acknowledge and mitigate these. Many recent pieces of research have been conducted and written with this aspect firmly in the centre of the work. This is further explained and discussed in Section 2.8 below.

Finally, there is a scalability challenge. Human-based approaches are inherently costly to scale, especially when a large amount of manual effort is required—to verify a censored website, for example. Agent (or node) based approaches also hit scaling issues where it is required that more agents are deployed in order to acquire more or better measurements. This is made worse if an agent, or set thereof, is required within each geographic territory to be measured. New methods to counter these scalability problems could allow for an order of magnitude more measurements to be acquired in shorter time frames.

2.7.8 Catalogue of Existing Approaches

Summarised in Table 2.1 is a set of important censorship research pertinent to this thesis. This shows the breadth of different and diverse approaches taken to detect censorship activity in the research community over the course of the past two decades.

Name	Author(s)	Date	Type(s) of approach	Country specific
Internet filtering in china	Zittrain et al. [212]	2003	DNS, HTTP, keyword analysis	China
Investigating large-scale Internet content filtering	Wolfgang [200]	2006	DNS, IP, TCP	
ConceptDoppler: a weather tracker for internet censorship	Crandall et al. [33]	2007	Keyword analysis	
Detecting Forged TCP Reset Packets	Weaver et al. [196]	2009	IP, TCP	
Quantitative analysis of culture using millions of digitized books	Michel et al. [96]	2010	Text analysis	
Empirical Study of a National-Scale Distributed Intrusion Detection System: Backbone-Level Filtering of HTML Responses in China	Park et al. [110]	2010	HTTP, probes	China
Greatfire.org	Anonymous [64]	2011	Probes, volunteers	China
Censorship and co-option of the internet infrastructure	Bailey et al. [15]	2011	DNS	
Analysis of Country-wide Internet Outages Caused by Censorship	Dainotti et al. [35]	2011	BGP, IP	
An anomaly-based censorship detection system for tor	Danezis [37]	2011	Tor metrics	
Automated Named Entity Extraction for Tracking Censorship of Current Events	Espinoza et al. [53]	2011	HTTP, keyword analysis	China
Censmon: A Web Censorship Monitor	Stakianakis et al. [184]	2011	DNS, HTTP, probes	
Internet Censorship in China: Where Does the Filtering Occur?	Xu et al. [206]	2011	IP, probes	China
Censorship and deletion practices in chinese social media	Bamman et al. [16]	2012	Social media	China
Open Observatory of Network Interference (OONI)	Filastò et al. [55]	2012	Probes, Tor metrics	
Detecting internet outages with precise active probing	Quan et al. [118]	2012	Probes	
Inferring Mechanics of Web Censorship Around the World	Verkamp et al. [193]	2012	DNS, HTTP, probes	
Tracking and quantifying censorship on a Chinese microblogging site	Zhu et al. [210]	2012	Social media	China
Detecting throttling as a mechanism of censorship in Iran	Anderson [10]	2013	SSL/HTTPS	Iran

Internet censorship in Iran: A first look	Aryan et al. [13]	2013	DNS, HTTP	Iran
A Method for Identifying and Confirming the Use of URL Filtering Products for Censorship	Dalek et al. [36]	2013	HTTP	
Web Content Monitoring Tool (WCMT)	Esnaashari et al. [52]	2013	Probes, volunteers	
Assessing censorship on microblogs in china: Discriminatory keyword analysis and the real-name registration policy	Fu et al. [58]	2013	Social media	
Detecting DNS censorship without an internal vantage point	Scott et al. [183]	2013	DNS	
Towards a censorship analyser for tor	Winter [199]	2013	Probes, volunteers	
The velocity of censorship: High-fidelity detection of microblog post deletions	Zhu et al. [211]	2013	Social media	China
Towards a Comprehensive Picture of the Great Firewall's DNS Censorship	Anonymous [11]	2014	DNS	China
Censorship in the Wild: Analyzing Internet Filtering in Syria	Chaabane et al. [27]	2014	DNS, IP, keyword analysis	
Detecting intentional packet drops on the internet via TCP/IP side channels	Ensafi et al. [50]	2014	Side-channel, TCP	
Automated detection and fingerprinting of censorship block pages	Jones et al. [81]	2014	HTTP	
Toward automatic censorship detection in microblogs	Morrison [101]	2014	Social media	
Regional variation in Chinese internet filtering	Wright [202]	2014	DNS, Tor metrics	China
Detecting routing anomalies with Ripe Atlas	Yakimov et al. [209]	2014	Probes	
User-Based Internet Censorship Analysis (UBICA)	Aceto et al. [2]	2015	Probes, volunteers	
Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests	Burnett et al. [26]	2015	JavaScript, probes, (uninformed) volunteers	
Known Unknowns: An Analysis of Twitter Censorship in Turkey	Tanash et al. [188]	2015	Keyword analysis, social media	Turkey
Analyzing internet censorship in Pakistan	Aceto et al. [1]	2016	DNS, probes	Pakistan
Satellite: Joint Analysis of CDNs and Network-Level Interference	Scott et al. [182]	2016	DNS	
Mending Wall: On the Implementation of Censorship in India	Gosain et al. [63]	2017	DNS, IP	India
Autosonda: Discovering Rules and Triggers of Censorship Devices	Jermyn et al. [76]	2017	HTTP	

Measuring Decentralization of Chinese Keyword Censorship via Mobile Games	Knockel et al. [86]	2017	Keyword analysis	China
Augur: Internet-Wide Detection of Connectivity Disruptions	Pearce et al. [111]	2017	Side-channel, TCP	
Global Measurement of DNS Manipulation	Pearce et al. [112]	2017	DNS	
Automatically Generating a Large, Culture-Specific Blocklist for China	Hounsel et al. [69]	2018	HTTP, keyword analysis	China
An analysis of automatic image filtering on WeChat Moments	Knockel et al. [87]	2018	Image analysis, social media	China
403 Forbidden: A Global View of CDN Geoblocking	McDonald et al. [93]	2018	HTTP	
Detecting Censorable Content on Sina Weibo: A Pilot Study	Ng et al. [105]	2018	Keyword analysis, social media	China
Quack: Scalable Remote Measurement of Application-Layer Censorship	VanderSloot et al. [192]	2018	DNS, IP	
On identifying anomalies in tor usage with applications in detecting internet censorship	Wright et al. [203]	2018	Tor metrics	
Measuring I2P Censorship at a Global Scale	Hoang et al. [68]	2019	DNS, probes	
Distributed Detection of Tor Directory Authorities Censorship in Mexico	Iszaevich [73]	2019	Tor metrics	Mexico
An Efficient Method to Determine which Combination of Keywords Triggered Automatic Filtering of a Message	Xiong et al. [205]	2019	Keyword analysis, social media	
Detecting and Evading Censorship-in-Depth: A Case Study of Iran's Protocol Filter	Bock et al. [23]	2020	DNS, HTTP, probes	Iran
ICLab: A Global, Longitudinal Internet Censorship Measurement Platform	Niaki et al. [106]	2020	DNS, HTTP, probes	
Censored Planet: An Internet-wide, Longitudinal Censorship Observatory	Raman et al. [121]	2020	DNS, HTTP, IP, probes, side-channel	
Investigating Large Scale HTTPS Interception in Kazakhstan	Raman et al. [120]	2020	HTTP, probes	Kazakhstan
Measuring the Deployment of Network Censorship Filters at Global Scale	Raman et al. [122]	2020	HTTP, probes	

Table 2.1: Catalogue of approaches for detecting Internet censorship
Rows shown in grey were published after the time of experimentation conducted in this thesis.
These have been added to show the breadth of more recent research in this space.

2.8 Ethical Issues with Research into Censorship

Censorship research is sensitive, and touches on a variety of ethical concerns that we must consider. Wright et al. [204] provides an early look at the legal and ethical concerns of mapping censorship events. They deal primarily with mapping where filtering is occurring and less on the content of what is being filtered; focusing on techniques that rely on citizen volunteers to discover and report filtered content. Two later papers both published in 2015 broaden the scope of this area. The first of these by Jones et al. [80] identifies three different approaches to measuring censorship: deploying researchers with software, deploying citizens with software, and co-opting existing software. The second by Crandall et al. [32] further differentiates censorship detection between direct observation and inference using side channel experiments. All three of these papers argue that it is important to understand the technology and motivations behind filtering, while aiming to reduce the ethical risk to individuals during the data gathering phase.

In the work presented by this thesis, we use existing, publicly available data sources. Instead of collecting data ourselves, using individuals or sensors, one of our approaches use data that has already been collected by search engines. This technique is non-invasive and—crucially—not focused on data gathered directly by researchers or citizen volunteers; thus, removing the risk to individuals. While Google or Microsoft have their own set of ethical considerations—both when initially collecting the data and when making it available to others—ours become more subtly nuanced. Rather than look at the individual, we must consider the ethics of performing academic analyses of deliberately closed systems, and of using publicly available data gathered for a specific purpose to be used for different and arguably political purposes.

The first consideration is that the censorship we are observing has been installed by nation states deliberately as closed systems. They generally do not publish the details of how such systems work, or what content they are filtering. Authorities may argue that these systems are in place for the benefit of their citizens, and that obscuring their details is important to their function, or just remain silent on the subject. Is it ethical for us to observe—and publish—data about such systems? This is considered by the aforementioned Wright [204] but he argues that these practices are of insight to us as researchers, not just to help with understanding the internet, but also to provide insight into social and political issues. Crandall et al. [32] agree, stating that research such as this provides empirical data on censorship around the world, and that this data may be of use to political scientists and sociologists, and even to the general public.

The second consideration is of using data for purposes outside its original intent. Search companies gather data—ostensibly—to provide search results, and—cynically—to gather data on users and direct them to monetised content. Is it ethical for us to use this data to infer the details of

the aforementioned closed censorship systems? Our approach relies upon the cooperation of search engines, and access through their API. We believe that as long as their terms of use are obeyed, and our approach does not pose any risk to individuals, the ethical concerns of data used for alternative means is minimal.

2.8.1 Risks of measuring censorship

Most currently deployed censorship detection machines require the use of volunteer time or physical computer hardware located within the geographical region being measured. This is so a test can be conducted on whether a specific website, webpage or resource is being actively blocked in the region, thus confirming the censorship state of the content. Whilst this can offer an accurate view of censorship activity, the use of “people on the ground” also presents a set of issues:

- The physical safety of the volunteer could be compromised if situated in a particularly aggressive regime if they are discovered to be accessing censored content
- Running censorship circumvention/detection software or devices may be explicitly illegal in the region to be measured
- A “human-based” detection mechanism scales linearly where additional volunteers will be required in order to increase the rate of measurements, this can also introduce additional costs into the process
- The use of volunteer hardware presents further cyber security issues where additional steps in securing the measurement process/software should be undertaken so as not to expose the volunteer to unintended security risks

These are high enough risk issues that they make the use of volunteers for measuring censorship a particular ethical and moral problem. The fact is that any attempt to scan, detect, collect or identify censored material in any region where the access to, or dissemination of, such material is illegal, meets an ethical issue. Yet, conducting this research is of great importance to help to understand the discrete censorship activity taking place in disparate regions across the world and the use of human volunteers offers a valuable approach for achieving this.

One method of reducing the risk of these issues is to take measurements directly from censorship infrastructure, that is to say: a computer component specifically designed to censor the internet communications. Our indirect measurement approach complements many existing approaches—those with or without the use of people on the ground—and is discussed further in Chapter 3.

2.8.2 Misuse by adversaries

This thesis proposes a system that can be used to find content filtered by many censorship regimes around the world. In essence, this is an attempt to reverse engineer the decision making and technical processes put in place by such nation states that block web resources. The methods we use aim to isolate particular topics or threads of censored content that may lead to similar types of material that is itself also censored. This is achieved using widely available services and infrastructure (search engines) that are designed to group topics together, controversial or not, to bring order to unstructured data (natural language) and then make it available to the public. Recent advancements in natural language processing allow modern search engines to perform this with great accuracy. Hence, the question of misuse of the proposed frameworks are at the forefront of attention.

Our aim is not to make available a system that can be used by adversaries to make further content on the Internet unreachable to certain groups of people. This is however, a potential application of it. We must therefore, consider if it is in the moral interest of the public for such frameworks to be widely published. Given the previous discussion on the ethics of censorship measurements, we feel that this study is of benefit to the wider research community due to the requirement for more advanced tooling for testing censorship. For researchers and future studies, our system will be of use in on-going work into Internet transparency and socio-political issues regarding censorship.

“What you do speaks so loudly I cannot hear what you say”
Ralph Waldo Emerson

CHAPTER 3

A Method For Determining The Censorship Status Of Webpages And Domains

The Internet has been developed to be fault-tolerant, available and consistent. The distributed nature of key services that form the backbone of this massive, global network provide *multiple routes* for IP traffic across the globe via numerous sub-networks. Original research networks, such as ARPANET, JANET and NFSNET, were extended to connect towns and cities over land, and continents via undersea cables. The outcome of these developments was the introduction of new protocols and services that allow the Internet to operate at such scale. Of major importance are the AS routing prefixes and BGP. Both of these underpin the majority of all internet traffic, especially that between different countries and Tier 1 networks. BGP aims to provide redundancy for routes between major ASs so that clients can connect to their desired destinations, even if parts of one or many routing networks become unavailable. Another cornerstone of the modern Internet is DNS—introduced in 1983 [100]. DNS has further pushed the decentralisation of wide-area-networking and makes the use of IP more user-friendly. The product of these advancements is a highly *inter-connected* and *failure resistant* system.

Yet, this global super-structure is incredibly fragile regarding its security. Many of the key routing systems that direct traffic are vulnerable to attacks from a host of different actors, from *nation states* to *script kiddies*¹. An example of such an attack was the Turkish Hijacking of public DNS Providers in 2014 [43][132][135] by advertising false BGP routes—an activity that can currently be performed by most large networks. While this was an attempt to censor connections to Twitter originating from within Turkey, it exhibits the insecurity of these systems since any machine connected to the

¹**script kiddie:** noun. (Hacker Lingo) One who relies on premade exploit programs and files ("scripts") to conduct his hacking, and refuses to bother to learn how they work.

Turkish Internet was also affected, including external ISPs. Moreover, had BGP routes outside of Turkey been altered, a major DNS blackout could have occurred. This shows that attacks on essential infrastructure such as key routing devices, DNS and major ISPs are becoming more common and the systems in place are not built to deal these kinds of exploits.

The security of DNS is of major concern currently. The basic system itself does not offer any authentication, verification or encryption of the servers, messages and transmission respectively. While secure implementations of DNS are available, they're *not* widely adopted and most systems do not use them. These traits allow censors around the world to interfere with DNS queries in a highly scalable manner within the infrastructure within their jurisdiction. This chapter will discuss examples like these and show how we can use the insecurities of internet systems to gain a vantage point for censorship activity. Further, designs for a novel method that uses DNS systems as test subjects for censorship measurement instrumentation are discussed, a key contribution of this thesis.

3.1 Detecting Blocked Domains

Interfering with DNS as a means to censor content on the Internet is known to be practised in numerous countries, such as: China [11][91][202], India [6][63][208], Indonesia [82][194], Iran [13], Pakistan [1][103], Thailand [59] and Turkey [8][43][156], among others. Given that almost all communication between peers on the Internet will make use of at least one DNS query as part of a request, if a DNS service becomes unavailable or is indeed blocked, many applications will cease to function. DNS therefore has been, and currently is, a key target for many censors as a comparatively low-cost and low-effort means of restricting access to Internet resources. As mentioned, this has been studied over the past decade at considerable length and depth, and covering various countries. The ubiquity of DNS services in use across the general internet *and* by censors makes it a useful channel for us to capture censorship measurements from. Through understanding how a DNS service looks in a “normal” and non-censored operating condition, we can create instrumentation to identify disparities in DNS responses which can be used to detect the censorship of an internet domain. This process, while lacking in the ability to target specific webpages, is highly accurate and robust when measuring at a domain/sub-domain level. Furthermore, we can take measurements against large-scale infrastructure, remotely and without the need for people on the ground.

3.1.1 Route 53: How DNS Works

Most ISPs operate their own DNS services which the majority of internet users will often use by default. These DNS servers are usually set automatically on the user machines by the Dynamic Host Configuration Protocol (DHCP) when they connect to the ISP's network. There are also many other free public DNS operators that provide access to reliable DNS services as alternative to an ISPs offering. These include:

Cloudflare 1.1.1.1

Dyn 216.146.35.35

Google Public DNS 8.8.8.8

Level3 209.244.0.3

OpenDNS 208.67.222.222

It is generally assumed that these public DNS services are secure and uncensored, yet, little research has been published directly on this topic. However, these geo-distributed services are widely used among the research community in controlled experiments for censorship measurements [38][39][112].

When a DNS query for domain resolution is made, the requesting machine will usually send it using UDP on port 53¹. Because of this, multiple responses can be generated by the server, since UDP does not provide a guarantee of delivery and the underlying network may be unreliable. Furthermore, a client will usually only accept the *first* response received, where subsequent responses for the same query are discarded. This is an important aspect of DNS blocking as will be explained later, especially with regards to DNS interception. Queries are not encrypted by default and DNS servers do not often provide a means of authenticating themselves to remote users. This means that servers can be spoofed by attackers in a transparent manner. Moreover, channels of communication between remote hosts during DNS queries can be sniffed by eavesdroppers since the transmission is in plaintext. Secure protocols for DNS do exist, such as DNSSEC [150], DNSCrypt [42] and DNSCurve [20], but they are not in widespread use as of 2019.

Due to the inherent security weaknesses in the standard DNS protocol, censors have the capability to manipulate DNS traffic at scale. Even though DNS only provides part of the network stack required for internet communications, it is highly integrated into the way networked applications operate. For example, Windows Update will make requests to *update.microsoft.com*, not to a specific IP address. Therefore, for the update system to operate, the application must be

¹The DNS server listen port can vary, however 53 is the default.

able to resolve the IP address for that domain using DNS. An adversary who can monitor and manipulate the transmission channel between a client and their configured DNS servers could alter the query or response when the domain name is being resolved. This forms the basis of *DNS interference*.

3.1.2 Different DNS Response Types

DNS servers can respond to queries with 10 different response codes or types. These range from protocol implementation errors to server timeouts. Nine of the 10 response codes indicate an error or inability to provide a result for the query. The single successful response code, 0: NOERROR, is returned for all queries that have been processed successfully *and* where results are available. The type of responses originating from censorship infrastructure will depend on the methods used to filter the traffic. Commonly, queries for blacklisted domains are responded to with either a 3: NXDOMAIN—the domain does not exist, or 0: NOERROR—with an incorrect record. Some other censored servers will simply drop the request resulting in a time out on the client where no response is received within a given time period. Table 3.1 shows the different result types that can be received and the causes for them.

Table 3.1: DNS response types when resolving an IP address for a domain

Response Type	Cause(s)
Error	<ul style="list-style-type: none"> » <i>Query was malformed</i> » <i>DNS server is unable to process the request</i> » <i>DNS server has purposefully responded with an error</i> » <i>Domain does not exist (NXDOMAIN)</i>
Request timeout	<ul style="list-style-type: none"> » <i>DNS server is offline or overloaded</i> » <i>Network is unreliable</i> » <i>DNS server has purposefully dropped the request</i>
Address (A) record returned	<ul style="list-style-type: none"> » <i>1 or more correct IP addresses for the domain</i> » <i>1 or more incorrect IP addresses for the domain</i>
No records returned	<ul style="list-style-type: none"> » <i>DNS server cannot make recursive lookups</i> » <i>DNS server purposely did not respond with records</i>

As we take DNS measurements during censorship research, it is important to note what kind of response is received. This, and the response data, are what give us the ability to determine if a particular DNS query or server is censored by an external actor. Given this, we must be able to obtain measurements from known, *uncensored* services in order to make comparisons. During this work, we largely utilised public DNS services that are operated by large organisations for public use.

These services are not known to be censored by any actor, however a future research avenue could be to lead a study on these services and identify how they differ (if at all) in the results they provide.

3.2 DNS Hijacking

Hijacking of DNS queries or servers can be achieved in many ways. The first and most common, is to simply operate or manipulate the server that users will send DNS queries to. This is often done in regions where ISPs are under government mandate to censor certain material from the Web. Turkish ISP *Türk Telekom* does exactly this [185] and will automatically set the DNS servers of its users to the company's filtered DNS service, and Indonesian ISP, *Telkom Indonesia*, also operates censored DNS servers that will block access to blacklisted sites.

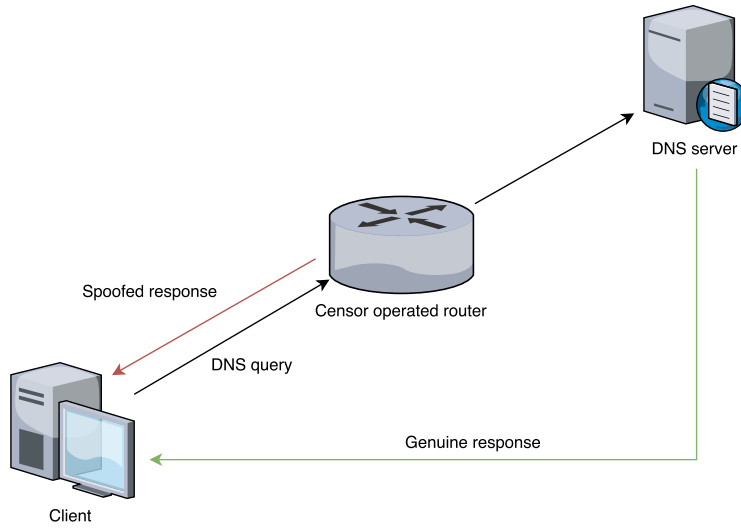
For the vast majority of internet users, this type of filtering will stop access to ostracised sites. Technical knowledge into how DNS blocking is achieved and how to circumvent it is not commonplace. The main methods used for bypassing internet censorship are *VPN tunneling* or use of tools such as *Tor* [173] or *Psiphon* [157]. A key feature of all of these approaches is that they protect DNS queries in layers of encryption so that they cannot be read or manipulated. Moreover, they will often be configured to anonymise DNS requests so that the original requester cannot be identified by the operator through a simple geo-IP lookup. Most Tor implementations will, by default, tunnel all DNS queries through the Tor network - thereby protecting them from adverse manipulation by censors¹.

If a censor chooses not to manipulate the DNS services themselves, they can perform a *man-in-the-middle* or *man-on-the-side* attack shown in Figures 3.1a and 3.1b respectively. This is an approach where a censor will operate machinery that passively monitors DNS requests that pass through key routing points or nodes on a network [47]. If a query for a blacklisted domain is identified, an action may be taken to attempt to halt access to it. Commonly, this will be done by spoofing a response to the query with an incorrect IP address or an error. Importantly, the injected response will be transmitted to the client as soon as the query passes the censor's node. This means that the spoofed response will likely be received by the client before the genuine response down simply to timing and distance (i.e. the injected packets will be en-route to the client before the request reaches the destination DNS server). Further, if key routing points are used as attack surfaces, the number of hops between the client and the router compared to the client and the DNS server will be lower, therefore the spoofed response will often be the first received.

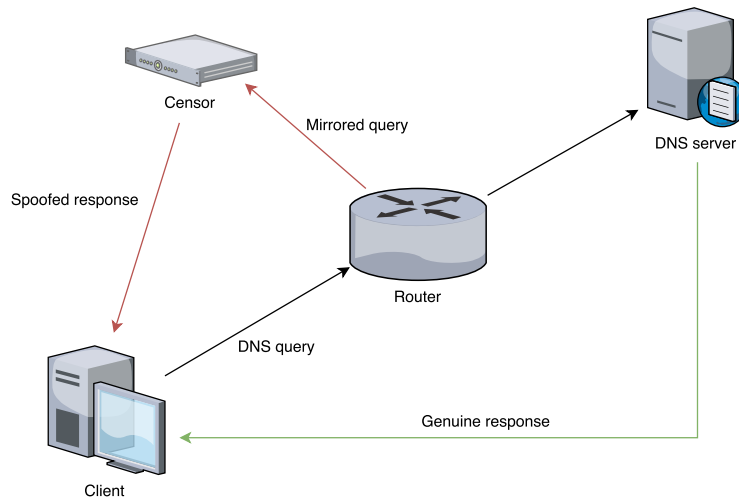
¹Of course, the validity of DNS responses is not guaranteed, since they will be resolved by the exit node. This means that if that node is in a censored region, a hijacked DNS response can still be received.

Figure 3.1: Direct DNS query interference

(a) Man-in-the-middle attack



(b) Man-on-the-side attack



Many networking stacks will track the incoming and outgoing packets to ensure duplicates are dropped. Given that spoofed packets often arrive first, this means the genuine ones will not be processed by the DNS implementation. We can demonstrate the way DNS responses are spoofed by making a query for a censored domain and a non-censored domain in a region where packet injection occurs. Figure 3.2 shows a simple experiment to obtain records for *baidu.com* and *twitter.com* using a DNS server located within China (61.155.18.36). To capture the query and responses, Wireshark [178] was used to sniff networking traffic from the test machine. In 3.2a, the query for *baidu.com*, we see only a single response (the second line in the list). Whereas in 3.2b, for *twitter.com*, we see four different responses - and 3 different IP addresses given for the queried domain. This shows that at least two responses for this query were spoofed by censorship infrastructure in China, however, all the received IP addresses for *twitter.com* appear to be incorrect—which implies that the destination DNS server was manipulated or poisoned. Importantly, the responses are all identical except for the differing IP addresses, this includes the DNS transaction counter which ensures the client will accept the injected responses as they will be read as genuine. Another important feature of this hijacking is that *multiple* spoofed responses are received per query. This suggests that the DNS request is handled by censorship infrastructure at several points along the route to the destination server. At each interception point, a fake response is generated and sent to the client.

Figure 3.2: Wireshark traffic capture when making queries to a DNS server in China

(a) DNS query for baidu.com

No.	Time	Source	Destination	Protocol	Length	Info
10	2.894727	10.207.121.58	61.155.18.36	DNS	69	Standard query 0x1691 A baidu.com
11	2.358986	61.155.18.36	10.207.121.58	DNS	133	Standard query response 0x1691 A baidu.com A 188.149.132.47 A 123.125.114.144 A 111.13.101.208 A 220.181.57.217

(b) DNS query for twitter.com

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.207.121.58	61.155.18.36	DNS	71	Standard query 0xc244 A twitter.com
2	1.067243	61.155.18.36	10.207.121.58	DNS	98	Standard query response 0xc244 A twitter.com A 93.46.8.89
3	1.067785	61.155.18.36	10.207.121.58	DNS	87	Standard query response 0xc244 A twitter.com A 159.106.121.75
4	1.067714	61.155.18.36	10.207.121.58	DNS	87	Standard query response 0xc244 A twitter.com A 93.46.8.89
5	1.067716	61.155.18.36	10.207.121.58	DNS	98	Standard query response 0xc244 A twitter.com A 46.82.174.68

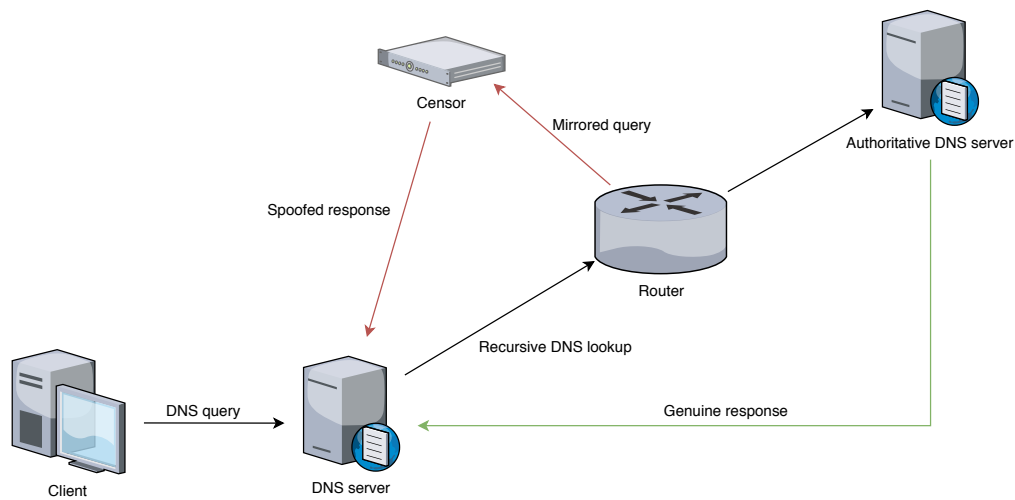
3.3 DNS Cache Poisoning

The majority of DNS servers use a cache to improve response times for users. When an unknown domain is queried, the server will make a recursive lookup to an authoritative server higher in the DNS hierarchy. Ultimately, the root server for the region will respond with the authoritative answer,

this can take numerous steps depending on the depth of the hierarchy. In order to reduce the latency for user queries, servers will store the answers for a short time¹—Time To Live (TTL)—specified by the original answer from the authoritative server. Any query for a cached answer will then be responded to within the TTL of the original request.

If a DNS server is operating with, or connects to, other services part of a censorship infrastructure, there is a high chance the cache of the server will be *poisoned* by an incorrect or erroneous result. This occurs when a DNS server makes a recursive lookup to an authoritative server for a locally unknown domain. The recursive query is treated by all infrastructure it passes (including other DNS servers) as a standard DNS query for a given domain. As such, if the query passes through censorship infrastructure, it can be intercepted as any other normal user query. Furthermore, if the end-server is hijacked or part of a censorship infrastructure too, the resulting answer could also be filtered or manipulated. The answer will then be stored in the cache for its TTL. This process can also affect multiple DNS servers in the route from low-level to authoritative server leading to the poisoning of numerous DNS server caches. Figure 3.3 shows how this process can occur within a censored region where the actor implements a *man-on-the-side* against DNS queries.

Figure 3.3: Poisoning of DNS server cache



Poisoning DNS caches is in active use by several censors worldwide. Farnan et al. found in 2016 that Chinese censorship infrastructure actively attempts to poison the caches of other DNS services to block access to blacklisted domains [54]. This shows that censorship of the Chinese internet is targeted towards underlying DNS infrastructure in addition to user queries.

¹This is commonly set to a default of 300 seconds or 5 minutes

Importantly, the outcome of such a system is that even if a DNS server is *outside* of a controlled or censored area within the country, a domain can still be filtered by hijacking any query that enters the authoritative region and thus poisoning the caches of outside services. This allows for a country like China, which has infrastructure spread over a large geographic area to continue to successfully block and filter domains through the features of recursive DNS lookups.

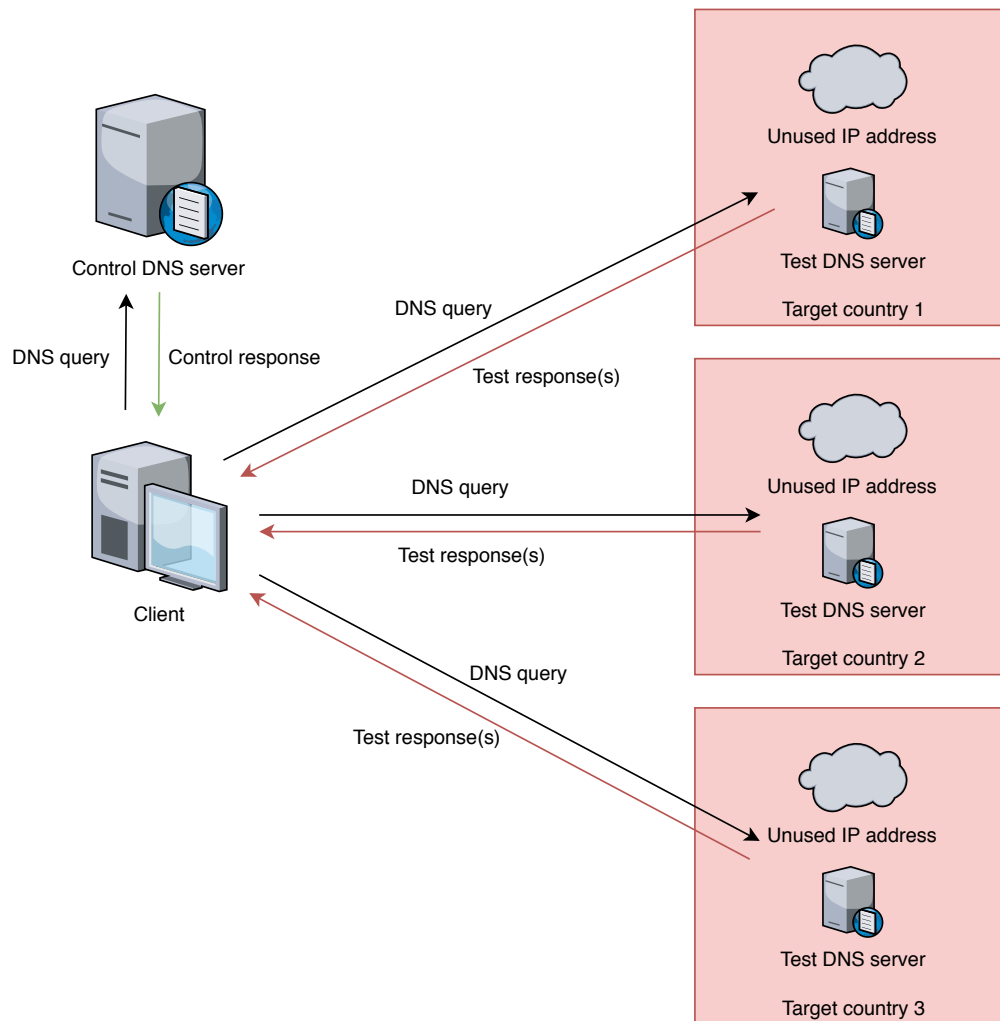
3.4 An Instrument for Filtered Domain Measurements

We define a multi-stage method for determining if a particular domain is filtered within a target country from a remote vantage point. This process uses two of the key features that allow DNS to be used to block domains by censorship actors: *lack of secure transport (encryption)* and *no server-side authentication*. Due to these, censorship infrastructure is able to *inject* or *manipulate* DNS responses in existing streams and *hijack* existing DNS services. We can use these actions to identify when DNS results are manipulated by using a “safe” control DNS service that provides non-censored results. Through comparing the “clean” answers with the potentially manipulated results, one can take precise measurements of censorship activity from an external vantage point. The method requires three key pieces of information in order to work:

- IP address for a control DNS server that is not filtered and returns true results
- IP address for a live DNS server within the target country
- A non-DNS server IP address (*with no host associated with it*) that is within the IP range for the 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, acting rogue or the query is intercepted by other censorship infrastructure. Figure 3.4 shows how the process is executed for multiple target countries.

Figure 3.4: Process for determining filter status for a domain in specific target countries



3.4.1 Finding Evidence of Censorship in DNS Responses

The responses from each test query are analysed for filtering activity. For each domain we are testing per target country, we make six checks to ascertain the filter status of domains and the means by which the target country applies the filter. Each of the following checks must pass in order for the domain to be considered unfiltered:

- (1) The DNS query receives a response from the target country when sent to a non-DNS server IP address

Result: **Query intercepted**

- (2) Measurement server does not resolve an IP address, control server does resolve an IP address

Result: **Server hijacked and/or manipulated**

- (3) Measurement server responds with a private IP address, control server responds with non-private IP address

Result: **Query intercepted and/or results manipulated**

- (4) Measurement server resolves an IP that times out on an HTTP GET request, whereas the resolved IP from the control server does not

Result: **Query intercepted and/or results manipulated**

- (5) Control server resolves an IP that times out on an HTTP GET request, whereas the resolved IP from the measurement server does not

Result: **Query intercepted and/or results manipulated**

- (6) Content length of the webpages from each resolved IP differs by more than a defined percentage amount

Result: **Query intercepted and/or results manipulated**

If the results from any of the above checks are found to be positive, we consider the domain to be filtered by the measurement DNS server or blocked through DNS interception. The complete process for the filtered domain measurement method is shown in Algorithm 3.1.

One of the key features of this algorithm is the use of inactive IP addresses within the target countries. If a response is received to any query to these, this is indicative of filtering activity since

it would have originated from filtering infrastructure: the query was intercepted and a poisoned result returned. If we receive a response to a query, we therefore mark the domain as filtered. If a response is not received after a given timeout, the filtering infrastructure was not triggered, and thus the domain is not marked as filtered (subject to a pass on the other 5 checks).

Algorithm 3.1: Pseudocode for Domain Filtering Check

```

resolveDNS  $\leftarrow$  function(dom, nameserver) {NULL on timeout}
httpGETRequest  $\leftarrow$  function(ip) {NULL on timeout}
isPrivateIP  $\leftarrow$  function(ip)
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 will be NULL on timeout}
1: mFakeIP  $\leftarrow$  resolveDNS(dom, mDNSFake)
2: mIP  $\leftarrow$  resolveDNS(dom, mDNS)
3: cIP  $\leftarrow$  resolveDNS(dom, cDNS)
4: mIPContent  $\leftarrow$  httpGETRequest(mIP)
5: cIPContent  $\leftarrow$  httpGETRequest(cIP)

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

    {check if a record was returned}
9: if mIP  $==$  NULL and cIP  $\neq$  NULL then
10:    return TRUE {no IP returned, mDNS is rogue server}
11: end if

    {check if mIP is private}
12: if isPrivateIP(mIP)  $==$  TRUE and isPrivateIP(cIP)  $==$  FALSE then
13:    return TRUE {mIP is private address, mDNS is rogue server}
14: end if

    {check if mIP points to NULL content}
15: if mIPContent  $==$  NULL and cIPContent  $\neq$  NULL then
16:    return TRUE {bad IP returned, mDNS is rogue server}
17: end if

    {check if mIP points to bad content}
18: if mIPContent  $\neq$  NULL and cIPContent  $==$  NULL then
19:    return TRUE {bad IP returned, mDNS is rogue server}
20: end if

    {check if mIP points to bad content}
21: if length(mIPContent) / length(cIPContent)  $>$  MAXDIFF then
22:    return TRUE {mIP points to incorrect content, mDNS is rogue server}
23: end if

    {check passed}
24: return FALSE {dom not filtered by mDNS}

```

3.4.2 Comparison with existing DNS measurement approaches

There have been numerous previous approaches for detecting Internet censorship that utilise DNS as a primary means for acquiring measurements. The method detailed here has a set of similarities with these, as well as introducing new underlying techniques to improve detection accuracy in some areas.

This method uses similar methods to previous work in the following ways:

- Analysis of DNS query response data where censorship is assumed if an A-record (IP address) is not returned [1][11][13][15][63][112][182][183][184][193][200][202]
- Identification of server timeout for DNS or HTTP requests where the server does not respond after a given length of time (e.g. 60 seconds) [13][63][184][193][202][212]

This method is novel from existing work in the following ways:

- Use of a non-DNS server IP address to detect DNS interference in a robust and accurate manner
- Use of a control DNS server to compare the webpages of hosts (IP addresses) that are contained with the query responses

3.4.3 Advantages

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—that is, infrastructure which has been specifically designed to handle DNS queries and, in the case where censorship takes place, block/manipulate them. As discussed in Section 2.8 is of utmost importance that our testing does not compromise the safety of individuals within censorship regimes [46]. By taking measurements directly from infrastructure, we do not impose danger on volunteers or unaware agents in the target country.

Secondly, it is an efficient mechanism that is scalable and yields faster results than manual (human-based) checks. It is possible for us to test several hundreds of domains a minute using commodity hardware since the DNS protocol is very lightweight¹. Whereas manually testing domains would be limited and reliant on constant interaction from a test agent. Furthermore, we are able to increase our accuracy by using the artefacts of DNS interception since we *only* receive a response to a query (to unused IPs) if it *was* intercepted by censorship infrastructure. This is a

¹DNS response times are commonly between 5ms and 200ms with over 95% responding within 100ms [5][133]. Giving a range of between 300 and 6,000 (nominally 600 with the average response time) possible DNS requests a minute. Even if throttled to 10%, we would yield a rate of 1 request a second, or 86,400 per 24 hours.

highly effective method for distinguishing filtering activity due to its binary nature—*query intercepted or query not intercepted*.

Thirdly, we can perform measurements from outside target countries, using several, remote vantage points, giving us the capability to analyse filtering within a wide array of censorship regimes. DNS is also rarely filtered by any ingress or egress firewalls due to its commonplace in many methods for internet communications.

Finally, the method is sensitive to multiple different DNS censorship tactics that are employed across different countries. We are aware that individual censorship regimes block DNS in alternative ways, for example one may enact via interception and another via re-routing users to block pages (by sending an incorrect IP address). We can use this instrument to detect in many different circumstances if a domain is indeed filtered.

3.4.4 Limitations

We must acknowledge a number of inherent limitations with this technique.

The biggest drawback 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.

We also require that measurement DNS servers respond to queries from remote countries in the same manner they do for queries made domestically—which may not be the case. For instance, a Chinese censor could respond to DNS queries originating in China differently from those originating in the UK (as identified from the source IP address on the query). That being said however, previous research has confirmed that, at the time of writing, detecting Chinese censorship of DNS is possible from remote location [54][91][202].

Whilst DNS censorship is widely practised it is difficult to know definitively if a web page is blocked through sole use of this technique. Many censorship regimes will make use of multiple methods to restrict access to the Internet. Use of alternative testing frameworks, such as ICLab [142] or OONI [171], would mitigate this limitation significantly and allow for more targeted testing. Furthermore, there are other techniques for detecting specific blocked webpages that have been published [55][202]. The use of DNS in our method is due to its scalability, speed and the ethical considerations explained previously. The downside of this is that we will not detect if URLs are filtered by other means, thereby limiting the “resolution” of the discovered filtered resources to a domain / sub-domain level.

Sophisticated censors may provision their infrastructure to specifically evade attempts to detect blocked domains. This could be achieved by mirroring responses from the major open DNS

providers, as described in Section 3.1.1. If this occurred, it would be very difficult to determine if a particular domain was being censored or not. There are also specific circumstances where the inferences, described in Section 3.4.1, which are drawn from DNS responses may be inaccurate:

For (1) A DNS query that is intercepted may actually return the "correct" IP address for the website. This could mean that the domain would technically be available in the censored region (although, as mentioned earlier the censor may just be returning an alternative response for our origin location), however, we still maintain evidence of DNS interference even in this case—since the query *was* intercepted.

For (4, 5) A query could be responded to with an IP address of a "locally available" server or Content Delivery Network (CDN)—this is to say that many internet services use servers located close to their users, to provide faster access speeds—we, from the UK may not be able to access these servers.

For (6) DNS queries from the control and measurement servers could point to different hosts of the website. A valid webpage from these hosts could be substantially different—if the content depends on the locality of the user—thereby creating a false positive (Type I error).

3.4.5 Summary

We to make a trade-off between effectiveness, efficiency and ethical considerations. The possible altercations that arise if volunteer agents are used to test for censorship and measure filtering activity are, in our opinion, beyond the level of morally acceptable research. Measurements must be taken against infrastructure where individuals are not at risk of being compromised. Furthermore, we ensure to only test against infrastructure that is built for the types of query we send. This is to say: we only make DNS queries against DNS services and the DNS filtering infrastructure in place in the target countries are *designed* to intercept or manipulate DNS. No other kind of crafted payloads are used, only regular DNS query packets are sent to any test target.

The methodology presented here extends and combines existing approaches for measuring DNS censorship, as well as introducing novel detection techniques, such that individual domains in any given target country can be tested. Previous to this, DNS measurements used attributes in the query responses to determine the blocked status of a domain. We show that a second, non-DNS server IP address can be used for accurate detection of DNS interference patterns.

This new method for remotely determining if a domain is censored will form an important component of the two system designs for discovering filtered domains described in Chapter 4 and Chapter 5. Both offer new approaches for using a set of known filtered URLs in order to find other

(previously unknown) filtered domains. Through experimentation of these, we will gather a better understanding of the reliability, efficiency and efficacy of this censorship measuring instrument.

“Don’t be evil”

Paul Buchheit, 2000,
Google.
At a meeting on company values

CHAPTER 4

A Framework For Modelling Characteristic Linguistic Connections Between Filtered URLs

Filtering of individual webpages within a site or domain, based on their content, is a well-known and practiced activity. The goals of this kind of censorship can vary and depend on the given regime. However, with the rise of social media and platforms for individual expression, the determination for certain censors to gain control over this area is increasing.

Currently, the methods, by which the censorship of URLs and domains can be discovered for any particular country or Internet region, are limited. There is a need for more efficient processes for building larger and more extensive URL block lists. These are vital datasets for almost all censorship research since we often require a starting set of censored material, resources or websites in order to test new monitoring functions and country specific filtering. In this Chapter, we explore a new approach for building more diverse and longer filter lists of Internet domains per country. Similar to most previous work, we also require a starting, or *seed*, list of already known filtered URLs to bootstrap the system. However, with this, we aim to increase the length of the starting lists of confirmed blocked web resources by at least an order of magnitude.

As explained in Section 2.5.3, the blocking of singular webpages is a technological challenge in of itself. It requires more complex and powerful infrastructure to maintain a good service, also, with the speed and throughput that many modern web services operate at, a censor must match this capability if they are not to increase the latency of use—the Chinese filtering infrastructure contained over 490¹ Intrusion detection system (IDS) instances in 2011 [207]. Furthermore, many large-scale services are geo-distributed across different regions of a country, for performance and

¹The number of IDS devices as of 2019 is likely much higher than this.

fault tolerance reasons. This provides extra challenges for a censor since their own infrastructure must also cover these localities. Given these requirements and potential limitations, a common mitigation is to block entire domains or IP address ranges using exercised control over DNS servers and key internet routers respectively. Filtering of this type is often used when websites or services contain overarching levels *sensitive* material or the genre or site category is undesired as a whole (common examples being religious sites, pornography or news sites).

By analysing language patterns on webpages, we can use the traits of censors to model what kind of content is being blocked in a certain jurisdiction and thus use the model as a tool to discover filtered URLs. As a consequence of *linguistic connectivity* [186], we hypothesise that the language patterns apparent in *one censored page or domain* could be used to link a *second censored page or domain*. This is a key assumption of the approach for discovering filtered URLs that is tested during this chapter.

4.1 Text Mining in Webpages

Deriving detailed and useful information for textual data is an important and highly active field. Capabilities to gain insights, parse, store and classify unstructured text are highly sought after by businesses, governments and academics alike. The demand for this kind of processing is due in part the sheer amount of textual data that modern businesses need to deal with — upwards of 80% [166] — and also because of the inherent preterition of clear structure within the raw data. This lack of clear mark-up makes processing raw text a challenging problem within computer science and Artificial Intelligence (AI). Natural language processing (NLP) is the study of interactions between computers and human language with an aim for machines to be able to process language corpora. A major part of this area is Natural language understanding (NLU) where the aim is to build systems and algorithms to allow for machine comprehension of natural language — with a focus on identifying the semantics of text. Recent advancements in this field utilise deep-learning techniques which have produced impressive results for language parsing and modelling especially [62]. Examples of the use of text mining outside of traditional Computer Science fields include patent analysis [191], biology [9] and medicine [195] among others.

Textual data on the Web is similar in order and structure to raw text in many of the above examples — in that it lacks any. While it's true that most text on the Internet is delivered in HTML, which is a formatted document, the text content is still simply made up of blocks of language. Our aim is to process this raw data so we can build *short summaries* of larger pieces of text to describe the content on *individual* webpages.

4.1.1 Descriptive Tags for Web content

We use the notion of tags to characterise small pieces of text that describe larger bodies of content. These could be a short descriptions or a list of individual words. Importantly, these are not necessarily keywords that are themselves filtered by a country, but instead linguistic patterns that may be likely to be present on other blocked web pages.

Example 4.1.1


If you own a virtual-reality headset, you've seen a few health and safety rules. Don't use your VR headset in a moving vehicle, for instance, or make sure to take frequent breaks. For most of us, these guidelines make sense: VR nausea is a very real problem, and limiting our time in artificial worlds is the easiest way to avoid getting simulator sickness. But what if you broke all the rules and decided to stay in virtual reality for 48 hours straight – eating, sleeping, working and living in a VR headset? Well, then you'd be Dean Johnson, head of innovation for Bandwidth and crazy man who spent two days blindfolded with technology.

Johnson has been challenging the rules of consumer VR from the beginning – when virtual reality hit the mainstream last year, he spent 24 hours immersed in a mix of Rift, Vive and Gear VR experiences, setting an unofficial record for longest time in virtual reality. This year, he doubled that effort, recruiting Sarah Jones from Coventry University to join him in two days of extreme VR immersion – breaking for only five minutes each hour to record vlogs and use the facilities.

The experiment was designed to question the arbitrary limits of VR-use time and help expose virtual reality to a wider consumer audience, but it wasn't a PR stunt for any specific headset manufacturer. "In fact, it was quite the opposite," he says. Every company he invited to participate in the project turned him down. "Mostly because they thought we'd die," he joked.

Taken from: <https://www.engadget.com/2017/06/07/two-people-spent-48-hours-in-nonstop-virtual-reality/>

Figure 4.1: Tags derived from text in Example 4.1.1



4.1.2 Information Extraction

Automatically extracting structured information from unstructured text is a vital task in text data mining. In order to derive useful features from information one must first process the unstructured data into a computer readable format. An important goal for any information retrieval system is to perform computations on this unstructured text. Tools and techniques for structuring natural language have been the subject of large efforts in academia and business. A key part of this is to be able to perform logical reasoning on data in order to interpret the information in a computable way. In doing so, we can create algorithms to draw inferences from the data to complete a task. One such example of this is natural language search. A well-designed search engine will create logical relationships between the query and bodies of text within the search index. A common method used to do this is keyword extraction and matching. Keywords are contextualised pieces of text that describe part of a larger piece of text (similar to our notion of descriptive tags) and are widely used in modern search.

4.1.2.1 arc90 Readability

The Readability project from arc90 [12] was an experimental approach to extract the main body of text from a webpage. The aim was to distinguish the main article text from the other text present, such as menu items, headers / footers, side bars, etc. By solely identifying the main body of text or paragraphs on a webpage, one can use this as part of an input to an other process or NLP pipeline — for instance, keyword extraction. If other pieces of readable text are used, performing subsequent NLP tasks may incur inaccuracies or inconsistencies. A common issue is the derivation of terms from a side bar text during keyword extraction, such as one for “*Similar articles*” or “*Recommendations*”. This could include keywords (from titles or links to other pages) for data not associated with the main article or body of text for the given URL one is processing.

4.1.2.2 Term Frequency — Inverse Document Frequency (TF-IDF)

TF-IDF [22][119][181] is a statistical method of determining how important a given word is within a document. It uses a corpus (for the given language) to offset the frequency of a word as it appears in the document. Essentially, the weighting for a particular term becomes larger proportionally with the number of times it appears in the document, offset by its frequency in the corpus. Taking the n terms with the highest weightings can give us a list of significant keywords that characterize the content of the document. TF-IDF is widely used in information retrieval for keyword extraction due to its proficiency at disregarding common words and favouring the importance of

less frequent terms. Moreover, TF-IDF can be seeded with any corpus of text, this allows for highly domain specific keyword extraction as well as a more generalised model.

4.1.2.3 TextRank

TextRank [98] is a graph-based, unsupervised method for keyword and sentence extraction. This algorithm generates a text graph from unstructured natural language placing “*text units*” as vertices, and then drawing edges between them. *Text units* can be any piece of text comprising of one or more words, sentences or collocations¹. Each edge represents a relationship between two text units, this could be syntax, semantics or contextual overlap. The simplest form of this process will use individual words for vertices and word placements for edge relationship. After the graph is constructed, an iterative ranking model is used to score each vertex for importance in the graph. This process executes until the score converge and there are no further changes—similar to the graph ranking model used for PageRank [109]. The vertices are then sorted by score and the keywords / phrases are derived from them. An important feature and advantage of TextRank is that it is language neutral, in that a corpus is not required. This means that a well designed graph model can be used to pull keywords from any unstructured text regardless of the source.

4.1.2.4 Rapid Automatic Keyword Extraction (RAKE)

RAKE [180] is an unsupervised, domain dependent method for extracting keywords or phrases from unstructured text. It uses a co-occurrence graph to build a mapping of the semantic proximity relationship between individual keywords. Candidate words are derived from the text after removing all stop words. Per candidate word, a score is calculated using the word frequency, the word degree and the ratio between both. To generate phrases, the algorithm will look for high scoring individual words that co-occur or are separated by a stop word and appear more than once. This allows for the dynamic extraction of phrases that may not exist in a training corpus for a supervised algorithm — such as *virtual reality headset* or *health and safety*.

4.2 Recursive Discovery for Filtered Webpages

The current methods for enumerating and identifying censored material on the Internet rely on large amounts of human effort or participation. These systems need constant updates and verification to ensure they continue to operate as required. This often results in old or outdated URL filter lists or such systems becoming non-operational and taken offline. Methods that require human interaction usually produce high quality results, albeit at a smaller scale—due to the

¹**collocation:** noun. the habitual juxtaposition of a particular word with another word or words with a frequency greater than chance

necessity for a large number of volunteers. Further, and discussed previously, any system that uses human participants can increase their potential risk for providing data or performing tasks.

With these factors in mind, a tool that is simple and inexpensive to operate (*in terms of time, money and effort*), doesn't require a constant update schedule and runs automatically, is highly desirable for the censorship research community.

4.2.1 System Overview

We propose a new approach for discovering filtered URLs within a target country. This technique uses large search engines to link patterns of language to web pages in order to find previously unknown filtered URLs. It is impractical and inefficient simply to crawl large portions of the internet for this task; as such, we propose a framework that uses existing infrastructure and services that already build a broad “view” of pages and documents on the Internet. Further, as documents are created or updated these changes will be reflected in these services.

The methods by which we can detect censorship remotely will determine what measurements we can capture, as do the technical means that a particular country uses in order to block access to websites. Different countries of course use different techniques, therefore any implementation of this system should be altered accordingly to the target.

Previous work [40] has shown that censors have used content analysis as a part of their censorship infrastructure, and will block resources or traffic containing *keywords* deemed to be sensitive. The aim here is to leverage this method in a similar manner in order to quickly identify blocked content and thus increase transparency on censorship activity. Known censored webpages are used as *seeds* to start the search, from each of these, we extract *descriptive tags* to use as the basis for search queries. This gives us a broad range of keyword terms that are semantically related to the content on each filtered webpage. The use of search engines allows us to exploit the linguistic, semantic and structured “links” that are created between search query and webpages on the Internet. An important purpose of search engines is to *rank* and *order* unstructured textual data so that they can be discovered via natural language queries. By using these services—that are developed, maintained and distributed by 3rd parties—we can perform highly accurate discovery tasks without a large infrastructure or development requirement. Furthermore, as the 3rd party search engines improve and add new documents over time, our system will also benefit without further internal development.

A key feature of this proposed framework is the recursive element: as more censored webpages are identified, we produce more *descriptive tags* that are used for further web searches. Importantly, this allows the system to continue to find censored material in an on-going fashion—second-order,

third-order and so on. Blocked webpages can be discovered numerous “hops” from the original *seeds*. A downside of this is the fact that we *rely* on both the quality of seed URLs and that search engines are producing accessible “links” between different filtered content. These issues can be partially mitigated through the use of well researched filtered URLs to seed the system.

4.2.2 Methodology

The framework does not aim to derive blocked keywords, which is a separate area of study, but instead identify terms that are likely to be shared by web pages that discuss key sensitive topics. The underlying assumption is that blocked pages publish patterns of language or key phrases that other filtered pages may themselves contain.

To begin the process, the system must be seeded with known blocked URLs for a target country T . For each of these, we download the web pages and extract the text from them - discarding any HTML code, scripts or CSS. We then extract *descriptive tags* from the body of text that will then be used as web search queries.

A web search will be made for each descriptive tag and the resulting candidate URLs will be stored and checked for evidence of filtering in target country T . The system is recursive and results are used for further discovery: for each newly discovered filtered URL, the tag extraction and search processes are run again to gain further candidate URLs for the next iteration. This framework is summarised in Figure 4.2.

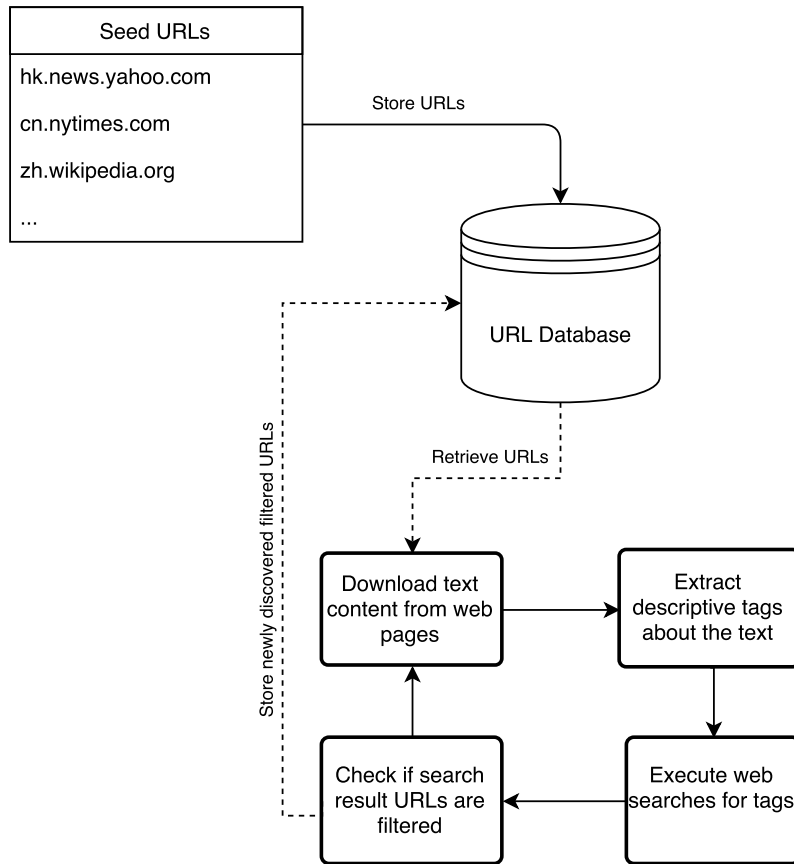
Further, given the dynamic nature of both web content and the behaviour of censors, any implemented approach could continually re-check filtered URLs for changes to acquire results over a longer time period. This is especially important for news sites that are continually updated and refreshed. Pseudocode detailing the approach is shown in Algorithm 4.1.

The intuition behind this approach is that filtered web pages may contain patterns of language that can then be used to discover new filtered URLs. Importantly, this framework does not identify filtered keywords, but instead uses key terms (tags) that are likely to be shared between sites discussing similar or according topics. It is important to note that this framework is modular and cleanly separated into three distinct parts:

1. **Tag extraction:** the method by which descriptive tags are derived from a body of text.
2. **Web search:** the search engine(s) that are used to find URLs that are related to descriptive tags.
3. **Filtering check:** the method for checking if a given URL is filtered in a certain country.

Given this, measurements of URL filtering can be made for alternative target countries through different means depending on requirements. Furthermore, there are various existing tools and

Figure 4.2: High-level overview of the recursive filtered URL search



Algorithm 4.1: Pseudocode sample for the framework

```

urlDB ← new Database(seed_urls)
tagDB ← new Database()
isFiltered ← function(url)
getTagsFromWebPage ← function(url)
doWebSearchForTag ← function(tag)
loop
  for url IN urlDB do
    if isFiltered(url) == TRUE then
      tags ← getTagsFromWebPage(url)
      ADD tags TO tagDB
    end if
  end for
  for tag IN tagDB do
    urls ← doWebSearchForTag(tag)
    ADD urls TO urlDB
  end for
end loop
  
```

platforms that provide functionality to test if certain web resources are blocked in different countries.

The framework is scalable since it uses openly-available and cost effective search engine APIs in order to find new candidate URLs. This means that we do not need to build or maintain large amounts of infrastructure to crawl the Internet because this requirement is fulfilled by the search engine. Further, the more specific we can be with our search queries, the more accurate and targeted results we will achieve.

The implementation of this framework uses the DNS method as a means for checking if URLs are filtered, described in Section 3.1. However, the censorship checking component is entirely modular and separable to the other parts of the system. One could therefore make use of alternative tools, data, or sensor networks such as GreatFire.org [64], OONI [55], ICLab [142] to take censorship measurements as part of the process.

4.2.3 Process for Extracting Descriptive Tags

To isolate descriptive tags, the documents downloaded from filtered web pages are first cleaned of any non-readable parts including HTML code, Javascript and image/binary data. The remaining text is then tokenized, and each word weighted using TF-IDF. This results in a ranked list of words that best characterise the content of those pages in contrast to typical text in that language. At this time we only consider words that contain letters from the ISO basic Latin alphabet.

According to best practice we remove the 1000 most common English words using the google-1000-english list [56]. The aim of this is reduce the number of tags that are too common and would yield generic web search results although in future implementations we can alter this to achieve a better result. The process used for extracting *descriptive tags* is shown in Algorithm 4.2.

Algorithm 4.2: Descriptive tag extraction process

```
removeHTML  $\leftarrow$  function(rawhtml) {removes all HTML, CSS & Javascript code}  
tfidf  $\leftarrow$  function(text) {returns list of tags}  
istop1000  $\leftarrow$  function(word)
```

```
1: tagList  $\leftarrow$  new List()  
2: rawhtml {rawhtml downloaded from URL}  
3: text = removeHTML(rawhtml)  
4: for tag IN tfidf(text) do  
5:   if istop1000(tag) == FALSE then  
6:     APPEND tag TO tagList  
7:   end if  
8: end for
```

4.2.4 Use of Existing Search Engine Infrastructure

As previously mentioned, we make extensive use of existing search engines and APIs that provide natural language search. We want to leverage the research and development effort that is expended in producing high quality search engines, along with the vast infrastructure required for crawling the billions of documents that constitute *the Web*. For the purposed of censorship research, we want to be efficient in our discovery. This means preventing our system from crawling large amounts of content on the Internet unnecessarily. We also want the process to identify filtered webpages as quickly as possible after they are published or blocked respectively. Although most search engine companies don't reveal details on how they organise or order search results, it is generally believed that "newer" content will be ranked higher, especially if it relates to recent news storied [24]. This means that we can incorporate the latest content into our own results with a high degree of relevance to breaking news.

4.2.4.1 Seeding the System

We need to seed this process using a set of known filtered URLs for the country we are working to identify other censored content in. Our seed list can be obtained via existing URL filter lists and detection tools or we can manually curate a small list via censorship reports and other censorship research pieces. The important factor is the used of *high quality* seeds. We need URLs that:

- Report as filtered via our chosen detection method (see **Filtering check** above)
- Contain textual data that can be used to derive the initial set of *descriptive tags* (URLs that point to images will not be useful here)
- Are filtered because of the content they contain rather than for other reasons (such as blocked keywords in the domain or URL path)

4.3 Chinese DNS Blocking Experiment

To determine the effectiveness, efficiency, and validity of this approach, we implement the framework into a tool that can be used to find domains that are actively filtered in China. This test is designed to provide insight into how we can use this kind of approach to produce constantly updated URL filter lists that can then be used for other censorship research and transparency of filtering activity. We choose China as the target country for this due to the known use of DNS as a filtering method and the large amount of already known blocked content and URLs. This gives us a good base for starting the recursive discovery. However, it is important to know that this is framework is general and can be applied to any target so long that we have a reliable filter check method for the given target.

4.3.1 Implementation

We produce as implementation following the framework using the largest search API provider at the time of writing—see Section 4.3.1.1 below. The test implementation, filter check and tag extractor were written in Python. All tests were performed from a single virtual machine node.

4.3.1.1 Web search

We use the Bing search engine, exposed through the Azure Cognitive Services API [97], to conduct queries for descriptive tags. This is important because we aim to exploit the sorting and relevance algorithms search companies implement to find appropriate candidate URLs for filtering measurements. Web search is a large and complicated business; most engines do not simply rank pages based on hyperlinks, but rather current trends and activity. This is beneficial because not only and web pages indexes for content, but interest is tracked by surveying the users of the engine. This gives us the ability to isolate sites and pages that are currently of particular interest to the public at any given time - which in turn could lead us to material that is on the forefront of filtering activity. At the time of experimentation, the Azure Cognitive Services API was the only available service to execute large scale web searches. Since then, several competitors have emerged which, if the experiment is repeated, would offer a wider search engine service choice for use.

An alternative to Bing is Common Crawl—an open data project that scrapes the web for pages. While this can give us more control over the search process, the project does not provide methods for processing, sorting or querying the data. Whilst some open-source search engine systems are available, *ElasticSearch* for example, the complexities and administration required to operate such a service for the entirety of the Common Crawl corpus (many petabytes of raw data) would be significant and costly. Hence why we aim to leverage the power of existing engines, with built-in discovery, aggregation and ranking algorithms, to find relevant links.

A third option is Baidu—a Chinese operated search engine, however, it is known that results from this service can be filtered; as such, this will reduce the number of filtered URLs we can find due to less being returned by the engine. A study by Jiang et al. showed that Baidu tends to drive traffic to well-known, major sites and its results raise questions about its impartiality [77]. Furthermore, Baidu will often only direct users to site with China even if a more relevant site, given the query, exists outside of the Chinese IP range.

4.3.2 Parameters

The implemented tool has numerous parameters that can be set before commencing the recursive discovery. Through preliminary testing of the tool, we identified a set of useful parameters as a starting point. After having run the experiment, we will evaluate these choices and determine any learnings or possible improvements for any further experimentation. As with any system like this, these parameters could likely be optimised for future implementations. To summarise, an overview of the parameters for the system; we use:

- TF-IDF for isolation of descriptive tags
- The top 5 descriptive tags from each filtered web page
- Azure Cognitive Services API to perform web searches
- The top 50 search result URLs from each web search
- DNS filtering measurements to determine if a domain is filtered—Section 3.4
- 14 day experiment time

We initialise our tool with URLs taken from the Citizen Lab's URL test list [29] for China. Out of the 204 URLs present in the list, we find that 44 of the domains are DNS filtered and these are used to seed the system.

We chose to use the top 50 results for each web search, this figure was chosen due to the nature of search engine rankings and click through rates (CTR). We find that the ranking algorithms optimise for the first 10 results and the vast majority of all result clicks for search engines occur in the top-10, with an exponential drop-off afterwards [74][131]. Furthermore, an analysis of organic CTRs in 2019 [128] showed that only 0.78% of search engine users clicked on any link beyond the second page of results, where each page contains 10 individual results. Given we know that search engines optimise for these initial results, it stands to reason that we can perform this experiment at a cut-off of 5 search pages, or 50 results, which captures over 99% of all result clicks by real users. It should also be noted that we only use the *organic result listings*, paid-for advertisements are *not* included when we capture the search results.

4.3.3 Results

We conducted approximately 54,000 web searches during a 14 day experiment which yielded over 2,500,000 search results total. From these, we found 1355 filtered domains and 115,337 individual filtered URLs in China. In total, our system crawled 1,113,653 unique URLs and 329,575 domains.

Table 4.1 depicts the number of filtered URLs and domains that were discovered and the hit-rate of the tool where we calculate the number of filtered domains discovered per 1000 URLs crawled. The rate of filtered URL/domain discovery is an important metric to consider, it offers a quantifiable

data point that can be used to evaluate the efficiency of the overall method. Currently, the alternative (published) censorship detection approaches in the literature do not offer this metric, so efficiency comparison between them is difficult. Nevertheless, we can see that roughly 10% of the URLs and 0.4% of domains we visited were indeed censored. This is consistent with the censored domain breakdown shown in Figure 4.3a, which contains 95% of all the filtered URLs visited from just these 15 domains. Outside of the Alexa top 1000 global domains, we find a greater spread of domains in the filtered URLs we visited - shown in Figure 4.3b.

It is important to note that when counting filtered domains, that each domain and sub-domain is counted separately as they may have different DNS entries. This is important because DNS filtering in China occurs on a subdomain by subdomain basis, *secure.example.com* may be filtered while *www.example.com* is not. Furthermore, we have counted all Tumblr pages as a single result due to the fact that there seems to be a blanket block on all sub-domains (Tumblr works by providing each of it's users with a distinct sub-domain for their page).

Table 4.1: Discovered filtered URLs

Metric	Counts
Searches executed	54,000
Unique URLs visited	1,113,653
Unique Domains visited	329,575
Unique Filtered URLs	115,337
Unique Filtered domains	1355
Filtered URLs / 1000 URLs visited	103.57
Filtered domains / 1000 domains visited	4.11


4.4 Evaluation of Approach

Our testing demonstrates that the framework we propose is sufficient in discovering a significant number of filtered domains that were not present in initial seed lists. These results indicate that there is an exploitable connection between the content of individual filtered web pages and other filtered domains. Further, existing search engines can demonstrably be used as tools to uncover censored material on a large, automated scale.


An advantage of this approach is that *only* filtered URLs are used as the basis of linguistic patterns for search requests. This increases the efficiency of the tool and prevents unnecessary crawling of large portions of the Internet to discover newly-filtered URLs.

Figure 4.3: Filtered domain distributions

(a) 15 most frequently occurring filtered domains



(b) 50 most frequently occurring filtered domains with Alexa top 1000 removed



The system has discovered a large number of domains that are currently being filtered in China. The hit rate shows that the tool has to crawl approximately 1000 URLs to find 4 *filtered* domains. In this instance, the approach could have been optimised by limiting the crawling to domains that were not already stored in the database. For this experiment however, it was necessary to crawl larger numbers of pages in order to fully test the usefulness of the framework. In a future version, these optimisations could increase the efficiency and reduce the cost of further discovery. Furthermore, we could start to reduce the crawls of larger more well-known sites—such as facebook.com and twitter.com. This would lower the total number of URLs crawled by a substantial amount, however, it is not clear if these social media services actually benefit the system by providing very new and constantly changing content—further studies into these mechanics is a potential avenue for future research.

In comparison to lists available via the Citizen Lab—shown in Table 4.2, we have found significantly more filtered domains. Given our tool checks for DNS filtering, we can have a high confidence in the numbers of blocked web pages we have found.

Table 4.2: Comparison with alternative filtered URL lists for China

	Recursive Discovery	Citizen Lab [29]	CensMon [184]
Filtered URLs	115,337	204	N/A
Filtered URLs - Top 1000 Removed	4153	166	N/A
Filtered Domains	1355	83	176
Filtered Domains - Top 1000 Removed	759	68	N/A

The list of censored domains discovered by *CensMon* is not publicly available, therefore we only have the number reported.

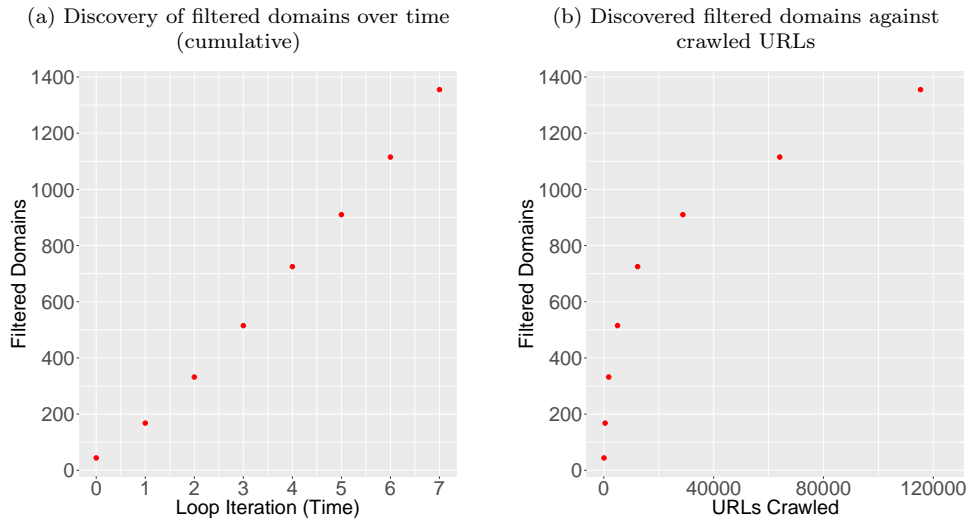
4.4.1 Efficiency of Discovery

Our tool is recursive which means we can quantify the rate of filtered domain discovery over time. This metric is useful in determining the effectiveness of the approach since we can compare the number of domains crawled against the number of filtered domains found. Figure 4.4a shows the number of filtered domains discovered over time, or each loop of the system; from this we can see that the discovery rate is relatively constant. Figure 4.4b shows the number of discovered filtered domains against the number of URLs crawled. From this we can see that the number of filter domains discovered per URL visited tapers off to form the start a plateau after the seven rounds we executed. This shows that the rate of discovery in terms of URL visit efficiency is initially very high, where we can visit relatively few censored sites to find many other sites that are also filtered. However, as the system runs into later rounds, this rate begins to drop off. By implication, we would need to crawl an exponentially higher number of URLs per round to maintain the same rate of discovery of filtered domains—or at least until every possible domain has been visited. Furthermore, due to the nature of this method, we will only ever find censored sites that do connect to our initial seed list—a limitation of these types of “snowball” techniques. This means that using alternative seed lists may result in different sets of end results. Though, it should be noted, that the seed lists used in our experiment was highly interconnected¹—as shown in Figure 4.7—which suggests that if we had started with a subset of these seed URLs, we could result in a similar end set of filtered domains. This having been said, the rate of discovery may well have been lower, at least initially in the earlier rounds, until we had more filtered sites to crawl.

In comparison to *CensMon* [184], which also reports results from an experiment ran over a 14 day, our recursive discovery machine visited 23 times more URLs (115,337 vs 4950) and found over seven times more filtered domains in China (1355 vs 176). Whilst this represents a significantly larger amount of work and processing of individual URLs, we can show that this does lead to a greater

¹Indirectly, they are connected through the language patterns used as search terms

Figure 4.4: Performance of filtered domain discovery



number of discovered filtered domains. However, this being said, *CensMon* used 174 individual agents across 33 countries (using the PlanetLab platform) to conduct measurements, this contrasts to the single remote agent used in our approach which would be notably more efficient to operate over a longer time period and likely at a reduced cost.

Out of the 329,575 domains we visited, 328,220 were *not-blocked* in China and 1355 were *blocked* (as shown above). Of these non-blocked domains, 94,168 appear in the *Majestic Million*¹[146] and of the filtered domains, 352 appear. From this we can calculate that the probability of selecting a filtered domain at random from the Majestic Million is 0.035% and to select a non-filtered domain, 9.4%. Whilst we can not have an exhaustive coverage over all Chinese filtered domains, it is easy to see how a “randomly” selected seed list (from the Majestic Million) would likely not yield good results here, again reiterating the importance of choosing a strong list of seed URLs for the system. Furthermore, the usefulness of censorship filter lists, gathered either manually or with discovery tools like this one, are fundamental for censorship research since we always need a starting point. Exhaustive lists of domains that exist are available publicly (over 350 million domains are registered as of 2020), however, frequently measuring and testing each domain for filtering activity is likely not a practical exercise for many researchers in the field. Therefore, more efficient methods for gathering lists of censored web content are of great benefit. The approach described here is significantly more effective than simply selecting random domains to test for evidence of censorship.

¹The Majestic Million is commonly used database of domains that is compiled and ordered based on which has the highest number of backlinks. The list contains 1 million unique domains and is updated daily based on web crawl data.

4.5 Further Analysis of Results

Further analysis of the database of URLs collected by our tool reveals a number of interesting artefacts about its operation, and the state of censorship in China. Firstly, we look at the most commonly occurring filtered domains found, from this we can see the larger and better known sites and social media services make up a large portion of the URLs processed. Approximately 95% of all filtered URLs found are from just 15 domains—these are shown in Figure 4.3a. Due to the fact that domains within the Alexa Top 1000 generally rank higher in search results and are the most frequently occurring within our database, we also produce Figure 4.3b that shows the 50 most common filtered domains outside of the Top 1000. These domains appear far less regularly in search results, but importantly, they can give us a deeper insight into what the Chinese government is currently blocking.

4.5.1 Enumeration for additional filtered suffixes

We run a scan over additional suffixes for each of the discovered filtered domains found by the system. The aim of this is to uncover evidence of further filtering through association by domain. Generally, a single company or organisation may own numerous top-level-domains (TLDs) for a one or more of their main domains - for example *google.com* and *google.co.uk* are both owned by Google Inc. However, this is not always the case since alternative TLDs of the same domain name may be held by different entities - for instance, *archive.org* and *archive.co.uk* are possessed by individual companies. If we know that *example.org* is filtered within China, we use this enumeration to discover if any other suffixes of *example.** are also filtered incidentally due to the fact *example.org* is filtered.

We use the Public Suffix list maintained by the Mozilla Foundation [57] to perform this scan. This contains all known, publicly available suffixes that can be used for different domains. It contains all TLDs (*.com*, *.org*, etc), as well as less well known or utilised suffixes such as *pvt.k12.ma.us* and *appspot.com* which can still be hosted on different servers or contain content possessed by distinct organisations. Further, as described previously, alternative subdomains can be filtered individually, so we must check them separately. For each filtered domain found previously, we check it with every suffix in the list for filtering within China, then resolve each of these¹ to determine if DNS records exist for that suffixed filtered domain. These results are shown in Table 5.4.

From this process, we discovered an additional 97,167 filtered domains which represents a near 72 times increase on the number of originally discovered domains. Of these, 5408 were successfully resolved with records that pointed to IP addresses. We therefore find nearly 3 times

¹Using Google Public DNS: 8.8.8.8 as a resolver

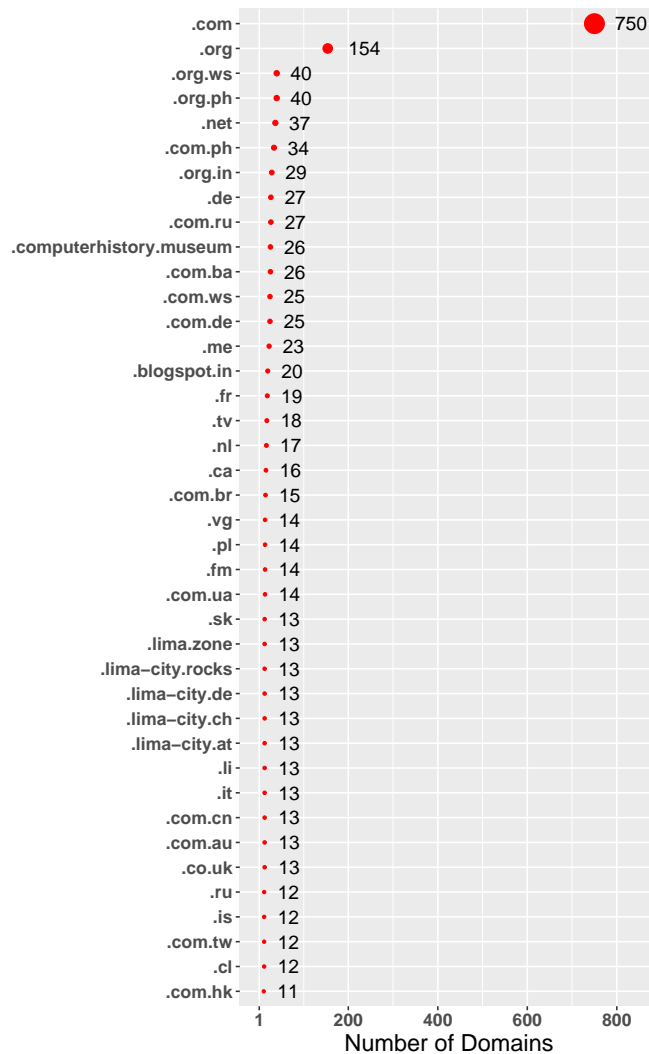


Figure 4.5: Breakdown of filtered suffixes

the number of valid filtered domains after this enumeration. It also shows us that the Chinese censorship regime filters many domains that do not have DNS records associated with them. A breakdown of the most filtered suffixes is shown in Figure 4.5. Unsurprisingly, the *.com* and *.org* along with numerous other country-level TLDs have the greatest number of filtered domains associated with them, likely due to the fact these rank amongst the most widely used TLDs for websites around the world [117]. More unexpected findings are the inclusion of suffixes relating to the Computer History museum of California and Lima City (a German domain registrar) within the top filtered suffixes.


Table 4.3: Filtered domain counts after suffix enumeration

Original filtered Domains	After suffix enumeration	Of which, DNS records exist
1355	97,167	5408

4.5.2 Locations of filtered hosts

The geographical locations of servers that host filtered domains can provide insight into the countries of origin being blocked. We infer the origin locations of servers hosting filtered domains by making a DNS query for each domain to the control server and using MaxMind GeoIP2 country database [92] to locate the resulting IP addresses by country of origin. We expect that websites originating in the USA, where over 50% of Internet servers are located [4], to be the most widely filtered. Diplomatic relations between the US and China continue to be strained over various international issues [72], which reinforces this presupposition that large amounts of US content may be censored. Figure 4.6 shows the breakdown of filtered hosts around the world. From this, it is clear that the US *is* the most blocked, where Belgium, Canada, Germany, Russia and the UK are the next most filtered.

Figure 4.6: Locations of filtered hosts



4.5.3 Seed domain connectivity

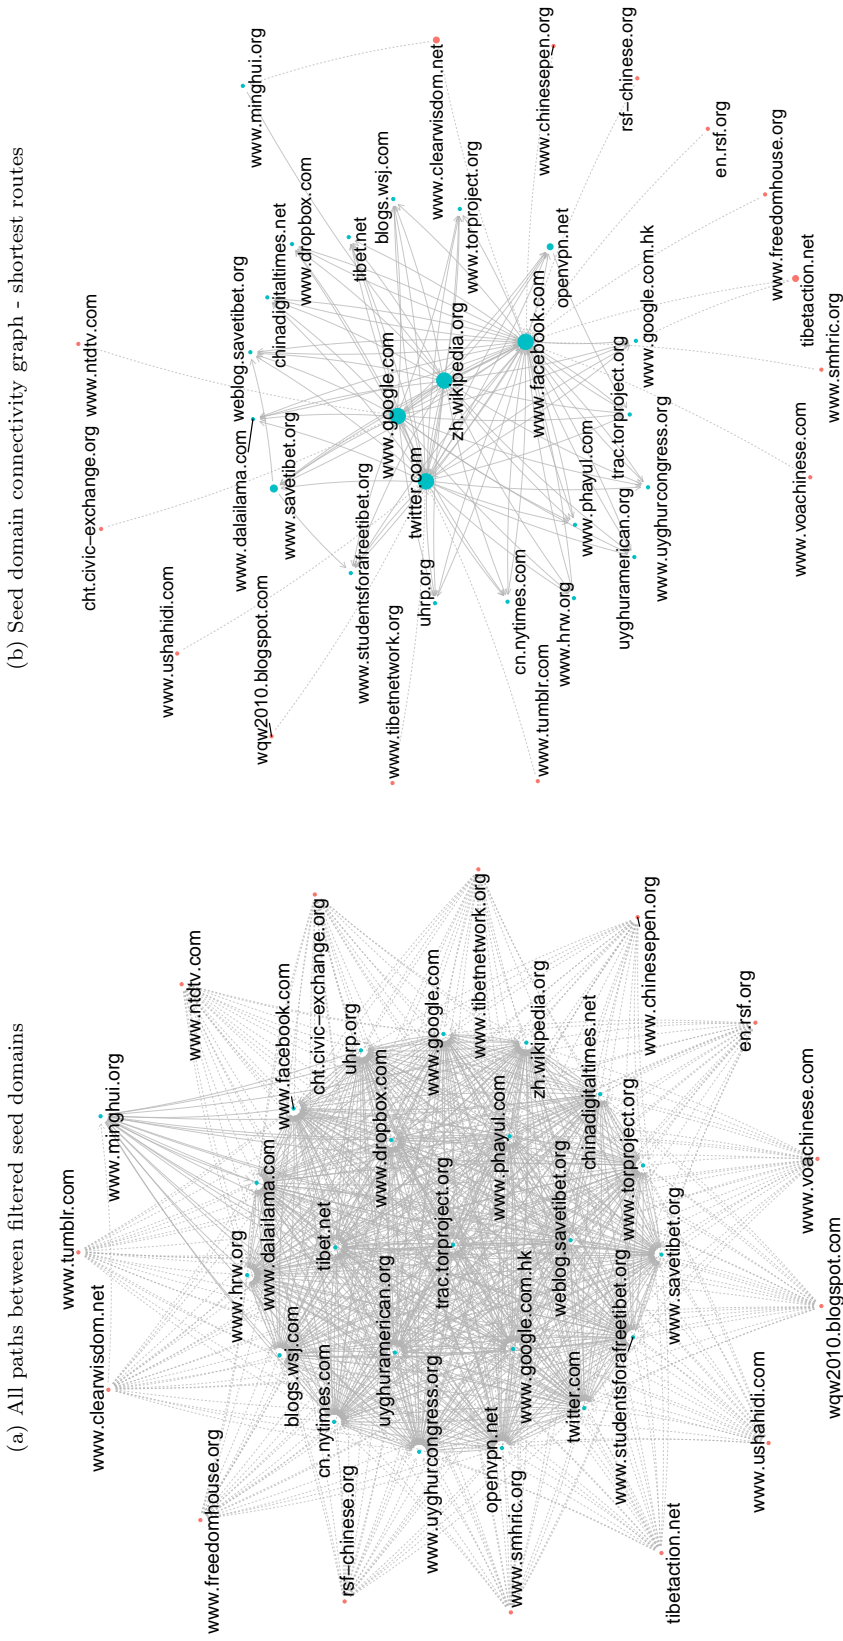
As part of the evaluation of the framework, we are interested in measuring how effective it is at discovering filtered material. This is not an easy test however, since we cannot possibly know how many blocked domains exist for any particular target country. An alternative is to initialise the system with a subset of the seed domains, and then find the recall metric on how many of the remaining seed domains are found. This isn't a perfect test given that the selection of filtered seed URLs likely have biases in how they were discovered originally¹, however we can still provide a notion of how well the approach can discover other filtered content that is known. To demonstrate the recall of seed URLs, we calculate the network of connections between them having run the experiment. This graph is shown in Figure 4.7 and depicts all the possible paths between the different seed domains that were traversed during the test. This is a highly connected network, within our dataset, 14 of the seed domains are not reachable from the remaining 25 and 5 were totally disconnected².

This demonstrates how the selection of seed URLs is indeed important, but also how this system provides a broad enough scope over filtered domains. If a 50% subset of the filtered seeds available were used, this approach would have found a *minimum* of 25 of all filtered seeds. This represents a base recall rate of 57%, in a longer test, this may likely be higher given the recursive nature of the system.

¹This is beyond our control since we use the filter lists maintained by the CitizenLab

²Did not point to or have inbound connections from any other filtered seed domain

Figure 4.7: Seed domain connections



4.5.4 Limitations & Issues

There are three main limitations we have identified with the implementation of this framework.

The first is that the act of filtering may result in search engines being unlikely to list blocked web pages highly in their rankings due to a lack of other sites linking to them. Similarly, for smaller websites, it may be the case that search engines do not rank these highly enough to be included in our approach. This could be mitigated by using one or more alternative search engines and combining the results. However, even given this, we can still achieve good results when using just the single engine.

A second limitation is use of “anti-SEO” techniques by websites. This could prevent a high-ranking listing and therefore, detection by our system. This is especially prudent since a site that contains sensitive content may aim to reduce its exposure to search engines in an explicit effort to reduce the chances of being discovered and blocked by a censorship regime.

Finally, we only consider English words when isolating descriptive tags. This means that a large chunk of potentially filtered (Chinese) sites are being missed by our tool. A future version of the system could include Chinese words and phrases as well as their translation in English.

4.6 Summary

This Chapter has presented a framework for the automated discovery of filtered websites without reliance on local, contextual knowledge or language specific information. The approach is fully automated, scalable and effective. The experiment has demonstrated that the technique can be used to discover previously unknown filtered content, which yields further meaningful results used recursively to continue the search. The system does not require complex infrastructure—beyond existing public services—nor significant computational or network requirements to operate, but instead uses existing public and cost-effective resources.

Further, this approach does not rely on local knowledge or participation from individuals in countries of interest, and thus avoids many significant ethical challenges in building URL filter lists. The technique is limited however, by the effectiveness of the keyword (*descriptive tag*) mining methods and filtering check method. While these present challenges, the real-world results demonstrate considerable improvement over existing systems and could be further improved in future iterations of the tool or technique.

The results produced are a significant contribution to the research field. The list of currently filtered domains within China is far more up to date and accurate than what is currently available and we expect to make this available in the near future. With the framework we have proposed,

maintaining URL filter lists can be done with minimal effort and resources. The experimental results demonstrate this as we present a filter list of filtered domains for China that is 30 times larger than the current most widely-used public list.

The ability to find URLs that can lead to other blocked material is substantial and has considerably increased the capability for monitoring and understanding internet censorship as it develops in an ethically sound manner.

4.6.1 Future Improvements

The current limitations of the implementation lend themselves to some future extension. The most obvious, and important, of these is to extend the targets of the tool and apply it to other states that implement network filtering. Clearly, this requires integration with other methods that can be used for automated checking of filtering activity beyond China, however, we have proposed a number of means to do this and are actively pursuing this potential.

More directly, further research into the incorporation and comparison of multiple search engines into the approach would yield a useful insight into the biases created, or not, of each individual engine. It would also be advantageous to quantify the extent to which using more search engines improves the efficacy and efficiency of the system. Furthermore, research into the possibilities of the use of multiple search engines aiming to reduce overall bias could allow for better operating parameters for the system.

The extraction of longer key phrases from filtered URLs as search terms could be investigated as a means to provide more specific search results. This may include the use of n-grams or other, more sophisticated language processing techniques. There are a number of existing services that can provide functionality for creating short descriptions from larger bodies of text, or indeed web pages. These models are usually based on machine learning and could be custom designed for this purpose. Even so, some existing services are able to leverage large amounts of data that has been scraped from the Internet.

Finally, incorporating a range of methods for checking the filtered status of resources, and the use of existing testing infrastructures such as OONI or ICLab, may increase the efficiency of the tool to build more comprehensive filtering lists.

“Market research is difficult because people don't think what they feel, don't say what they think and don't do what they say”

David Ogilvy

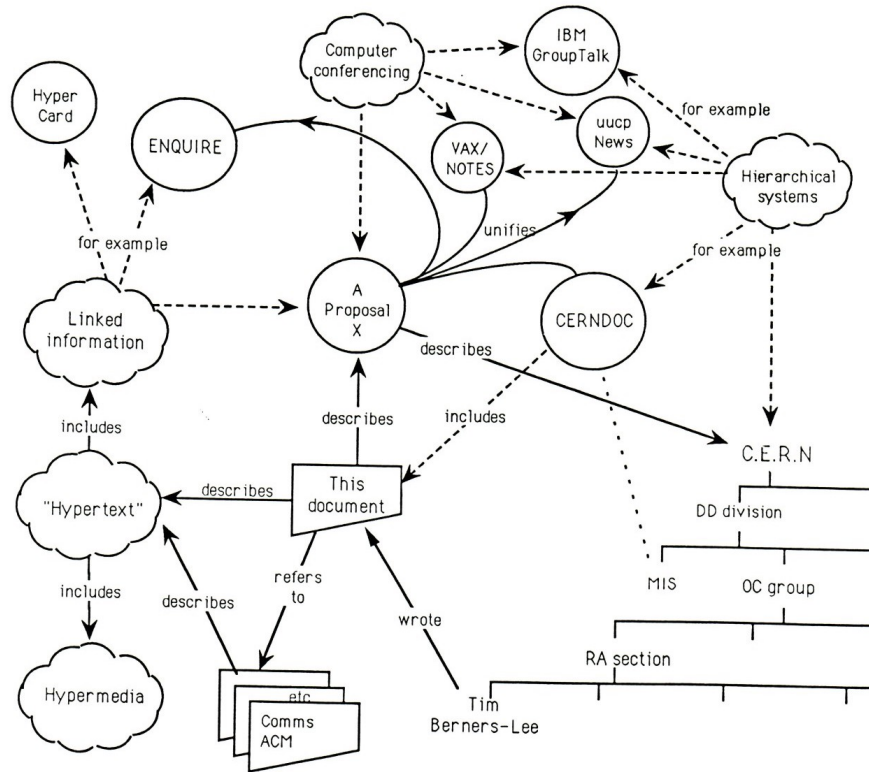
CHAPTER 5

Automated Filtered Resource Discovery Using Hyperlink Traversal within the Webgraph

The World Wide Web (the Web) is a logical network built to exploit the physical network of cables, routers, servers, satellites and cell towers of the Internet. This infrastructure was being developed and expanded throughout the 1980's and early 1990's at fast pace. By 1995 there were approximately 16 million Internet users, in 2000 there were 360 million [187]. Adoption of the technologies that made up the early Web were pivotal in this enormous growth. Key to this were the Uniform Resource Locator (URL), Hyper Text Transfer Protocol (HTTP), Hypertext Mark-up Language (HTML) and the web browser—some of which can be found in the original proposal for the Web sent to the managers of CERN by Tim Berners-Lee—Figure 5.1.

At its heart, the Web is a collection of documents embedded with referral links between them. This allows humans and machines to traverse between different resources and document groups in a semantically meaningful way. The semantics are derived by *how* the links are used and placed. For instance, if an uncommon word or technical term is used within some text, a *link* can be placed to a second document, or other piece of text in the same document, containing its meaning or description. This will provide a *“describes relationship”* between the term and the other document (piece of text). The real value of such a system arises when links are used to direct between large groups of documents, or websites. Furthermore, and patently key to the Web's expansion, is the fact links can direct users or machines to documents hosted by *other* servers that are *physically* located in different places, potentially on opposite ends of the globe. This is the reason why the Internet has grown and become distributed widely around the world—dispersal of services. Importantly, an Internet user on the Web can roam between hundreds of different

Figure 5.1: Original diagram for the Web in a proposal to CERN’s management—Information Management: A Proposal [18]



webpages (documents) and websites (groups of documents) hosted by globally separated sources (servers) in minutes. The implication of this phenomenon is that the same user can be exposed to vast ranges of different ideologies, thoughts, opinions & information as they traverse.

There have been few efforts to use such a reality for discovering censored web resources. The first use of linking filtered URLs via language connections in search terms were found in [39]. Discovering other blocked content via the *Webgraph*—traversing semantically relevant links between documents—is another new approach to monitoring Internet censorship [38]. The premise is that people and organisations will link to other materials that they refer to or deem important for their own content. Using traditional web crawling / scraping techniques will allow for us to discover these and find pieces of content that were previously unknown to be filtered.


5.1 The Connectedness of the Web

The fundamental basis of the Web are hyperlinks. Hyperlinks are usually expressed as URLs (see Section 2.3.2), though they are simply abstract references to data so can be of numerous formats. Links between webpages on the Internet allow for easy traversal between pieces of data for humans and machines (or crawlers). The original use of hyperlinks was to *refer* a user to definitions of uncommon terms or abbreviations written into textual data—see Figure 5.2.

From simple referrals to deep, complex graphs of interconnected documents, this is key operating principle of the World Wide Web. As the network has grown, it has been named as the Webgraph—set of directed hyperlinks (*edges*) between webpages (*nodes*). This graph has been studied and exploited by many researchers and businesses, particularly the *degree distribution* over the nodes and their connections—the proportion of links to and from nodes in the network. The precise distribution is unclear and not fully understood

[95], however, it is generally accepted that the Webgraph does not follow the *random graph model* [51]. It appears to more readily described as a *lognormal* degree distribution that follows a power law through a *scale-free network* [30][17]. This follows that new pages added to the web are more likely to create further links to other webpages with an already proportionally higher degree than the rest of the graph, which in turn increases their likelihood to be linked again in the future¹. This is also known as *cumulative advantage*. This occurrence has led to the development of numerous applications of the Webgraph for data analysis—such as PageRank [109] and Cyber-community detection [88]. These systems exploit the connections between webpages and sites to compute interesting and useful weightings between documents. The aforementioned PageRank is still a fundamental part of Google's system for ranking web search results. *Link-analysis* has been pivotal to the development of search engines where the importance and relevancy of webpages regarding search terms is a key feature. This allows for the identification of pieces of content "similar" to another piece and particularly the detection of communities of interest on the web. These communities can be expressed in numerous ways, the most relevant being forums, blogs or sites that share common topics of discussion. For the purposes of censorship detection, finding and

Figure 5.2: Original diagram for hyperlinks in a proposal to CERN's management—WorldWideWeb: Proposal for a HyperText Project [19]



¹Akin to the "rich get richer" generative model

exploring these communities is a key research task since we may identify new pieces of filtered content as well as uncovering the specific communities that censors are targeting.

5.1.1 Links between Online Communities

Online communities (or *internet communities*, *web communities*) are *virtual communities* that contain members who interact via the Internet. These can manifest in numerous ways, such as information / bulletin boards, forums, blogs or social networks. It is often the case that the members of online communities share common interests or ideas which fuel the content that is disseminated within them. Examples of such groups are: special interest forums (e.g. angling or coffee), chat rooms and virtual worlds (such as Second Life or other online multiplayer games).

“A virtual community is defined as an aggregation of individuals or business partners who interact around a shared interest, where the interaction is at least partially supported and/or mediated by technology and guided by some protocols or norms.”

Constance Elise Porter—A Typology of Virtual Communities: A Multi-Disciplinary
Foundation for Future Research [114]


Large online communities can also have *inner communities* or cliques that form within the wider network. This is especially true for global social networks such as Facebook or Twitter which both have many thousands of private groups and circles that share content internally. Wikipedia, a large online encyclopædia, also has many inner groups (although they are publicly accessible). These consist of different topics and sets of articles that are largely written, edited and managed by defined groups of authors. This has the affect of creating a *scale-free network* within Wikipedia itself, where popular articles are more likely to be linked in the future. Interestingly the “graph” structure of Wikipedia is very similar in structure to the Webgraph itself, as shown in Figure 5.3.

The ubiquitous nature of scale-free networks within human derived social graphs is widespread. We find that these phenomena exist across almost all activity where people communicate ideas and information—similar to the commonality of the Pareto principle (80/20 rule). Figure 5.4 shows how technical documents written for the Request for Comments (RFC), published by the Internet Engineering Task Force (IETF), also produces a network readily described by a log-normal degree distribution in regards to the referral links between individual pieces of content.

Given the recurrence of scale-free networks in online communities, it seems that Internet users share information in a relatively predictable manner. This is to say that if we follow links or referrals between pieces of content, we will likely find a set of *hub nodes* that act as disseminators for wider sets of content. Furthermore, we can conceptualise these *hub nodes* as weak authorities on the

Figure 5.4: Graph of RFCs (source: Björn Höhrmann)

The colour of the nodes indicate the subject cluster where the size and location depend on the number of requests and inbound-links for each document respectively.



This technique works on a simple premise—filtered webpages contain links to other filtered webpages. Similar to [39], we begin the discovery 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 Figures 5.5 & 5.6 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 go to 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.

Figure 5.5: High-level overview of filtered webpage traversal


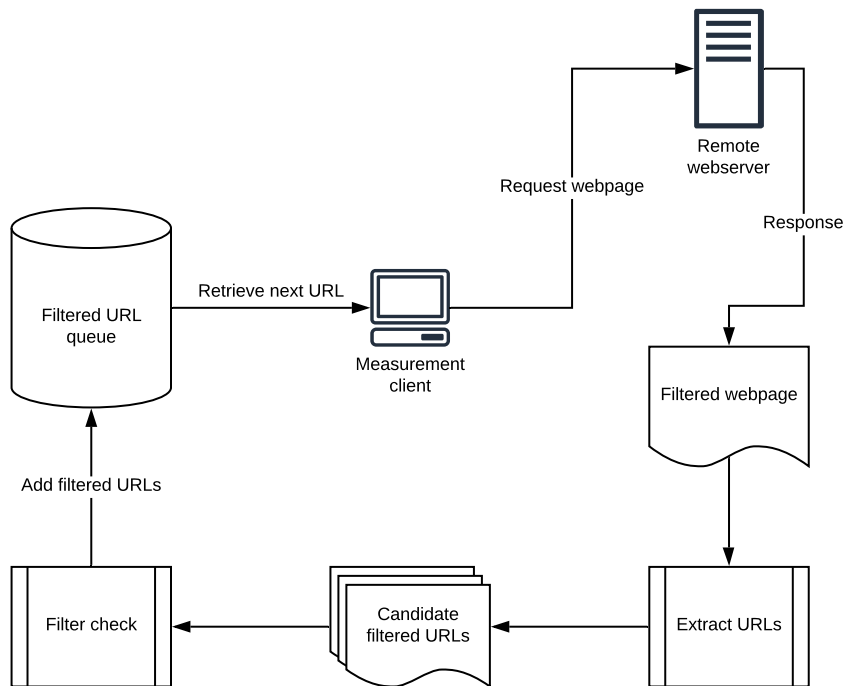


Figure 5.6: Filtered URL traversal process: High-level methodology



5.2.1 URL Extraction from Webpages

Extracting URLs from raw HTML webpages is a key part of this approach. This process must accomplish three tasks:

1. Identify the set of all *valid* URLs within a webpage (via URL search regular expression below)
2. Remove URLs that point to a static asset (javascript, CSS, images, binary content, etc.)
3. Remove URLs that point / refer to the parent URL of the webpage

It is important that this procedure is efficient and effective since it would be detrimental to the approach as a whole if the URL extraction process is slow or inaccurate. Self-looping in web crawling is desirable in certain cases, however for these purposes we need only visit a particular URL once per *session*—a session lasting until the crawl has reached a set limit of depth from the seed URLs. This regex identifies URLs by searching for top level domains across web pages, then extrapolating the URL path to find complete URLs. Domains in email addresses are not captured with this and will be excluded in the derived candidate URLs for each webpage scanned.

```
URL search regular expression
(?:i)\b((?:https?:({1,3}|[a-z0-9%])|[a-z0-9.\-]+[.](?:com|net|org|edu|gov|mil|
aero|asia|biz|cat|coop|info|int|jobs|mobi|museum|name|post|pro|tel|travel|xxx|ac|
ad|ae|af|ag|ai|al|am|an|ao|aq|ar|as|at|au|aw|ax|az|ba|bb|bd|be|bf|bg|bh|bi|bj|bm|
bn|bo|br|bs|bt|bv|bw|by|bz|ca|cc|cd|cf|cg|ch|ci|ck|cl|cm|cn|co|cr|cs|cu|cv|cx|cy|
cz|dd|de|dj|dk|dm|do|dz|ec|ee|eg|eh|er|es|et|eu|fi|fj|fk|fm|fo|fr|ga|gb|gd|ge|gf|
gg|gh|gi|gl|gm|gn|gp|gq|gr|gs|gt|gu|gw|gy|hk|hm|hn|hr|ht|hul|id|ie|il|im|in|io|iql|
ir|is|it|je|jm|jo|jp|ke|kg|kh|ki|km|kn|kp|kr|kw|ky|kz|la|lb|lc|li|lk|lr|ls|lt|lu|
lv|ly|ma|mc|md|me|mg|mh|mk|ml|mm|mn|mo|mp|mq|mr|ms|mt|mu|mv|mw|mx|my|mz|na|nc|ne|
nf|ng|ni|nl|no|np|nr|nu|nz|om|pa|pe|pf|pg|ph|pk|pl|pm|pn|pr|ps|pt|pw|py|qa|re|ro|
rs|ru|rw|sa|sb|sc|sd|se|sg|sh|si|sj|Ja|sk|sl|sm|sn|so|sr|ss|st|su|sv|sx|sy|sz|tc|
td|tf|tg|th|tj|tk|tl|tm|tn|to|tp|tr|tt|tv|tw|tz|ua|ug|uk|us|uy|uz|va|vc|ve|vg|vi|
vn|vu|wf|ws|ye|yt|yu|za|zm|zw)/)(?:[^\s()<>{}|\[\]]+|\\([^\s()]*?\([^\s()]+\)|
)*?)|\\([^\s()+?\\)]+)|(?:[^\s()]*?\([^\s()]+\)|)*?)|\\([^\s()+?\\)]+)|[^\s!()\\
{};:'.>?«»‘’‘’]|(?:(<!@)[a-z0-9]+(?:[.-][a-z0-9]+)*[.](?:com|net|org|edu|gov|
mil|aero|asia|biz|cat|coop|info|int|jobs|mobi|museum|name|post|pro|tel|travel|xxx|
ac|ad|ae|af|ag|ai|al|am|an|ao|aq|ar|as|at|au|aw|ax|az|ba|bb|bd|be|bf|bg|bh|bi|bj|
bm|bn|bo|br|bs|bt|bv|bw|by|bz|ca|cc|cd|cf|cg|ch|ci|ck|cl|cm|cn|co|cr|cs|cu|cv|cx|
cy|cz|dd|de|dj|dk|dm|do|dz|ec|ee|eg|eh|er|es|et|eu|fi|fj|fk|fm|fo|fr|ga|gb|gd|ge|
gf|gg|gh|gi|gl|gm|gn|gp|gq|gr|gs|gt|gu|gw|gy|hk|hm|hn|hr|ht|hul|id|ie|il|im|in|io|
iql|ir|is|it|je|jm|jo|jp|ke|kg|kh|ki|km|kn|kp|kr|kw|ky|kz|la|lb|lc|li|lk|lr|ls|lt|
lu|lv|ly|ma|mc|md|me|mg|mh|mk|ml|mm|mn|mo|mp|mq|mr|ms|mt|mu|mv|mw|mx|my|mz|na|nc|
ne|nf|ng|ni|nl|no|np|nr|nu|nz|om|pa|pe|pf|pg|ph|pk|pl|pm|pn|pr|ps|pt|pw|py|qa|re|
ro|rs|ru|rw|sa|sb|sc|sd|se|sg|sh|si|sj|Ja|sk|sl|sm|sn|so|sr|ss|st|su|sv|sx|sy|sz|
tc|td|tf|tg|th|tj|tk|tl|tm|tn|to|tp|tr|tt|tv|tw|tz|ua|ug|uk|us|uy|uz|va|vc|ve|vg|
vi|vn|vu|wf|ws|ye|yt|yu|za|zm|zw)\b/(?(!@))
```

5.3 Experimental Analysis

This approach lends itself to thorough censorship detection experiments since it is both scalable and an ethically sound approach (see Sections 2.8 & 5.3.1). We aim to derive useful research data from this method as with the previous in Chapter 4 as well as determining the effectiveness of such as censorship monitoring system. To achieve this, we again take censorship measurements regarding the blocking of URLs in regimes known to filter [parts of] the Internet from its citizens. Taking from the previous approach in Section 4.3, in this round, we apply the URL traversal technique to four different countries in order to build a more in-depth map of censorship of the Web across a wider base of Internet users.

5.3.1 Further Ethical Consideration

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 and motivation. This can be casual in many ways, not least because testing of internet filtering often requires sending network traffic to and from censorship regimes [32]. Certain studies within the field have required the use of aware 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 researches 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 [204].

We must also further 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.

5.3.2 Country Specific Testing

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 depth¹: 100
- Seed URLs obtained from the CitizenLab filter lists [29], each pre-processed using the filter check
- All testing performed from a single virtual machine node over a 14 day period, from a UK IP address

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 [3]. 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 already been seen—it will be counted in the statistics gathered, but not checked again.

The target countries tested were: China, Indonesia, Iran and Turkey. We choose these due to the known DNS filtering activity occurring in these countries [112] as to provide a good test bed for the measurements. 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 5.1. 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 as mentioned previously, we aim to take measurements on mass infrastructure within the target countries rather than smaller organisations or individuals.

The use of “fake” DNS servers—IP addresses that do not point at a live DNS service—is to improve accuracy of the filter check, as described in Section 3.4.1. When a DNS query is sent to a non-DNS server, or unused IP address, into the IP range of a particular country, say China, any response we receive is a clear indication of an interception of the DNS request. The query was responded to by an intermediary, if no interception occurs, we will not receive a response.

¹This is the maximum depth of recursion from the seed URLs

Table 5.1: DNS servers used for experiments

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

5.3.3 Results

Table 5.2 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 [7], and Indonesia, a similar ban in 2010 [71] & 2017 [115].

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—shown in Table 5.3. 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.

Table 5.2: Results from experimental analysis

	<u>URLs</u>		<u>Domains</u>		
	Extracted	Filtered	Filtered	Filtered <i>Top 1000 removed</i>	Filtered <i>No adult domains</i>
China	33,082,217	2,098,264	1576	1454	1454
Indonesia	12,580,357	835,395	47,143	47,065	1280
Iran	15,381,873	1,868,852	651	576	576
Turkey	19,250,931	913,213	39,725	39,614	513
<i>Totals:</i>	80,295,378	5,715,724	89,095	88,709	3823

Table 5.3: Comparison of results to CitizenLab filter lists
CitizenLab figures accurate as of 1st Sept 2017

	Filtered Domains <i>Alexa Top 1000 & adult domains removed</i>		
	CitizenLab	Hyperlink traversal	<i>Difference</i>
China	127	1454	<i>11.4X</i>
Indonesia	124	1280	<i>10.3X</i>
Iran	351	576	<i>1.6X</i>
Turkey	131	513	<i>3.9X</i>

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. In total, we extracted over 80 million URLs over the 14 day period and visited 5.7 million of these that were filtered. We may well have found more filtered domains had we visited all extracted links, however, this would incur a significant time penalty—16x from the numbers in our tests. Furthermore, each layer of depth we enter during a web crawl adds a significant number of new URLs to crawl—for instance, if each of the 80 million extracted gives us one additional new URL to visit, we have doubled our amount of work. This is why an efficient mechanism for reducing the candidate URL list is important. In this case, we use a filter check, however, it could be modified in a future version depending on the circumstances or requirements.

5.4 Evaluation of Approach

Our results demonstrate that this approach is effective at finding previously unknown filtered domains. A major advantage of this technique is that *only* filtered URLs are visited, meaning that we can achieve efficient web crawling. Furthermore, we are able to track the paths that lead to filtered content by analysing routes taken by the crawler. Also, by identifying the backlinks of filtered URLs and the outbound links to other filtered URLs, we can analyse the network of filtered sites.

We observe that the results found in Turkey and Indonesia contain large numbers of adult sites—something that these two nations are known to be currently censoring [7][115]. This 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 [201]. 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.

5.4.1 Top-Level-Domain Enumeration

To extend the list of discovered filtered domains, we perform an enumeration of all publicly available top-level-domains (TLDs) (similar to that in Section 4.5.1) that can be attributed to different domains; and therefore, different DNS records. The purpose of this process is to find and quantify the extent of which the available suffixes of a filtered domain are themselves filtered—for example, if *facebook.com* is filtered, are *facebook.co.uk*, *facebook.ie* or *facebook.org* also filtered.

For this task, we use the Public Suffix List maintained by the Mozilla Foundation [57]. This list of TLDs contains all known public suffixes, common examples such as *.com* and *.org*, and less well-know 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 sub-domains, 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 5.4.

Table 5.4: Filtered domain counts after TLD enumeration

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

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 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.

5.4.2 Geographical location of blocked hosts

We infer the locations of the servers hosting filtered domains in each test country using the same method as used in Section 4.5.2. In this instance, the process is executed across the discovered filtered domains found across all four of the test countries. The breakdown of the origin of hosts for filtered domains across each test country is shown in Figure 5.7. A normalised version of this data is shown in Figure 5.8 and a comparison between the target countries in Figure 5.9. We normalise by dividing the number of servers found in each location hosting blocked content for a particular country by the total number of servers found for that country. To compare between the target countries we have followed the same method but divided per location rather than by each of the target countries themselves.

Unsurprisingly, and consistent with the findings in Section 4.5.2, 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 [4].

Figure 5.7: Location breakdown of hosts serving filtered domains for each target country




Figure 5.8: Location breakdown of hosts serving filtered domains for each target country (normalised)





Figure 5.9: Location comparison of hosts serving filtered domains between each target country



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 subdomain 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 data

centres 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.

5.4.3 Limitations & Drawbacks

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 web masters 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 [116].

5.5 Summary

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 or special access to third-party systems and APIs to operate. 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 is of significant benefit to current and future studies concerning research within this field; and, for organisations that build circumvention tools. 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.

5.5.1 Future Improvements

The approach described here 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. However, given the inherent issues with ethically taking censorship measurements for research, this is a somewhat tricky issue so may well not be feasible given the requirements.

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 [39] (Chapter 4) would integrate well with this approach. A combined system of this type could improve both 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.

*"Fifteen hundred years ago everybody knew the Earth was the center of the universe.
Five hundred years ago, everybody knew the Earth was flat,
and fifteen minutes ago, you knew that humans were alone on this planet.
Imagine what you'll know tomorrow."*

Agent K, MiB

CHAPTER 6

Using Filtered URLs To Gain A Deeper Understanding Of Internet Censorship in the Webgraph

As previously discussed, regimes that operate censorship infrastructure are usually reluctant to disclose details on *what* is being censored and *how* the given material is discovered or identified. The *why* of censorship implementation often comes down to a strand of political or social ideology which supports the reduction of spread of certain information. For our purposes, we will try to steer clear of the political debate on *why* censorship is implemented by certain regimes since that is a complex and tangled web of arguments that has *no clear* right or wrong answer. It seems safe to say that the motivations for censorship, filtering or blocking of the Internet in the West are somewhat different to those in the Middle East and Asia in general; and, while I personally support a doctrine of freedom of speech and expression (as do most of my colleagues), I would not posit that my "*Western*" views should necessarily trump those of other individuals around the world. Taking a hard line on either side of the censorship debate appears to have non-obvious and unintended consequences for all those involved. Furthermore, and on that note, it also seems that many democratic, Western societies are increasing censorship activity across a variety of platforms on the Internet. This may be collateral from the numerous political movements that have taken place over the past decade (such as LGBTQ+ rights, Feminism and BAME¹), or a new taste of semi-conscious, controlling-type mentality that arises from increased anti-terror, social justice and security positions. It is widely reported that many social media accounts of high-profile individuals and organisations, particularly on the far sides of the political spectrum are routinely

¹BAME is a term long used in the UK to refer to black, Asian and minority ethnic people. Its origin derives from "political blackness", an idea that various ethnic groups united behind to fight against discrimination in the 1970s.

disconnected or removed from the platforms that host them (commonly Facebook, Twitter & YouTube). Now this can be seen as a positive move in the reduction of *hate speech* or potentially dangerous *radical political* views, but it is unquestionably censorship of those voices; whether this is positive or negative for a society as a whole remains to be seen.

With this in mind, we will analyse censorship activity within countries that: 1. already censor large internet platforms that originated outside of their geography; 2. use unknown sources or methods to identify sensitive content to block. This will involve taking an objective view of the types of content being filtered and the links between them (hard connections—hyperlinks & semantic—language). The aim of this is to build a better understanding of the efficacy of our new censorship discovery tools such that we can identify where they could be improved upon to acquire further unknown filtered content and improve the efficiency of collection for either increased scalability or more focused discovery tasks.

6.1 Filtered Webpage Category Breakdown


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 [197]. This returns a list of categories attributed to each domain and allows us to isolate from a high level different genres the websites that are being blocked in each target country¹. Figure 6.1 shows the overall category list and counts across all the target countries that were measured for censorship in Chapter 5. The observed categories relate to numerous subjects; and, each domain checked can belong to multiple categories.

Website categorisation is used widely across different industrial, academic and government sectors for a variety of reasons. One such use is within corporate web filters that are used across the entire network of large organisations—such as businesses or schools. The systems will be configured to block, or filter, any webpages / domains that are classified within a certain set of blacklisted categories. Rules may well be greedy or indifferent, that is to say that if a website is classified with three unique categories and one of them appears in the blacklist, the website could be blocked or allowed depending on the respective ruling. Large nation state censors are known to use website / webpage categorisation as part of their infrastructure [83], this type of filtering allows for widespread censorship of material that matches to a known blacklist of categories. Given this, we can observe how the distribution of filtered website categories across each test country. Further, we can identify if, at a particular point in time, a censor is filtering one type of content

¹Webshrinker was the service chosen to perform this task due to its widespread use in DNS filtering products for corporate cyber security.

Figure 6.1: Category counts over discovered filtered domains across all of the measurement (target) countries

Pornographic domains removed




more than others. This information could be used in support of other evidence—blogs, news / media reports, official reports, social media—to ascertain what the current filtering targets of a censorship regime are.

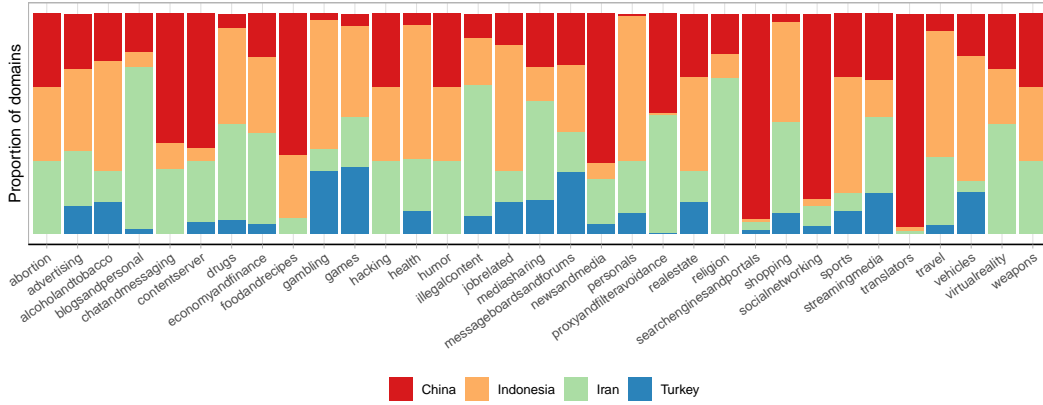
Figure 6.2a shows the category breakdown for each country. From this we can see that certain types of site are overwhelmingly being blocked over others. Of particular interest is the filtering of news and media, search engines and translators by China, personals and shopping by Indonesia

Figure 6.2: Filtered website category comparisons

(a) Category breakdown of filtered domains per target country



(b) Category comparison of filtered domains between target countries



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 [104] and a similar circumstance around the Iranian presidential election in 2013 [190].

Figure 6.2b 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 which different types of content that are being blocked across the test countries. 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.

6.1.1 Co-occurring Categories

To build on the observations of different filtered content types per test country, we calculate the co-occurrence matrix between categories of filtered domains using the Pearson product-moment correlation coefficient [113]. This gives further information as to the way different blocked websites “cross” different subject topics in regards to their content and target users. The Pearson coefficient gives the linear correlations between each category that is identified per filtered website.

The calculated category co-occurrences per target country, from Chapter 5, are shown in Figure 6.3. The contrasting differences of co-occurring categories of filtered domains between the test countries is indicative of different targets of censorship. We know that Islamic countries largely illegalise all pornographic and gambling web content due to the *role of Hisbah* as part of Islamic law [107]. This is shown by the category co-occurrences in Figures 6.3b (Indonesia), 6.3c (Iran) & 6.3d (Turkey), where the adult category co-occurs numerous times with other categories of filtered domains. The converse of this can be seen in Figure 6.3a (China), where adult and gambling categorised sites did not correlate largely with any other particular category. We can see clear differences in the co-occurring categories between these countries giving us a good view into which types of content the respective regimes are blocking. We also notice some categories don't co-occur at all: *abortion, contentserver, hacking, humour & weapons* never coincide with other categories.

It is likely that methods for censor regime derived categories of websites relies, at least in part, on keyword semantic analysis [85], as shown to be the case in China. The calculation of co-occurring categories is a relatively simple method to reverse-engineering this process. Since most censors do not publish which types of material is blocked, or indeed the methods used to do so, the censorship research community has used numerous techniques to attempt to describe the systems in place. The set of known blocked website categories for a country is a useful top-down view of this space due to the way that specifics of automated *to block or not to block* decisions are coalesced into the overall classified category for a webpage or website. In contrast, a look into specific keywords that are found on blocked pages provide a bottom-up approach into determining the types of content, this is discussed further in the next section.

Knowing which content categories co-occur with each other could be useful for a future version of the systems described in Chapters 4 and 5. In order to run these approaches efficiently, it's important to reduce the size of the candidate URL list. Looking at highly co-occurring website categories could prove a practical mechanism for doing so and would be worthy of further investigation.

Figure 6.3: Category correlation heatmaps per country

(a) China




Figure 6.3: Category correlation heatmaps per country (continued)

(b) Indonesia




Figure 6.3: Category correlation heatmaps per country (continued)

(c) Iran





Figure 6.3: Category correlation heatmaps per country (continued)

(d) Turkey



6.2 Descriptive Tag Distribution in URLs

We use the notion of *descriptive tags* (as discussed in Section 4.1.1) to define a piece of language that can be used to compare the topics and content between different webpages. Given this, we can model the language used across filtered websites to derive the patterns of content that are indicative of sensitive material. With this information, we can identify specific themes, topics or semantics that are homogeneous across different sets of censored web content. Furthermore, we can identify the *corresponding linguistic connections* that are apparent through the dataset of filtered webpages that have been collected.

6.2.1 Searching for Descriptive Tags in Filtered URLs

An interesting use of the collection of descriptive tags is that we can search for them *ex post facto* on each of the filtered URLs discovered. From this we can see which tags feature most prominently in filtered domains within China. This is by no means a search for necessarily blocked keywords, however, it could be used as a method to find patterns of language or topics that may have caused a page or domain to be blocked. Figure 6.4 shows the top 75 descriptive tags across all the filtered URLs in our database after we removed the Alexa Top 1000¹. These are the 75 most frequently occurring tags found across all filtered web pages per test country. It is unsurprising that the top 8 tags found across all filtered URLs are related to social media and mobile devices—we would expect these terms to potentially appear several times per web page since they represent the largest social networks on the Internet. Further down the list we find tags relating to freedom of speech movements, democracy, and filtering avoidance technology. Rather interestingly, we also find many references to Trump (US president in office) and Obama (previous president in office)—we believe this is due to the number of articles on filtered news sites and social media platforms that have mentions of them. At the time of experimentation, the volume of news relating to the US presidential activities was large, so this is to be expected.

¹We do this because it was found that many of the top 1000 have similar and generic language, we are more interested in the subtler and potentially more sensitive language.

Figure 6.4: Top 75 descriptive tags found in filtered webpages





Figure 6.4: Top 75 descriptive tags found in filtered webpages (*continued*)



6.2.2 Descriptive tag word embeddings

Word vectors, or embeddings, allow us to mathematically represent the semantic relationships between terms in natural language text. A word vector is a high dimensional model of an individual word—essentially a matrix of numbers that can have additional mathematical operations applied to it.

Word embeddings can be generated using numerous techniques, from bag of words to deep neural networks. However, generally the process for pre-processing the text is the same—in that this is simply the splitting of all documents in the corpus to individual sentences.

We can gain interesting insight through this process by identifying tags that are semantically related. This provides us the opportunity to find groups of tags, or patterns of language, that are represented across the filtered web pages. Furthermore, we can distinguish links between filtered content *by content* rather than direct links—such as hyperlinks. This is akin to how the mechanism used in Chapter 4 operated, except we can directly observe the content based links without the inherent bias involved with a search engine’s ranking algorithm¹. This approach could be extended to produce further search queries for the approach in Chapter 4 or to generate blocked keyword lists for other censorship measurement systems.

6.2.2.1 Calculating word vectors for descriptive tags

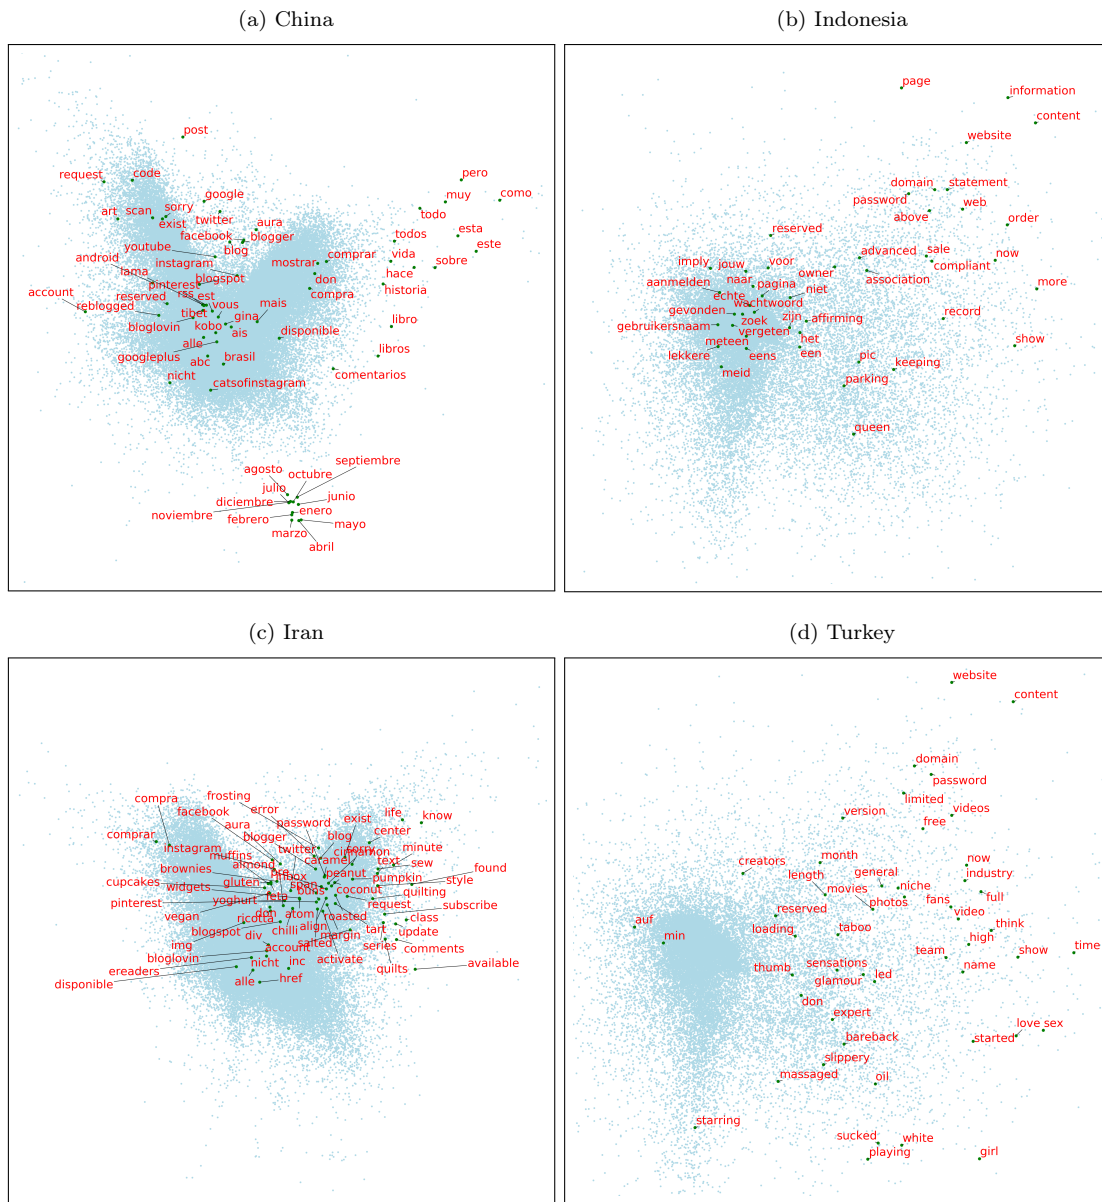
We calculate the word vectors for all descriptive tags found per test country. The algorithm used for this is Word2Vec [99], created by Google in 2013. For each filtered webpage discovered, we take the text negated from any HTML or structural code. These texts represent the documents in our final corpus. For each of the documents, we split the text into individual sentences. These sentences are then used as raw input to the Word2Vec algorithm. The implementation used for this work is from the Python Gensim NLP framework.

We plot the top 75 descriptive tags per test country, from those in Figure 6.4, amongst all word vectors calculated per test country—shown in Figure 6.5. To create these 2-dimensional plots, the high-dimensional word vectors were processed using principle component analysis (PCA), and the first two components were plotted. The X & Y values for each tag plotted can be conceptually viewed as a *distance* between the words, thereby presenting a simplified visual for the semantic relations between individual words as found in the context of these filtered websites.

These word embeddings show how the distribution of semantics of the language found in blocked websites differs between each test country. The density of the plotted tags for each

¹Search engine rankings are affected by over 100 different factors, many of which are unrelated to the content of a web page

Figure 6.5: Descriptive tag word embeddings per country



country contrasts substantially, showing how different linguistic patterns are observed for their respective censorship activities. Specifically, the sparse nature of the top 75 tags plotted for Indonesia and Turkey (Figures 6.5b & 6.5d) contrasts to those in China and Iran (Figures 6.5a & 6.5c). This shows that many of the top 75 tags *co-occur across numerous* Iranian filtered websites as the *semantic distance* between them is relatively short. We also observe that the plot for China, Figure 6.5a, displays 3 sub-clusters. These again, are groups of tags that are a *short semantic*

distance apart, and are frequently occurring across filtered content for China.

In all of the word embedding plots, we see a number of different languages other than English, showing that the censors we have monitored are language agnostic in the types of content they filter. This suggests that their systems are highly diverse in terms of coping with numerous languages, or that they have systems capable of determining sensitive content without requiring deep knowledge or data on a specific language. As the web is continually evolving, the case for the latter is potentially stronger.

6.2.3 Connection through descriptive tags

We model the co-occurrences of tags between filtered webpages, similar to the process for web page category. The purpose of this is to demonstrate the way that language, or content, is a key link between different censored websites. Further, this may indeed be a fundamental predictor in the likelihood to whether a particular web page or site becomes filtered in the future. While we don't focus on the keywords (read tags) themselves, we can demonstrate the extent to which the language used is pertinent to online censorship.

Figure 6.6 shows the language links between filtered domains in China for a single hop between any two websites¹. These are instances where the same descriptive tags are found on a parent filtered page to a child filtered page—where the parent and child are connected by a hyperlink. For clarity, only the top 30 domains—that is the those with the *highest number of child filtered domains*—are shown on the graph.

This graph displays numerous and dense clusters of tags that are effective links, or “*linguistic connections*”, between large sites and services. The clusters suggest that certain pieces of language are apparent on certain websites; and, furthermore, could potentially be grouped as predictors for potential filtering activity. This is a key learning from the data collected since we can identify “hubs” of censorship activity and the associated terms, a future implementation of a censorship detection system/tool could make use of this to: 1. increase efficiency, by prioritising the hubs as central points of discovery; and, 2. improve the efficacy of identifying terms that are indicative of blocked content, by collecting text from these hubs.

6.3 Content Citations

There has so far been a main focus on the *linguistic connections* between filtered websites within this Chapter. These are less obviously present than a form of harder, or syntactical connection, such as hyperlinks. In this respect links via language can be thought of as *soft-relationships* and

¹A hop in this instance is a single co-occurrence of a tag on two distinct domains

direct webpage citations, hyperlinks, can be thought of as *hard-relationships*. That is, a hard-relationship is visible as a published connection to another webpage, whereas a soft-relationship is more nuanced and subtle. It is somewhat more difficult to ascertain a soft-relationship, yet when achieved, these uncover a wide array of insight about censorship activity. Yet, we must not discount hard-relationship as published citations. These show a different kind of connection, ones that a site or service has made public, or indeed a user on such a service has. In fact, we find that we can discover more related censored domains by following hyperlinks—web crawling—than simply using linguistic connections, as shown in Chapter 5.

6.3.1 Routes of discovery for Chinese censorship

Identifying the routes by which filtered content was found can give a view on the *linked* nature of censored websites. We know that blocked content that is published references other blocked content—given the method used for discovery in Chapter 5.

Our results show that the networking effect between different filtered domains is vast for Chinese censorship. The original premise for this work was that filtered webpages will contain language that links them to other filtered sites. We can demonstrate this by visualising a graph of filtered domains and how they link to each other similar to the graph of seed domains in Chapter 4. Many of the domains that link to the most other filtered content are major sites that are widely used on a mass scale—as shown in Figure 4.3a. This is unsurprising since these sites often appear high in search engine rankings for popularity reasons. While we find that the top domains do point to smaller sites, many of the sites found in search results are themselves within the set of top domains. Figure 6.7 shows the connectivity between smaller, less widely known sites, the Alexa Top 1000 are removed and we plot the top 50 filtered domains that link to the most other filtered domains.

In this diagram, we can see clear collections of sites that contain distinct topics, as we expect. Notable examples is the cluster of domains belonging to VPN providers found from *www.torvpn.com* and the links between twitter related sites *twitpic.com* and *twiends.com*. Furthermore, this indicates how certain sites group together, especially those that are within the subject of human rights and democracy. Topics around the Uyghur human right projects and governmental sites for East Turkistan are tightly linked in this network which correlates with a significant amount of media scrutiny and investigation into tensions between the people of these areas and Beijing [66][21].

6.3.2 Backlinking of filtered webpages


For a more in-depth look into the networking effect between blocked websites, we calculate the number of filtered backlinks to and filtered forward links from each blocked webpage¹. 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.

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 6.8, 6.10, 6.12 and 6.14 show the backlinks of filtered domains for each target country. We remove the Alexa Top 1000 sites from each of these to provide a higher-definition look at the block sites with less traffic, shown in Figures 6.9, 6.11, 6.13 and 6.15.

Notable observations in Figures 6.10 and 6.14 are that the top sites that link to other filtered domains appear to be adult link collections which supports the findings in [201]. We can also see in Figures 6.9 and 6.13 that many of the linkers to filtered content are freedom of expression and independent news sites, both of which often contain political criticism.

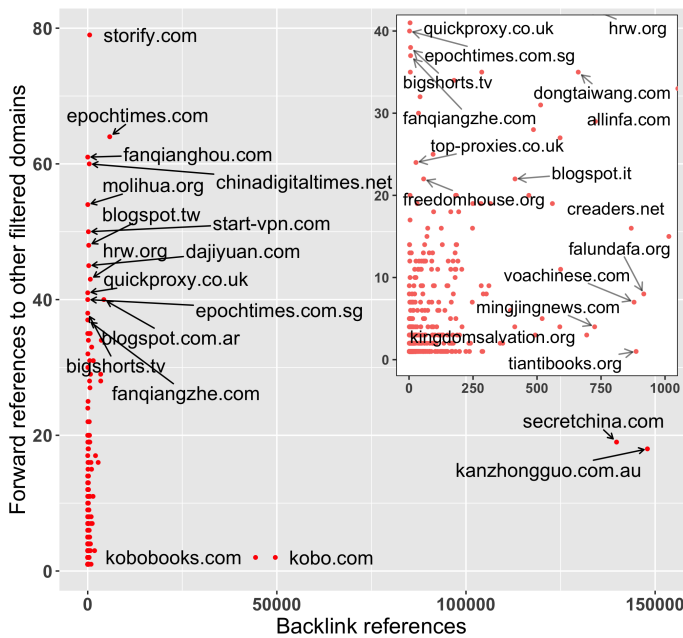
¹Note that the backlinks and forward links are also themselves filtered in the given target country

Figure 6.8: Backlinks of discovered filtered domains
China




We find that the largest outward linkers to other filtered content in China are the major western social networks. In particular content on *blogspot.com* and *blogger.com* are both referenced by other blocked websites (over 2 million backlink references) as well as linking to over 500 other filtered domains.

Figure 6.9: Backlinks of discovered filtered domains
(Top 1000 sites removed)
China




With the top visited sites in the world removed, we can see better the other, non-social network, types of sites that are blocked. As also shown in Figure 6.2a, we see numerous news portals and web proxy / VPN sites in the plot here. *secretchina.com* and *kanzhongguo.com.au*, which are both news sites that report on China's internal politics, are both heavily linked to by other filtered sites.

Figure 6.10: Backlinks of discovered filtered domains
Indonesia




All of the top forward linkers and back referenced sites for Indonesia are adult sites. This conforms with earlier findings where we see high co-occurrences between the “adult” category with others (Figure 6.3b). This diagram also very clearly shows the major link-hubs (those with the highest number of forward references) that our crawler system used to find numerous other filtered sites.

Figure 6.11: Backlinks of discovered filtered domains
(adult sites removed)
Indonesia




Here we see a large number of webhosting and dating sites in the figure. We also find large social networking and blogging sites—which we know are major targets for many censors. *asacp.org* is charity that works to protect children online, after further investigation, it appears to be associated with a large number of adult sites (which often link to it) which cooperate to reduce the exploitation of children in the adult film industry. This is why we see it backlinked from many other filtered content.

Figure 6.12: Backlinks of discovered filtered domains
Iran




The majority of sites here are social networking and news sites. We also see that *citizenlab.org*, a research institute which deals with political and human rights issues; and, also from whom we derived our seed lists, appears relatively highly as a forward linker to other filtered content in Iran.

Figure 6.13: Backlinks of discovered filtered domains
(Top 1000 sites removed)
Iran




With the top 1000 site removed we can see a wide variety of sites in this plot. Particularly there are numerous VPN/web proxy list sites. *tribunezamaneh.org*, a news portal and broadcaster that advocates for freedom of speech for the Iranian media, is highly back referenced here.

Figure 6.14: Backlinks of discovered filtered domains
Turkey



Similar to the results found for Indonesia, the majority of the filtered sites shown in this plot are adult sites. We also find that the top referrers to other filtered sites are also link repositories or directories pointing to adult sites.

Figure 6.15: Backlinks of discovered filtered domains
 (adult sites removed)
Turkey



With the adult sites removed, we can see that *instagram.com* is the domain with the highest number of backlinks from other filtered sites. There are also numerous gambling sites here in the top outward linkers, with around 5 outward links each. The sites in question: *888.com*, *888casino.com*, *888poker.com*, *888poker.es* and *888games.com* are all associated with each other by the same ownership.

6.3.3 Filtered Domain Discovery Power

The basic assumption of this work is that filtered web pages contain linguistic patterns that could lead to further filtered URLs. The framework in Chapter 4 achieves this by deriving descriptive tags from a base URL which are used as search queries to link to other URLs—where the tag is the connection.

To convey how strongly certain domains provide descriptive tags that lead to the discovery of other filtered material, we introduce a measure: *discovery power*. This represents the number of filtered URLs from different domains that are found by searching with tags isolated from a given base domain. We calculate the discovery power by looking at the tags that, when used as search queries, result in the discovery of at least one other poisoned domain. (We do not count instances within the search results where the discovered domain is the same as the base domain—that is to say we don't count self-discovery.)

Comparing the discovery power with the number of crawled URLs for each domain provides insight into how much potentially sensitive language is present on those domains. This is not only interesting in the understanding of censorship activity, but also to optimise our tool to consider only higher power domains when deriving tags.

In Figures 6.8, 6.10, 6.12 and 6.14, we can conceptualise *discovery power* as the ratio between *forward references* to *backlink references*—that is the number of filtered domains pointed outbound divided by the number of filtered domains pointed inbound from any given website or domain. The more filtered domains a particular website links to against the number that point to it, give us a sense of how *powerful* the given websites links are.

Notable examples of domain discovery power in Figure 6.9, for China, are: *storify.com*, *epochtimes.com*, *fanqianghou.com* and *chinadigitaltimes.net*, all of which produced “*high-power*” tags that resulted in a higher number of filtered URLs than the number URLs crawled for each of those domains.

6.4 Summary

This chapter has provided a first look at a new dataset of censored websites. Through the separation of network topology and language/content, we have studied numerous aspects of censorship activity among four different nations. While it is difficult to identify the specific mechanisms in use by these regimes—in regard to *how* or *why* they find content to block—we have been able to analyse the *results* of such systems. In particular, the categorisation breakdown of filtered websites provides a key insight into the types of content being blocked per

measurement country. A direct result of this analysis is the identification of different sets of content categories that are filtered in each country. A second finding are the category co-occurrence distributions for each country. These show how deemed sensitive content, per censor, crosses different category boundaries. This strand of study could of course also be extended through the use of more distinct categories.

From a network topological point of view, we have explored the way that censored webpages link to one-another and how this affects discovery paths in finding them. This is likely key to how censorship mechanisms identify new content to block. The results of this work give a high-level perspective of how the linking nature among blocked websites could be exploited in this way. In addition, we can identify specific websites domains that are *highly* connected via back and forward links to other filtered websites. These could be viewed, conceptually, as “link repositories” of sensitive webpages/websites. We coin the term—*domain discovery power*—to describe a measure that is a result of direct network hyperlinks *and* linguistic connections derived from content. This provides a method for determining discrete comparisons between different filtered domains.

We have modelled linguistic patterns across filtered websites using word embeddings and descriptive tag connections. This shows how the direct network connections (or hyperlinks) can be partnered with language connections in order to bolster the overall conceptual picture of how censorship applies across filtered websites. We find that the individual measurement countries each have different, and sometimes contrasting, models of language in regard to the websites they censor. Furthermore, these models can could change over time, offering a chance to monitor the language of censorship as it changes and develops.

6.4.1 Future Research

Having reviewed the data outputs gathered from the two conducted experiments, we can identify several new avenues for further research. In particular, there are many ways the discovery approaches could be altered to possibly yield better or more efficient results. With the experience of the performed experimentation and analysis, we are placed in a stronger position to design more robust future approaches and techniques for censorship detection.

Identifying more relevant keywords from filtered content

Several pieces of past research have focused on building higher quality blocked keyword lists, particularly: [33][53][86]; and, [69] which builds upon the work detailed in this thesis from our earlier publication [39]. This being said, we can now identify several new approaches for conducting this type of study. Firstly, we have the concept of our descriptive tags which can be

collected from filtered webpages then weighted by the manner in which they provide patterns of linguistic connectivity—as shown in Sections 6.2 & 6.3. Using the webgraph as means for identifying and ranking keywords could offer a fruitful mechanism for identifying the trend of filtered topics and content in future work. Secondly, we have built new language models for analysing patterns of censored content using word embeddings. These could offer another method for determining the importance of particular keywords. Furthermore, if integrated into an automated system, such as the types presented in the earlier chapters, it could offer more efficient and dynamic mechanisms for discovering newly filtered content through continual analysis and processing of textual data from web pages.

Using filtered website categories as a discovery feature

The experiments we conducted have focused on identifying relationships between filtered content using hyperlinks and language patterns. Looking at the data collected concerning the website categories, we could use this to better identify suitable candidate URLs to add to the discovery/check phase in the frameworks. This could be based on findings around which website categories are more prevalent in a particular target country's filtered domain lists—if we find a certain category of website leads us to more filtered content, we can prioritise these to be more efficient during the discovery. Again, this feature would be useful if used dynamically as to reduce the number of URLs we need to visit/check per filtered domain found, thereby potentially allowing the system to find more results for a target country over the same time frame. Moreover, we could also begin to express whether filtered sites in one category follow to more filtered sites in the same or other categories.

Links between filtered sites as predictive indicators

We have made extensive use of inter-website relationships in order to *discover* censorship, however, there could be additional predictive power in these links. A future study built upon the work presented here, but aimed at using the links between websites to predict which are or will be censored in the present or future respectively. Furthermore, in a new version of our discovery system, one could use the linking nature of filtered content to dynamically predict routes to further filtered content in a more efficient manner. The extent to which censored websites could be tiered in their importance is open for further research beyond what has been presented in this thesis.

"Time spent doing nothing is rarely wasted"

Paraphrase of John Lennon's: "Time you enjoy wasting,
was not wasted."

CHAPTER 7

Conclusion

The academic communities that study of censorship across the globe are broad and active. There are numerous strands of work, research groups and organisations that provoke scientific curiosity into the complex, widespread, human, technical, social and political world that is censorship. New works in this field are consistent, applicable and useful inclusions to the public knowledge about the phenomenon. Technical pieces account for a large subset of these works; and, while it is sometimes easy to think of many censorship activities as excessively technical problems, it remains the case that censorship is a socio-political event. The technical means of exercising control have always been at the forefront of thought / information restriction, through either the destruction of printing presses, or indeed the filtering of internet communications. At the end of each of these activities however, therein sit the people, government, organisation or authoritarian regime that makes decisions on what to block and what means to implement.

This thesis has dealt with several strands of work that are prevalent in the wider academic community that studies censorship. To summarise, the discussed contributions are methods for: *accurately determining if a website is censored, discovering previously undocumented censored websites and using the derived data for insight into censorship activities*. These works have had direct impact on the community, its literature and future studies. Two peer-reviewed papers were published as a direct result of these [38][39], and our paper: *"FilteredWeb: A framework for the automated search-based discovery of blocked URLs."* [39] has been extended by other researchers in the field [69]. In addition, the papers have received numerous citations in more recent studies into mapping censorship.

The means to accurately detect if a website is filtered from a remote vantage point is a key contribution of this work. Numerous approaches have been used in the past for censorship research, however, these usually suffer from potential ethical issues. The ability to determine if a web service or domain is being censored *without the use of human volunteers* is important. Further, the approach discussed here has been designed to take censorship measurements from large scale infrastructure that is *purposely designed to filter web traffic*. It therefore does not take the targeted systems outside of their known operating remit.

We use this new method for filter status checks in the first known use of search engines to uncover previously unpublished blocked domains in China. This experiment was predicated on the hypothesis that censored webpages will contain language that can be used as search queries to discover more censored pages. The results of the test were conclusive and we are able to identify an order of magnitude more filtered websites than the given seed list of filtered webpages—which was the most frequently cited and used filter lists available at this time. The approach is scalable since it leverages existing services and ethically sound given it uses the filter check process stated above. This test also shows that language is a form of relationship that exists between filtered websites and essentially forms a graph of topical links between them.

To further study the effects of connectedness between censored content on the Web, we introduced a second method for filtered domain discovery. Web crawling, the process by which one traverses hyperlinks between websites, is shown to also produce a stout set of unpublished censored domains from the same seed list as used above. The experiment to validate this method ran over four individual test countries: China, Indonesia, Iran & Turkey—which resulted in a better performance than the previous approach, in terms of censored domains found. The direct linking nature between filtered sites is a fundamental discovery of this work, where, we are able to demonstrate the extent to which certain online communities cite internally and externally to each other.

The results of these experiments yielded a dataset of newly known filtered URLs from a set of domains censored in the test countries. With these, we also obtained the direct links between them—hyperlinks or citations; and, the linguistic connections—descriptive tags. Through analysis of these features, we are able to determine certain characteristics of the content that was found to be filtered and the nature of the relationships between different censored web resources. We find that certain categories of site are particular targets for filtering in certain countries, and that these indeed differ between each of the test countries. Further, we are able to look at the textual keywords found on each filtered webpage to ascertain the language patterns used. These are both known to be key factors of how censorship machines operate at a technical level, so this work

provides a broad perspective of these features. Further, the analysis is concluded with an outlook over how we are able to identify sites that are *highly connected* to other filtered domains, through backlinks and forward citing links from and to them respectively. While we cannot affirm the mechanism, or part thereof, by which censors themselves identify content they wish to filter, the idea that these systems will use “*seed websites*” from which sensitive content can be found, is a likely candidate.

7.1 Resolution of Research Questions

This thesis originally proposed four key research questions. Here, details are given to describe the extent to which these have been addresses with the results produced.

To what extent can we determine if Internet resources are censored/blocked whilst keeping ethical concerns minimal?

In Chapter 3 we discuss a new instrument for measuring censorship of domains using new techniques to increase accuracy and scalability. We are fervent in the non-use of remote volunteers or agents to reduce potential ethical concerns, however, do acknowledge that we create limitations when resources are not checked from or within the country who’s censorship machines we are measuring. Beyond the approaches shown here, there are numerous others that have been offered by the research community over the past decade, some more ethically sound than others. To this extent it, and with the results in this thesis, it is possible to measure censorship of Internet resources without protruding onto morally or ethically insecure grounds.

How can we use automated censorship discovery methods to monitor Internet censorship?

In addition to the wide range of censorship discovery techniques presented in Section 2.7.8 (some of which are automated), Chapters 4 and 5 offer two new approaches for monitoring Internet censorship without direct manual involvement throughout their operation. We use lists of known filtered websites to “seed” the systems and then go on to identify large numbers of other filtered content, not in the original lists. Whilst no long-running computer system is completely automated—maintenance and updates are required at regular intervals—we show that the process of discovering and monitoring Internet censorship can be conducted without direct human involvement. This offers a more scalable and cost-effective system, in terms of time and money spent, we have shown that large numbers of websites can be monitored for censorship from a relatively small set of resources, in this case, a single virtual machine running on a laptop.

Can patterns and structure in one set of filtered content be used to find further filtered content?

In Chapter 4 we present a technique that leverages existing search engines using queries generated from web page text to traverse between different filtered domains. In Chapter 5 we used an alternative approach where we follow hyperlinks on web pages to identify other filtered domains. Both methods displayed results which show that we can indeed use a set of known filtered sites to find other filtered sites, where we find up to 11 times more in the experiments performed.

Are there homogeneous links between different filtered content that are apparent through the analysis of linguistic patterns?

In Chapter 4 we explicitly use and test the notion that textual patterns can be used to link different filtered content. The medium which was used to create these links were commercial search engines, which we use to traverse from one filtered site to another by identifying useful search queries on web pages through TF-IDF. This approach yielded a list of filtered domains over 30 times larger than the seed list used to start the search process. In Chapter 6 we discuss how language models can be built from text scraped filtered web pages to uncover what topics are apparent and how small text snippets (*descriptive tags*) can be used to form relationships between nodes in the webgraph.

7.2 Future Work

The techniques and studies in this thesis lend themselves to future work, and indeed already have as mentioned above. Since the methods are agnostic frameworks, there are numerous places where extensions and further study could be performed. Firstly, there is the continuation of the experiments to cover more countries. As the Internet becomes more prevalent across the world, especially in third-world countries, there is a rising likelihood that internet filtering will be enacted. The approaches in this work provide direct means for monitoring this. Secondly, an improvement to the filter check mechanism would be beneficial. While the trade-off between ethical censorship measurements and granularity has been discussed and made here, it would be more ideal if *individual webpages* could be determined as “filtered” or not. The rise of larger scale networks in third-world countries may provide an opportunity for this, however, the use of these will need to be studied to ensure individuals within censorship regimes *are not* endangered due to their use. Finally, the use of one or both of the techniques in an operational setting could provide a valuable dataset that can be used to study the temporal effects of internet censorship over time. An on-going and consistent database of filtered URLs per country, the entry points and drop-off times would be very useful for this task. As a means to measure the status of online transparency

and censorship activity, these data points may prove invaluable. An existing service such as OONI [171], PlanetLab [155] or Ripe Atlas [163] may well have the capability required to track these events going forward.

7.2.1 Summary of future research threads

Below is a summation of different strands of future research that have been identified as part of the work conducted in this thesis:

Domain filtering measurements

Improve result coordination with CDNs

There has been a small amount of previous research into measuring network interference and CDN biases, notably Allison et al. [93] and Scott et al. [182]. However, these are now over two and four years old respectively and the technical landscapes of CDNs are continually changing. Furthermore, and as noted in Section 3.4.4, a censorship measurement instrument that can consistently account for changes in locally specified results (i.e. a DNS resolver responding with an IP address of a server based on distance from the client) could produce more accurate results going forward. This is becoming more pertinent since CDNs now account for the majority of Internet traffic. A future study of differing CDN responses and on-going measurements of these would be a significant use to the censorship research community.

Better identification of website locality

Valid results from a DNS server could send users to different versions of the same site based on their location, for example, a French language site if the user is located in France. Jones et al. [81] created techniques to automatically fingerprint block pages in 2014, however, since more and more sites are being made specific for the locality of users, their work could be extended to include fingerprinting of more general web pages. This would improve the filter check presented in this paper through the ability to reject a differing IP address result from a test DNS server from a control if the end-website was indeed correct for its location.

Automatic censorship discovery tools

Incorporation with other censorship discovery tools/infrastructure

Section 2.7.8 lists 58 approaches for measuring censorship published over the past 17 years. The efficacy of the censorship detection tools presented in this thesis could be improved by integrating some of the systems or data available from other areas of this research community. Particularly, OONI [55] and ICLab [106] which already have significant networks of agents and measurement

infrastructure in place. Both OONI and IClab produce large quantities of data which, if incorporated, could be used for filtered website candidate identification or testing—both of which may improve the efficiency and breadth of the discovery of filtered content for particular target countries. Furthermore, the use of multiple search engines for the discovery system presented earlier would be worthy of further investigation.

Improved keyword generation for search terms

In Sections 6.2 & 6.3 we explore the way that language can connect filtered content. This methods used in this process could also be used to generate candidate keywords for block lists or web search queries. In building more complex language models, one would be able to extract potentially viable topics, trending or otherwise, which could be used in place of the TF-IDF method used in Chapter 4. A future version of this system would benefit from research into different keyword generation approaches.

Dynamic ranking/prediction of candidate filtered websites/URLs

Currently the approaches presented here test all candidate domains for evidence of censorship, a more efficient system would prioritise specific domains for testing based on those that may yield more fruitful links to other filtered content or indeed have a higher likelihood of being filtered. Whilst it would be preferable to test every domain across every country, there are over 350 million registered, a system to test all of these would likely be prohibitively expensive to operate over a wider time scale. As such, if we are aiming to test more content, we must reduce the candidate list. This thesis has presented a number of methods which could be used to perform this, notably inter-website relationships, content categories and language models. All three of these could provide the basis of future studies into building more scalable censorship measurement systems and the research community could benefit from further interrogation into these proposals.

7.3 Final thoughts

The main goal of the work presented in this thesis was to acquire new knowledge surrounding the possibilities for improved automated censorship discovery. To that end, the academic community now has access to two new and novel approaches for achieving this, both offering substantially improved results over the current state-of-the-art. Whilst there are several routes for improvements, extensions and further testing/experimentation, this remains a significant contribution to the field. Moreover, the described techniques place considerable emphasis on the ethical concerns surrounding censorship research, offering methods that can reduce the impact of obtaining measurements.

The landscape of censorship has changed dramatically over the past two decades, as have the regimes and technologies that enact it. The research field is constantly producing new sets of methodologies and data for studying and interrogating the actions of censors and their systems. There remains a broad set of future research themes which are open for further investigation.

To conclude—censorship is a societal event. It's not the means to which censorship is administered that really matters, it's the intent and motivation. Though, we can use these means as a second order process in order to study the activity. Researchers, academics and organisations that have been involved in this subject have known and accomplished censorship investigations for decades and indeed centuries. This is a thesis written by a Western author, in a Western institution. Most people in this world, my world, have disparate thoughts towards censorship, though most condemn the exploits that censorship can, *and is*, used to realise. This is of course not *globally ubiquitous*. It would be difficult to argue that censorship cannot have its place as a moral endeavour in any sense or context. Yet, it seems that for the most part, censorship is used for arbitrary control, rather than empirical social benefit.

List of References

- [1] Giuseppe Aceto, Alessio Botta, Antonio Pescapé, M Faheem Awan, Tahir Ahmad, and Saad Qaisar. Analyzing internet censorship in pakistan. In *Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI), 2016 IEEE 2nd International Forum on*, pages 1–6. IEEE, 2016.
- [2] Giuseppe Aceto, Alessio Botta, Antonio Pescapé, Nick Feamster, M Faheem Awan, Tahir Ahmad, and Saad Qaisar. Monitoring internet censorship with ubica. In *International Workshop on Traffic Monitoring and Analysis*, pages 143–157. Springer, 2015.
- [3] Giuseppe Aceto, Giorgio Ventre, Niccolo Rinaldi, and Antonio Pescape. Monitoring internet censorship: the case of ubica, 2014.
- [4] U.S. Central Intelligence Agency. *Internet hosts, CIA World Factbook*, Accessed Oct 2017. <https://www.cia.gov/library/publications/the-world-factbook/rankorder/2184rank.html>.
- [5] Bernhard Ager, Wolfgang Mühlbauer, Georgios Smaragdakis, and Steve Uhlig. Comparing dns resolvers in the wild. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 15–21, 2010.
- [6] Vani Agrawal and Priyanshi Sharma. Internet censorship in india. *Available at SSRN 3309268*, 2019.
- [7] Mustafa Akgul and Melih Kirlidog. Internet censorship in turkey. *Internet Policy Review*, 4(2):1–22, 2015.
- [8] Sayim Aktay. Teacher perspective on internet censorship in turkey. *Universal Journal of Educational Research*, 6(2):296–306, 2018.
- [9] Sophia Ananiadou, Douglas B Kell, and Jun-ichi Tsujii. Text mining and its potential applications in systems biology. *Trends in biotechnology*, 24(12):571–579, 2006.
- [10] Collin Anderson. Dimming the internet: Detecting throttling as a mechanism of censorship in iran. june 2013. URL <http://arxiv.org/abs/1306.4361>.
- [11] Anonymous. Towards a comprehensive picture of the Great Firewall’s DNS censorship. In *Free and Open Communications on the Internet*. USENIX, 2014.
- [12] arc90. *Readability*, 2009 (accessed March, 2019). <https://web.archive.org/web/20101107162633/http://lab.arc90.com/2009/03/02/readability/>.

- [13] Simurgh Aryan, Homa Aryan, and J Alex Halderman. Internet censorship in iran: A first look. In *FOCI*, 2013.
- [14] UN General Assembly. Universal declaration of human rights. *UN General Assembly*, 1948.
- [15] Michael Bailey and Craig Labovitz. Censorship and co-option of the internet infrastructure. *Ann Arbor*, 1001:48104, 2011.
- [16] David Bamman, Brendan O’Connor, and Noah Smith. Censorship and deletion practices in chinese social media. *First Monday*, 17(3), 2012.
- [17] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [18] Timothy J Berners-Lee. Information management: A proposal. Technical report, 1989.
- [19] Timothy J Berners-Lee and Robert Cailliau. Worldwideweb: Proposal for a hypertext project. 1990.
- [20] Daniel J. Bernstein. *DNSSCurve: Usable security for DNS*, Accessed 8th June, 2017. <https://dnscurve.org/>.
- [21] Preeti Bhattacharji. Uighurs and china’s xinjiang region. *Council on Foreign Relations*, 29, 2012.
- [22] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.
- [23] Kevin Bock, Yair Fax, Kyle Reese, Jasraj Singh, and Dave Levin. Detecting and evading censorship-in-depth: A case study of Iran’s protocol filter. In *Free and Open Communications on the Internet*. USENIX, 2020.
- [24] Backlinko Brian Dean. *Google’s 200 Ranking Factors: The Complete List (2019)*, 2019 (accessed April, 2019). <https://backlinko.com/google-ranking-factors>.
- [25] Sam Burnett and Nick Feamster. Making sense of internet censorship: a new frontier for internet measurement. *ACM SIGCOMM Computer Communication Review*, 43(3):84–89, 2013.
- [26] Sam Burnett and Nick Feamster. Encore: Lightweight measurement of web censorship with cross-origin requests. In *ACM SIGCOMM Computer Communication Review*, volume 45, pages 653–667. ACM, 2015.
- [27] Abdelberi Chaabane, Terence Chen, Mathieu Cunche, Emiliano De Cristofaro, Arik Friedman, and Mohamed Ali Kaafar. Censorship in the wild: Analyzing Internet filtering in Syria. In *Internet Measurement Conference*. ACM, 2014.
- [28] China.org.cn. *The internet in china.*, 2018 (accessed Oct, 2018). http://china.org.cn/government/whitepaper/node_7093508.htm.
- [29] citizenlab.org. *citizenlab/test-lists*, 2014 (accessed Feb, 2017). <https://github.com/citizenlab/test-lists>.
- [30] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

- [31] Richard Clayton, Steven J Murdoch, and Robert NM Watson. Ignoring the great firewall of china. In *Privacy Enhancing Technologies*, pages 20–35. Springer, 2006.
- [32] Jedidiah R. Crandall, Masashi Crete-Nishihata, and Jeffrey Knockel. Forgive us our SYNs: Technical and ethical considerations for measuring Internet filtering. In *Ethics in Networked Systems Research*. ACM, 2015.
- [33] Jedidiah R Crandall, Daniel Zinn, Michael Byrd, Earl T Barr, and Rich East. Conceptdoppler: a weather tracker for internet censorship. In *ACM Conference on Computer and Communications Security*, pages 352–365, 2007.
- [34] David Croteau, William Hoynes, William D Hoynes, et al. *The business of media: Corporate media and the public interest*. Pine forge press, 2006.
- [35] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapè. Analysis of country-wide Internet outages caused by censorship. In *Internet Measurement Conference*, pages 1–18. ACM, 2011.
- [36] Jakub Dalek, Bennett Haselton, Helmi Noman, Adam Senft, Masashi Crete-Nishihata, Phillipa Gill, and Ronald J. Deibert. A method for identifying and confirming the use of URL filtering products for censorship. In *Internet Measurement Conference*. ACM, 2013.
- [37] George Danezis. An anomaly-based censorship detection system for tor. 2011.
- [38] Alexander Darer and Joss Farnan, Oliver & Wright. Automated discovery of internet censorship by web crawling. In *Proceedings of the 10th ACM Conference on Web Science*, pages 195–204. ACM, 2018.
- [39] Alexander Darer, Oliver Farnan, and Joss Wright. Filteredweb: A framework for the automated search-based discovery of blocked urls. In *2017 Network Traffic Measurement and Analysis Conference (TMA)*, pages 1–9. IEEE, 2017.
- [40] Ronald Deibert. The geopolitics of internet control: Censorship, sovereignty, and cyberspace. *The Routledge handbook of internet politics*, pages 323–336, 2009.
- [41] Ronald J Deibert, John G Palfrey, Rafal Rohozinski, and Jonathan Zittrain. *Access denied: The practice and policy of global internet filtering (information revolution and global politics)*. 2008.
- [42] Frank Denis and Yecheng Fu. *DNSCrypt*, Accessed 8th June, 2017. <https://dnscrypt.org/>.
- [43] Andrea Di Florio, Nino Vincenzo Verde, Antonio Villani, Domenico Vitali, and Luigi Vincenzo Mancini. Bypassing censorship: a proven tool against the recent internet censorship in turkey. In *Software Reliability Engineering Workshops (ISSREW), 2014 IEEE International Symposium on*, pages 389–394. IEEE, 2014.
- [44] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976.
- [45] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, Naval Research Lab Washington DC, 2004.
- [46] David Dittrich, Erin Kenneally, et al. The menlo report: Ethical principles guiding information and communication technology research. *US Department of Homeland Security*, 2011.

- [47] Haixin Duan, Nicholas Weaver, Zongxu Zhao, Meng Hu, Jinjin Liang, Jian Jiang, Kang Li, and Vern Paxson. Hold-On: Protecting against on-path DNS poisoning. In *Securing and Trusting Internet Names*. National Physical Laboratory, 2012.
- [48] William H Dutton. *Freedom of connection, freedom of expression: the changing legal and regulatory ecology shaping the Internet*. UNESCO, 2011.
- [49] Tariq Elahi, George Danezis, and Ian Goldberg. Privex: Private collection of traffic statistics for anonymous communication networks. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1068–1079, 2014.
- [50] Roya Ensafi, Jeffrey Knockel, Geoffrey Alexander, and Jedidiah R Crandall. Detecting intentional packet drops on the internet via tcp/ip side channels. In *International Conference on Passive and Active Network Measurement*, pages 109–118. Springer, 2014.
- [51] P. Erdős and A Rényi. On the evolution of random graphs. In *PUBLICATION OF THE MATHEMATICAL INSTITUTE OF THE HUNGARIAN ACADEMY OF SCIENCES*, pages 17–61, 1960.
- [52] Shadi Esnaashari, Ian Welch, and Brenda Chawner. Wcmt: web censorship monitoring tool. In *Telecommunication Networks and Applications Conference (ATNAC), 2013 Australasian*, pages 183–188. IEEE, 2013.
- [53] Antonio M. Espinoza and Jedidiah R. Crandall. Automated named entity extraction for tracking censorship of current events. In *Free and Open Communications on the Internet*. USENIX, 2011.
- [54] Oliver Farnan, Alexander Darer, and Joss Wright. 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*, pages 95–98. ACM, 2016.
- [55] Arturo Filastò and Jacob Appelbaum. OONI: Open observatory of network interference. In *Free and Open Communications on the Internet*. USENIX, 2012.
- [56] first20hours. *google-10000-english*, Accessed 2017. <https://github.com/first20hours/google-10000-english>.
- [57] Mozilla Foundation. *Public Suffix List*, 2017 (accessed Sept, 2017). <https://publicsuffix.org/>.
- [58] King-wa Fu, Chung-hong Chan, and Michael Chau. Assessing censorship on microblogs in china: Discriminatory keyword analysis and the real-name registration policy. *IEEE Internet Computing*, 17(3):42–50, 2013.
- [59] Genevieve Gebhart, Anonymous Author, and Tadayoshi Kohno. Internet censorship in Thailand: User practices and potential threats. In *European Symposium on Security & Privacy*. IEEE, 2017.
- [60] Philippa Gill, Masashi Crete-Nishihata, Jakub Dalek, Sharon Goldberg, Adam Senft, and G Wiseman. Characterizing censorship of web content worldwide: Another look at the opennet initiative data. *Published online through Stony Brook University*, 2013.
- [61] Philippa Gill, Masashi Crete-Nishihata, Jakub Dalek, Sharon Goldberg, Adam Senft, and Greg Wiseman. Characterizing censorship of web content worldwide, 2008.
- [62] Yoav Goldberg. A primer on neural network models for natural language processing. *Journal of Artificial Intelligence Research*, 57:345–420, 2016.

- [63] Devashish Gosain, Anshika Agarwal, Sahil Shekhawat, H. B. Acharya, and Sambuddho Chakravarty. Mending wall: On the implementation of censorship in India. In *SecureComm*. Springer, 2017.
- [64] GreatFire.org. *Online Censorship in China | GreatFire Analyser*, Accessed 2017. <https://en.greatfire.org/analyzer>.
- [65] RADICATI GROUP et al. Inc., email statistic report, 2015–2019, 2015.
- [66] Ziad Haider. Sino-pakistan relations and xinjiang’s uighurs: Politics, trade, and islam along the karakoram highway. *Asian Survey*, 45(4):522–545, 2005.
- [67] J Hall and M Aaron. A survey of worldwide censorship techniques. In *NETWORK WORKING GROUP*. IETF, 2015.
- [68] Nguyen Phong Hoang, Sadie Doreen, and Michalis Polychronakis. Measuring I2P censorship at a global scale. In *Free and Open Communications on the Internet*. USENIX, 2019.
- [69] Austin Hounsel, Prateek Mittal, and Nick Feamster. Automatically generating a large, culture-specific blocklist for china. In *8th {USENIX} Workshop on Free and Open Communications on the Internet ({FOCI} 18)*, 2018.
- [70] Ling Huang, XuanLong Nguyen, Minos N Garofalakis, Joseph M Hellerstein, Michael I Jordan, Anthony D Joseph, Nina Taft, et al. Communication-efficient online detection of network-wide anomalies. In *INFOCOM*, volume 7, pages 134–142, 2007.
- [71] Open Net Initiative. *ONI Country Profile - Indonesia*, Accessed Jan 2018. <http://access.opennet.net/wp-content/uploads/2011/12/accesscontested-indonesia.pdf>.
- [72] The National Interest. *A War Between the U.S. and China Would Be World War III*, Accessed Oct 2017. <http://nationalinterest.org/blog/the-buzz/war-between-the-us-china-would-be-world-war-iii-might-be-19287>.
- [73] Gunnar Eyal Wolf Iszaevich. Distributed detection of Tor directory authorities censorship in Mexico. In *International Conference on Networks*. IARIA, 2019.
- [74] Bernard J Jansen and Amanda Spink. Analysis of document viewing patterns of web search engine users. In *Web mining: Applications and techniques*, pages 339–354. IGI Global, 2005.
- [75] Timothy Jay. *Why we curse: A neuro-psycho-social theory of speech*. John Benjamins Publishing, 1999.
- [76] Jill Jermyn and Nicholas Weaver. Autosonda: Discovering rules and triggers of censorship devices. In *Free and Open Communications on the Internet*. USENIX, 2017.
- [77] Min Jiang. Search concentration, bias, and parochialism: A comparative study of google, baidu, and jike’s search results from china. *Journal of Communication*, 64(6):1088–1110, 2014.
- [78] Michael Johnston. Good governance: Rule of law, transparency, and accountability. *New York: United Nations Public Administration Network*, 2006.
- [79] Ian Jolliffe. Principal component analysis. In *International encyclopedia of statistical science*, pages 1094–1096. Springer, 2011.
- [80] Ben Jones, Roya Ensafi, Nick Feamster, Vern Paxson, and Nick Weaver. Ethical concerns for censorship measurement. In *Ethics in Networked Systems Research*. ACM, 2015.

- [81] Ben Jones, Tzu-Wen Lee, Nick Feamster, and Phillipa Gill. Automated detection and fingerprinting of censorship block pages. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 299–304. ACM, 2014.
- [82] Arturo Filastò, Khairil Yusof, Tan Sze Ming, Kay Yen Wong, and Maria Xynou. *The State of Internet Censorship in Indonesia*, 2017 (accessed May, 2017). <https://ooni.torproject.org/post/indonesia-internet-censorship/>.
- [83] Gary King, Jennifer Pan, and Margaret E Roberts. How censorship in china allows government criticism but silences collective expression. *American Political Science Review*, 107(02):326–343, 2013.
- [84] Gary King, Jennifer Pan, and Margaret E Roberts. Reverse-engineering censorship in china: Randomized experimentation and participant observation. *Science*, 345(6199):1251722, 2014.
- [85] Jeffrey Knockel, Masashi Crete-Nishihata, Jason Q Ng, Adam Senft, and Jedidiah R Crandall. Every rose has its thorn: Censorship and surveillance on social video platforms in china. In *5th {USENIX} Workshop on Free and Open Communications on the Internet ({FOCI} 15)*, 2015.
- [86] Jeffrey Knockel, Lotus Ruan, and Masashi Crete-Nishihata. Measuring decentralization of Chinese keyword censorship via mobile games. In *Free and Open Communications on the Internet*. USENIX, 2017.
- [87] Jeffrey Knockel, Lotus Ruan, and Masashi Crete-Nishihata. An analysis of automatic image filtering on WeChat Moments. In *Free and Open Communications on the Internet*. USENIX, 2018.
- [88] Ravi Kumar, Prabhakar Raghavan, Sridhar Rajagopalan, and Andrew Tomkins. Trawling the web for emerging cyber-communities. *Computer networks*, 31(11-16):1481–1493, 1999.
- [89] Anukool Lakhina, Mark Crovella, and Christophe Diot. Diagnosing network-wide traffic anomalies. In *ACM SIGCOMM Computer Communication Review*, volume 34, pages 219–230. ACM, 2004.
- [90] John Locke. *An essay concerning human understanding*. 1841.
- [91] Graham Lowe, Patrick Winters, and Michael L Marcus. The great dns wall of china. *MS, New York University*, 21, 2007.
- [92] MaxMind. *GeoIP2 Databases*, Accessed May 2019. <https://www.maxmind.com/en/geoip2-databases>.
- [93] Allison McDonald, Matthew Bernhard, Luke Valenta, Benjamin VanderSloot, Will Scott, Nick Sullivan, J. Alex Halderman, and Roya Ensafi. 403 forbidden: A global view of cdn geoblocking. In *ACM Internet Measurement Conference*, 2018.
- [94] Marshall McLuhan, W Terrence Gordon, Elena Lamberti, and Dominique Scheffel-Dunand. *The Gutenberg galaxy: The making of typographic man*. University of Toronto Press, 2011.
- [95] Robert Meusel. The graph structure in the web analyzed on different aggregation levels. *Journal of Web Science*, 1:33–47, 08 2015.
- [96] Jean-Baptiste Michel, Yuan Kui Shen, Aviva P Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al. Quantitative analysis of culture using millions of digitized books. *science*, page 1199644, 2010.

- [97] Microsoft. *Microsoft Azure Bing Search API*, Accessed March, 2019. <https://azure.microsoft.com/en-gb/services/cognitive-services/>.
- [98] Rada Mihalcea and Paul Tarau. Texttrank: Bringing order into texts. Association for Computational Linguistics, 2004.
- [99] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [100] Paul V Mockapetris. Domain names-concepts and facilities. 1987.
- [101] Donn Morrison. Toward automatic censorship detection in microblogs. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 572–583. Springer, 2014.
- [102] James Moscola, John Lockwood, Ronald Prescott Loui, and Michael Pachos. Implementation of a content-scanning module for an internet firewall. In *Field-Programmable Custom Computing Machines, 2003. FCCM 2003. 11th Annual IEEE Symposium on*, pages 31–38. IEEE, 2003.
- [103] Zubair Nabi. The anatomy of web censorship in pakistan. *arXiv preprint arXiv:1307.1144*, 2013.
- [104] Bloomberg News. *China Tells Carriers to Block Access to Personal VPNs by February*, Accessed May 2019. <https://www.bloomberg.com/news/articles/2017-07-10/china-is-said-to-order-carriers-to-bar-personal-vpns-by-february>.
- [105] Kei Yin Ng, Anna Feldman, and Chris Leberknight. Detecting censorable content on Sina Weibo: A pilot study. In *Hellenic Conference on Artificial Intelligence*. ACM, 2018.
- [106] Arian Akhavan Niaki, Shinyoung Cho, Zachary Weinberg, Nguyen Phong Hoang, Abbas Razaghpanah, Nicolas Christin, and Phillipa Gill. ICLab: A global, longitudinal internet censorship measurement platform. In *Symposium on Security & Privacy*. IEEE, 2020.
- [107] Helmi Noman. In the name of god: Faith-based internet censorship in majority muslim countries. In *Routledge Handbook of Media Law*, pages 261–276. Routledge, 2013.
- [108] Astrid Nordin and Lisa Richaud. Subverting official language and discourse in china? type river crab for harmony. *China information*, 28(1):47–67, 2014.
- [109] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- [110] Jong Chun Park and Jedidiah R. Crandall. Empirical study of a national-scale distributed intrusion detection system: Backbone-level filtering of HTML responses in China. In *Distributed Computing Systems*, pages 315–326. IEEE, 2010.
- [111] Paul Pearce, Roya Ensafi, Frank Li, Nick Feamster, and Vern Paxson. Augur: Internet-wide detection of connectivity disruptions. In *IEEE Security and Privacy*, 2017.
- [112] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global measurement of {DNS} manipulation. In *26th {USENIX} Security Symposium ({USENIX} Security 17)*, pages 307–323, 2017.
- [113] K Pearson. Notes on regression and inheritance in the case of two parents proceedings of the royal society of london, 58, 240-242, 1895.

- [114] Constance Elise Porter. A typology of virtual communities: A multi-disciplinary foundation for future research. *Journal of computer-mediated communication*, 10(1):JCMC1011, 2004.
- [115] The Jakarta Post. *Indonesia blocks 800,000 websites*, Accessed May 2019. <http://www.thejakartapost.com/news/2017/01/07/indonesia-blocks-800000-websites.html>.
- [116] Maria Praetzellis. *Identify and avoid crawler traps*, Accessed May 2019. <https://support.archive-it.org/hc/en-us/articles/208332943-Identify-and-avoid-crawler-traps->.
- [117] Q-Success. *Usage of top level domains for websites*, 2017 (accessed Oct, 2017). <https://w3techs.com/technologies/overview/>.
- [118] Lin Quan, John Heidemann, and Yuri Pradkin. Detecting internet outages with precise active probing (extended). *USC/Information Sciences Institute, Tech. Rep*, 2012.
- [119] Anand Rajaraman and Jeffrey D Ullman. *Mining of massive datasets*, volume 1. Cambridge University Press Cambridge, 2012.
- [120] Ram Sundara Raman, Leonid Evdokimov, Eric Wustrow, J. Alex Halderman, and Roya Ensafi. Investigating large scale HTTPS interception in Kazakhstan. In *Internet Measurement Conference*. ACM, 2020.
- [121] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored Planet: An Internet-wide, longitudinal censorship observatory. In *Computer and Communications Security*. ACM, 2020.
- [122] Ram Sundara Raman, Adrian Stoll, Jakub Dalek, Reethika Ramesh, Will Scott, and Roya Ensafi. Measuring the deployment of network censorship filters at global scale. In *Network and Distributed System Security*. The Internet Society, 2020.
- [123] Access Now. *Venezuela blocks access to the Tor network*, Accessed Oct, 2018. <https://www.accessnow.org/venezuela-blocks-tor/>.
- [124] Amnesty International. *Russia: Move to block Telegram the latest blow in government assault on freedom of expression online*, Accessed Oct, 2018. <https://www.amnesty.org/en/latest/news/2018/04/russia-move-to-block-telegram-the-latest-blow-in-government-assault-on-freedom-of-expression-online/>.
- [125] Annebritt Dullforce. *FT 500 2015*, Accessed Oct, 2018. <https://www.ft.com/ft500>.
- [126] BBC News. *Italy Wikipedia shuts down in protest at EU copyright law*, Accessed Oct, 2018. <https://www.bbc.co.uk/news/world-europe-44696302>.
- [127] Beacon for Freedom. *The Long History of Censorship*, Accessed Oct, 2018. http://www.beaconforfreedom.org/liste.html?tid=415&art_id=475.
- [128] Brian Dean, Backlinko. *WE ANALYZED 5 MILLION GOOGLE SEARCH RESULTS*, 2019 (accessed November, 2020). <https://backlinko.com/google-ctr-stats>.
- [129] Business Review Romania. *Certain Facebook users in Romania see content related to street protests reviewed under community standards*, Accessed Oct, 2018. <http://business-review.eu/news/certain-facebook-users-in-romania-see-content-related-to-street-protests-reviewed-under-community-standards-153255>.

- [130] Clifford Coonan. *Google set to pull out of China over censorship*, Accessed Oct, 2018. <https://www.independent.co.uk/news/world/asia/google-set-to-pull-out-of-china-over-censorship-1925052.html>.
- [131] Dan O'Leary, IMN. *GOOGLE ORGANIC CLICK THROUGH STUDY: COMPARISON OF GOOGLE'S CTR BY POSITION, INDUSTRY, AND QUERY TYPE*, 2017 (accessed November, 2020). <https://www.internetmarketingninjas.com/blog/google/announcing-2017-click-rate-study/>.
- [132] Dan York. *Turkish hijacking of dns providers shows clear need for deploying bgp and dns security*, Accessed 15th May, 2017. <http://www.internetsociety.org/deploy360/blog/2014/04/turkish-hijacking-of-dns-providers-shows-clear-need-for-deploying-bgp-and-dns-security/>.
- [133] Dean McDonald, Akamai. *WHY YOU SHOULD CARE ABOUT DNS LATENCY*, 2017 (accessed November, 2020). <https://blogs.akamai.com/2017/06/why-you-should-care-about-dns-latency.html>.
- [134] dorinlazar.ro. *The strange case of Facebook censoring people protesting a corrupt government*, Accessed Oct, 2018. <https://dorinlazar.ro/the-strange-case-of-facebook-censoring-people-protesting-a-corrupt-government/>.
- [135] Earl Zmijewski. *Turkish internet censorship takes a new turn*, Accessed 15th May, 2017. <http://dyn.com/blog/turkish-internet-censorship/>.
- [136] Elizabeth C Economy. *The great firewall of China: Xi Jinping's internet shutdown*, Accessed Oct, 2018. <https://www.theguardian.com/news/2018/jun/29/the-great-firewall-of-china-xi-jinpings-internet-shutdown>.
- [137] ExpressVPN. *Apple removes VPN Apps from China App Store*, Accessed Oct, 2018. <https://www.expressvpn.com/blog/china-ios-app-store-removes-vpns/>.
- [138] Fortune, Time Inc. *Judge Decides To Unblock WhatsApp In Brazil*, Accessed Oct, 2018. <http://fortune.com/2016/05/03/judge-unblock-whatsapp-brazil/>.
- [139] Freedom House. *Populists and Autocrats: The Dual Threat to Global Democracy*, 2017 (accessed May, 2018). <https://freedomhouse.org/report/freedom-world/freedom-world-2017>.
- [140] Friends Of Fort Point Channel. *Countries Where WhatsApp is Banned!*, Accessed Oct, 2018. <http://friendsoffortpointchannel.org/countries-where-whatsapp-is-banned/>.
- [141] Golden Frog. *Iran Increases Internet Censorship, Blocks Social Media and Messaging App*, Accessed Oct, 2018. <https://www.goldenfrog.com/blog/iran-censorship-social-media-messaging>.
- [142] ICLab. *Internet censorship lab*, Accessed 4th August, 2016. <http://internetcensorshiplab.com/>.
- [143] International Monetary Fund. *Report for Selected Country Groups and Subjects (PPP valuation of country GDP)*, 2018 (accessed July, 2018). <https://www.imf.org/external/pubs/ft/weo/2018/01/weodata/weorept.aspx>.
- [144] Lotus Ruan, Jeffrey Knockel, Jason Q. Ng, and Masashi Crete-Nishihata. *ONE APP TWO SYSTEMS*, Accessed Oct, 2018. <https://citizenlab.ca/2016/11/wechat-china-censorship-one-app-two-systems/>.

- [145] Mada Masr. *Signal unstable: Alternatives to the encrypted messaging application*, Accessed Oct, 2018. <https://madasr.com/en/2016/12/19/feature/politics/signal-unstable-alternatives-to-the-encrypted-messaging-application/>.
- [146] Majestic-12. *The Majestic Million*, (accessed July, 2018). <https://majestic.com/reports/majestic-million>.
- [147] Mike Isaac. *Facebook Said to Create Censorship Tool to Get Back Into China*, Accessed Oct, 2018. <https://www.nytimes.com/2016/11/22/technology/facebook-censorship-tool-china.html>.
- [148] Motherboard, Vice. *A Brief History of YouTube Censorship*, Accessed Oct, 2018. https://motherboard.vice.com/en_us/article/59jgka/a-brief-history-of-youtube-censorship.
- [149] National People's Congress (NPC) of the People's Republic of China. *CONSTITUTION OF THE PEOPLE'S REPUBLIC OF CHINA*, Accessed Oct, 2018. http://www.npc.gov.cn/englishnpc/Constitution/node_2825.htm.
- [150] Network Working Group, IETF. *RFC 4033: DNS Security Introduction and Requirements*, Accessed 8th June, 2017. <https://tools.ietf.org/html/rfc4033>.
- [151] Network Working Group, IETF. *RFC 1738: Uniform Resource Locators (URL)*, Accessed April, 2019. <https://tools.ietf.org/html/rfc1738>.
- [152] Network Working Group, IETF. *RFC 882: DOMAIN NAMES - CONCEPTS and FACILITIES*, Accessed Oct, 2018. <https://tools.ietf.org/html/rfc882>.
- [153] Network Working Group, IETF. *RFC 883: DOMAIN NAMES - IMPLEMENTATION and SPECIFICATION*, Accessed Oct, 2018. <https://tools.ietf.org/html/rfc883>.
- [154] ONI. *OpenNet Initiative*, 2018 (accessed Oct, 2018). <https://opennet.net/about-oni>.
- [155] PlanetLab. *PlanetLab*, 2018 (accessed Oct, 2018). <https://www.planet-lab.org/>.
- [156] Proton Technologies AG. *ProtonMail is being blocked in Turkey. Here's how to bypass Turkey's online censorship.*, Accessed Oct, 2018. <https://protonmail.com/blog/turkey-online-censorship-bypass/>.
- [157] Psiphon Inc. *Psiphon*, Accessed Oct 2018. <https://psiphon.ca/>.
- [158] Radio Free Asia. *China Orders Xinjiang's Android Users to Install App That Deletes 'Terrorist' Content*, Accessed Oct, 2018. <https://www.rfa.org/english/news/china/china-orders-xinjiangs-android-users-to-install-app-that-deletes-terrorist-content-07142017102032.html>.
- [159] Reuters. *Cuba government filtering mobile text messages, dissidents say*, Accessed Oct, 2018. <https://www.reuters.com/article/us-cuba-censorship-idUSKCN11B265>.
- [160] Reuters. *Here's How China Is Able to Censor WhatsApp and Other Chat Apps*, Accessed Oct, 2018. <https://www.inverse.com/article/36814-china-advances-whatsapp-blocking>.
- [161] Reuters. *Vietnam unveils 10,000-strong cyber unit to combat 'wrong views'*, Accessed Oct, 2018. <https://www.reuters.com/article/us-vietnam-security-cyber/vietnam-unveils-10000-strong-cyber-unit-to-combat-wrong-views-idUSKBN1EK0XN>.
- [162] Riot Games. *RIOT'S APPROACH TO ANTI-CHEAT*, 2018 (accessed July, 2018). <https://engineering.riotgames.com/news/riots-approach-anti-cheat>.

- [163] RIPE. *RIPE Atlas*, 2018 (accessed Oct, 2018). <https://atlas.ripe.net/>.
- [164] Ryan Gallagher. *GOOGLE PLANS TO LAUNCH CENSORED SEARCH ENGINE IN CHINA, LEAKED DOCUMENTS REVEAL*, Accessed Oct, 2018. <https://theintercept.com/2018/08/01/google-china-search-engine-censorship/>.
- [165] Schneier on Security. *How Signal Is Evading Censorship*, Accessed Oct, 2018. https://www.schneier.com/blog/archives/2016/12/how_signal_is_e.html.
- [166] Seth Grimes. Unstructured data and the 80 percent rule, Accessed 15th May, 2017. <https://breakthroughanalysis.com/2008/08/01/unstructured-data-and-the-80-percent-rule/>.
- [167] Simon Denyer. *Apple CEO backs China's vision of an 'open' Internet as censorship reaches new heights*, Accessed Oct, 2018. <https://www.washingtonpost.com/news/worldviews/wp/2017/12/04/apple-ceo-backs-chinas-vision-of-an-open-internet-as-censorship-reaches-new-heights/>.
- [168] Sui-Lee Wee and Li Yuan. *China Censors Bad Economic News Amid Signs of Slower Growth*, Accessed Oct, 2018. <https://www.nytimes.com/2018/09/28/business/china-censor-economic-news.html>.
- [169] The federal Council, The Portal of the Swiss Government. *Federal Act on Copyright and Related Rights*, 2018 (accessed July, 2018). <https://www.admin.ch/opc/en/classified-compilation/19920251/index.html>.
- [170] The Guardian. *To censor or not to censor? YouTube's double bind*, Accessed Oct, 2018. <https://www.theguardian.com/technology/2017/mar/21/youtube-advertisers-censorship>.
- [171] The OONI Project. The open observatory of network interference, Accessed 18th February, 2016. <https://ooni.torproject.org/>.
- [172] The Register. *Google, AWS IPs blocked by Russia in Telegram crackdown*, Accessed Oct, 2018. https://www.theregister.co.uk/2018/04/17/russia_blocks_google_aws_ip_addresses_to_get_telegram/.
- [173] The Tor Project. *The Tor Network*, Accessed 9th June, 2017. <https://www.torproject.org/>.
- [174] The Tor Project. *Tor Metrics*, Accessed Oct, 2018. <https://metrics.torproject.org/>.
- [175] The Tor Project. *Tor metrics: top-10 countries by possible censorship events*, Accessed Oct, 2018. <https://metrics.torproject.org/userstats-censorship-events.html>.
- [176] The Tor Project. *Tor partially blocked in China*, Accessed Oct, 2018. <https://blog.torproject.org/tor-partially-blocked-china>.
- [177] United Nations Statistics Division. *GDP and its breakdown at current prices in US Dollars*, 2017 (accessed July, 2018). <http://unstats.un.org/unsd/snaama/dntransfer.asp?fid=2>.
- [178] Wireshark Foundation. *Wireshark*, Accessed 12th June, 2017. <https://www.wireshark.org/>.
- [179] Yvette Tan. *China just banned livestreaming because it's too hard to censor*, Accessed Oct, 2018. <https://mashable.com/2017/06/23/china-bans-livestreaming/?europa=true&from=groupmessage#QdGH7gR3yqqI>.
- [180] Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. Automatic keyword extraction from individual documents. *Text mining: applications and theory*, pages 1–20, 2010.

- [181] Gerard Salton and Michael J McGill. Introduction to modern information retrieval. 1986.
- [182] Will Scott, Thomas Anderson, Tadayoshi Kohno, and Arvind Krishnamurthy. Satellite: Joint analysis of CDNs and network-level interference. In *USENIX Annual Technical Conference*. USENIX, 2016.
- [183] Will Scott, Sujit Packiaraj, and Arvind Krishnamurthy. Detecting dns censorship without an internal vantage point. 2013.
- [184] Andreas Sfakianakis, Elias Athanasopoulos, and Sotiris Ioannidis. Censmon: A web censorship monitor. *Free and Open Communications on the Internet*. USENIX, page 113, 2011.
- [185] Efe Kerem Sozeri. *Ban against a single blog post leads Turkish ISPs to censor all of WordPress*, 2015 (accessed March, 2019). <https://www.dailydot.com/layer8/turkey-wordpress-censorship-block/>.
- [186] Edward P Stabler. 13 the finite connectivity of linguistic structure. *Perspectives on sentence processing*, 2015.
- [187] Internet World Stats. *INTERNET GROWTH STATISTICS*, Accessed May 2019. <https://www.internetworldstats.com/emarketing.htm>.
- [188] Rima S. Tanash, Zhouhan Chen, Tanmay Thakur, Dan S. Wallach, and Devika Subramanian. Known unknowns: An analysis of Twitter censorship in Turkey. In *Workshop on Privacy in the Electronic Society*. ACM, 2015.
- [189] Jingrong Tong. Press self-censorship in china: A case study in the transformation of discourse. *Discourse & Society*, 20(5):593–612, 2009.
- [190] Yeganeh Torbati. *Iran blocks use of tool to get around Internet filter*, Accessed Oct 2017. <https://www.reuters.com/article/us-iran-internet/iran-blocks-use-of-tool-to-get-around-internet-filter-idUSBRE9290CV20130310>.
- [191] Yuen-Hsien Tseng, Chi-Jen Lin, and Yu-I Lin. Text mining techniques for patent analysis. *Information Processing & Management*, 43(5):1216–1247, 2007.
- [192] Ben VanderSloot, Allison McDonald, Will Scott, J. Alex Halderman, and Roya Ensafi. Quack: Scalable remote measurement of application-layer censorship. In *USENIX Security Symposium*, 2018.
- [193] John-Paul Verkamp and Minaxi Gupta. Inferring mechanics of web censorship around the world. In *Free and Open Communications on the Internet*. USENIX, 2012.
- [194] Barney Warf. Geographies of global internet censorship. *GeoJournal*, 76(1):1–23, 2011.
- [195] Pernille Warrer, Ebba Holme Hansen, Lars Juhl-Jensen, and Lise Aagaard. Using text-mining techniques in electronic patient records to identify adrs from medicine use. *British journal of clinical pharmacology*, 73(5):674–684, 2012.
- [196] Nicholas Weaver, Robin Sommer, and Vern Paxson. Detecting forged TCP reset packets. In *Network and Distributed System Security*. The Internet Society, 2009.
- [197] WebShrinker. *WebShrinker Categories API*, Accessed May 2019. <https://www.webshrinker.com/>.

- [198] Zachary Weinberg, Mahmood Sharif, Janos Szurdi, and Nicolas Christin. Topics of controversy: An empirical analysis of web censorship lists. *Proceedings on Privacy Enhancing Technologies*, 2017(1):42–61, 2017.
- [199] Philipp Winter. Towards a censorship analyser for tor. In *3rd USENIX Workshop on Free and Open Communications on the Internet*. USENIX-The Advanced Computing Systems Association, 2013.
- [200] Sebastian Wolfgarten. Investigating large-scale Internet content filtering. Technical report, Dublin City University, 2006.
- [201] Gilbert Wondracek, Thorsten Holz, Christian Platzer, Engin Kirda, and Christopher Kruegel. Is the internet for porn? an insight into the online adult industry. In *WEIS*, 2010.
- [202] Joss Wright. Regional variation in chinese internet filtering. *Information, Communication & Society*, 17(1):121–141, 2014.
- [203] Joss Wright, Alexander Darer, and Oliver Farnan. On identifying anomalies in tor usage with applications in detecting internet censorship. In *Proceedings of the 10th ACM Conference on Web Science*, pages 87–96. ACM, 2018.
- [204] Joss Wright, Tulio Souza, and Ian Brown. Fine-grained censorship mapping: Information sources, legality and ethics. In *Free and Open Communications on the Internet*. USENIX, 2011.
- [205] Ruohan Xiong and Jeffrey Knockel. An efficient method to determine which combination of keywords triggered automatic filtering of a message. In *Free and Open Communications on the Internet*. USENIX, 2019.
- [206] Xueyang Xu, Z. Morley Mao, and J. Alex Halderman. Internet censorship in China: Where does the filtering occur? In *Passive and Active Measurement Conference*, pages 133–142. Springer, 2011.
- [207] Xueyang Xu, Z Morley Mao, and J Alex Halderman. Internet censorship in china: Where does the filtering occur? In *International Conference on Passive and Active Network Measurement*, pages 133–142. Springer, 2011.
- [208] Tarun Kumar Yadav and Sambuddho Chakravarty. Trends and patterns of internet censorship in india. 2018.
- [209] Todor Yakimov, JJ van der Ham, and Barry van Kampen. Detecting routing anomalies with ripe atlas, 2014.
- [210] Tao Zhu, David Phipps, Adam Pridgen, Jedidiah R Crandall, and Dan S Wallach. Tracking and quantifying censorship on a chinese microblogging site. *arXiv preprint arXiv:1211.6166*, 2012.
- [211] Tao Zhu, David Phipps, Adam Pridgen, Jedidiah R Crandall, and Dan S Wallach. The velocity of censorship: High-fidelity detection of microblog post deletions. In *USENIX Security Symposium*, pages 227–240, 2013.
- [212] Jonathan Zittrain and Benjamin Edelman. Internet filtering in china. *IEEE Internet Computing*, 7(2):70–77, 2003.

Appendices

Lists of Discovered Filtered Domains

All lists correct as of July 2018

A.1 China

10bestvps.co	archive.org	blogspot.lu	chinachange.org.ws
10bestvps.com	arkhangelsk.su	blogspot.md	chinadigitaltimes.net
10bestvps.info	armenia.su	blogspot.mx	chinesepen.org
10bestvps.org	ashgabad.su	blogspot.pe	chirurgiens-dentistes-en-france.fr
12hp.at	asiaexchange.org	blogspot.pt	ciudadblogger.com
12hp.ch	ask.com	blogspot.qa	civic-exchange.org
12hp.de	astrill.com	blogspot.rs	clan.rip
12vpn.com	atozproxy.com	blogspot.se	cidmail.ru
12vpn.net	avaaz.org	blogspot.si	clearwisdom.net
12vpn.org	avira.com	blogspot.sk	cloudapp.net
1337-pictures	azerbaijan.su	blogspot.sn	cloudcontrolapp.com
lkapp.com	backchina.com	blogspot.ug	cloudfunctions.net
lstyoutube.com	backplaneapp.io	blogspot.cn	cloudns.asia
2ix.at	bajaryoutube.com	bloomberg.com	cloudns.biz
2ix.ch	balashov.su	bloomberg.com	cloudns.cc
2ix.de	baq.kz	bloxcms.com	cloudns.club
4archive.org	bary.net	bmoattachments.org	cloudns.eu
4everproxy.com	baryonline.com	bolevpn.net	cloudns.in
4lima.at	bashkiria.ru	boomla.net	cloudns.info
4lima.ch	bashkiria.su	booyoutube.com	cloudns.org
4lima.de	bestvideoseonyoutube.com	boxun.com	cloudns.pro
4shared.com	bestvpn.com	boxun.com.ba	cloudns.pw
4u.com	bir.ru	boxun.com.cn	cloudns.us
6park.com	bitballoon.com	boxun.com.ph	cn.com
abkhazia.su	bitcointalk.org	boxun.com.ws	co.com
acevpn.com	biz.at	boxun.computerhistory.museum	co.nl
adygeya.ru	blackvpn.com	boxun.info	co.pl
adygeya.su	blockless.com	bplaced.com	codespot.com
ae.org	blogger.com	bplaced.de	cogsoz.com
airvpn.org	bloglovin.com	bplaced.net	com.de
aktyubinsk.su	blogspot.al	br.com	com.ru
amazonaws.com	blogspot.am	braceroarchive.org	com.se
amazonaws.com.cn	blogspot.ba	britishbeautyblogger.com	comprarevisualizzazioniyoutube.com
amnesty.org	blogspot.be	britishrecordsshoparchive.org	convyoutube.com
amnesty.org.hk	blogspot.bg	browsec.com	cookieparts.com
anchorfree.com	blogspot.ca	browsersafetymark.io	copyblogger.com
apiyoutube.com	blogspot.cl	bryansk.su	corvetteblogger.com
apk-dl.com	blogspot.co.at	bukhara.su	coursehero.com
apk-dl.com.ba	blogspot.co.il	buy-instagram.com	creaders.net
apk-dl.com.ph	blogspot.co.nz	bv.nl	crvdcntrl.net
apk-dl.com.us	blogspot.co.za	byen.site	cupcake.is
apk-dl.computerhistory.museum	blogspot.com	c.la	cyberghostvpn.com
apkpure.com	blogspot.cz	cactusvpn.com	cyon.link
apnayoutube.com	blogspot.dk	chg.ru	daddyblogger.com
appchizi.com	blogspot.fi	cdn77-secure.org	dagestan.ru
applinzi.com	blogspot.gr	change.org	dagestan.su
appsport.com	blogspot.hr	channelmeter.com	dailymotion.com
archive.is	blogspot.hu	childrenspoetryarchive.org	dalailama.cl
archive.li	blogspot.in	chimeforchange.org	dalailama.com
archive.limited	blogspot.is	chimkent.su	daplie.me
archive.link	blogspot.li	chinachange.org	dd-dns.de
archive.live	blogspot.lt	chinachange.org.ph	ddns.net

ddns.de	feste-ip.net	google.com.co	google.se
de.com	fhapp.xyz	google.com.cu	google.sh
de.cool	filedir.com	google.com.cy	google.sk
definima.net	filegear.me	google.com.do	google.sm
deletefacebook.com	filmesdoyoutube.com	google.com.dz	google.sn
desuarchive.org	filterbypass.me	google.com.ec	google.tm
developerblogger.com	finchvpn.com	google.com.ee	google.to
disconnect.me	finnishtwitter.com	google.com.eg	google.tt
disqus.com	firebaseapp.com	google.com.et	google.vg
dnsalias.net	fkyoutube.com	google.com.ge	google.vu
dnsfor.me	flickr.com	google.com.gh	google.ws
dongtaivang.com	flushyoutube.com	google.com.gi	googlelepis.com
dotvpn.com	fogf.org	google.com.gp	googlecode.com
dotvpn.com.ba	free-proxysite.com	google.com.gr	gotpantheon.com
dotvpn.com.ph	freeddns.org	google.com.gt	gr.com
dotvpn.com.us	freedomhouse.org	google.com.gy	greatfire.org
dotvpn.computerhistory.museum	freemusicarchive.org	google.com.hk	grozny.ru
downladeryoutube.com	freetibet.org	google.com.ht	grozny.su
dragonblogger.com	freevpn.be	google.com.iq	gumroad.com
drop-dropbox.com	freevpn.ca	google.com.ja	gumroad.com.ba
dropbox.com	freevpn.cc	google.com.jo	gumroad.com.cn
drud.io	freevpn.center	google.com.kz	gumroad.com.ph
drud.us	freevpn.co	google.com.lb	gumroad.com.us
duckdns.org	freevpn.org	google.com.lv	gumroad.computerhistory.museum
duckduckgo.com	freevpn.tv	google.com.ly	gunviolencearchive.org
dvrDNS.org	frootvpn.com	google.com.mt	harunyahya.com
dw.com	fujiyoutube.com	google.com.mv	hepforge.org
dxyoutube.com	futuremailing.at	google.com.mx	herokuapp.com
dynalias.net	gamer.com.tw	google.com.my	hide.io
dyndns-office.com	gamesforchange.org	google.com.na	hide.me
dyndns-remote.com	gb.com	google.com.ng	hideipvpn.com
dyndns-web.com	gb.net	google.com.ni	hideman.net
dyndns.info	georgia.su	google.com.nr	hidemyass.biz
dyndns.org	get-link-youtube.com	google.com.om	hidemyass.com
dyndns1.de	getcloak.com	google.com.pa	hidemyass.com.ba
dynu.net	geti2p.net	google.com.pe	hidemyass.com.ph
dynv6.net	geti2p.net.cm	google.com.pg	hidemyass.com.ws
dynvpn.de	geti2p.net.ph	google.com.ph	hidemyass.computerhistory.museum
earthcam.com	geti2p.net.ws	google.com.pk	hidester.com
east-kazakhstan.su	getlinksoundcloud.com	google.com.pl	hk.com
eastturkistan.net	getmediayoutube.com	google.com.pr	hola.org
eastturkistan-gov.org	git-repos.de	google.com.ps	home-websserver.de
eastturkistaninfo.com	github.com	google.com.pt	homeip.net
echofon.com	github.io	google.com.py	homelink.one
eksisozluk.com	githubarchive.org	google.com.qa	homelinux.net
eltuboadventista.com	githubusercontent.com	google.com.sa	homelinux.org
elyoutube.com	gitlab.io	google.com.sb	homeunix.com
emojistwitter.com	goldenfrog.ch	google.com.sg	homeunix.net
epochtimes.co.il	goldenfrog.com	google.com.sl	hongkongfp.com
epochtimes.co.kr	google.ae	google.com.sv	hootsuite.com
epochtimes.com	google.am	google.com.tj	hopto.me
epochtimes.com.au	google.as	google.com.tn	hopto.org
epochtimes.com.ba	google.at	google.com.tr	hotspotshield.com
epochtimes.com.br	google.az	google.com.tw	hrv.org
epochtimes.com.hk	google.ba	google.com.ua	hu.net
epochtimes.com.mx	google.be	google.com.uy	hulkshare.com
epochtimes.com.my	google.bg	google.com.vc	hurgokbayrak.com
epochtimes.com.ph	google.bi	google.com.ve	ibvpn.com
epochtimes.com.pt	google.bs	google.com.vn	ibvpn.net
epochtimes.com.sg	google.ca	google.computerhistory.museum	icij.org
epochtimes.com.tw	google.cd	google.de	info.at
epochtimes.com.ua	google.cg	google.dj	instagram.com
epochtimes.com.us	google.ch	google.dk	int.ru
epochtimes.computerhistory.museum	google.ci	google.dm	ipvanish.com
epochtimes.de	google.cl	google.es	ironsocket.com
epochtimes.jp	google.co.bw	google.fi	ironsocket.com.ba
epochtimes.ru	google.co.cr	google.fm	ironsocket.com.ph
epochtimes.se	google.co.id	google.fr	ironsocket.com.ws
eu.com	google.co.hu	google.gg	ironsocket.computerhistory.museum
evemnode.com	google.co.il	google.gl	is-lost.org
exnet.su	google.co.im	google.gm	islamhouse.com
expressvpn.asia	google.co.in	google.gr	issuu.com
expressvpn.biz	google.co.je	google.hn	ivanovo.su
expressvpn.cc	google.co.jp	google.hr	ixquick.com
expressvpn.center	google.co.kr	google.ht	jambyl.su
expressvpn.club	google.co.ls	google.ie	jihadology.net
expressvpn.co	google.co.ma	google.is	jokeblogger.com
expressvpn.com	google.co.nz	google.it	jp.net
expressvpn.cool	google.co.th	google.jo	jpn.com
expressvpn.digital	google.co.uk	google.kg	justpaste.it
expressvpn.expert	google.co.uz	google.kz	kalachakra2017.com
expressvpn.hk	google.co.ve	google.li	kalachakranet.org
expressvpn.info	google.co.vi	google.lk	kalmykia.ru
expressvpn.la	google.co.za	google.lt	kalmykia.su
expressvpn.mobi	google.co.zm	google.lu	kaluga.su
expressvpn.net	google.com	google.lv	karacol.su
expressvpn.solutions	google.com.af	google.mn	karaganda.su
expressvpn.tv	google.com.ag	google.ms	karelia.su
expressvpn.works	google.com.ai	google.mu	khakassia.su
expressvpn.ws	google.com.ar	google.mw	kobo.com
facebook.com	google.com.au	google.nl	kobobooks.com
faceless.me	google.com.ba	google.no	kpopinstagram.com
falun-dafa.net	google.com.bd	google.nr	kproxy.asia
falun-ny.net	google.com.bh	google.nu	kproxy.com
falundafa-nc.org	google.com.bi	google.off.ai	krasnodar.su
faluninfo.de	google.com.bo	google.pl	kurgan.su
faluninfo.net	google.com.br	google.pn	kustanai.ru
fanqianghou.com	google.com.by	google.pt	kustanai.su
fastly.net	google.com.bz	google.ro	lair.io
favoritewords.com	google.com.cn	google.ru	lataayoutube.com
fbsbx.com		google.rw	le-vpn.com
fedorapeople.org		google.sc	lenug.su

lima-city.at	opencraft.hosting	openvpn.sk	rsf.org.uk
lima-city.ch	opendemocracy.net	openvpn.solutions	rsf.org.ws
lima-city.de	openvpn.ac.cn	openvpn.space	ru.com
lima-city.rocks	openvpn.aip.ee	openvpn.systems	ru.net
lima.zone	openvpn.am	openvpn.tci	sa.com
line.me	openvpn.android	openvpn.tk	safervpn.asia
liquidvpn.com	openvpn.anquan	openvpn.top	safervpn.cc
listentoyoutube.com	openvpn.aquila.it	openvpn.us	safervpn.com
live.com	openvpn.arab	openvpn.vg	safervpn.mobi
loginto.me	openvpn.asia	openvpn.vip	safervpn.net
logmein.com	openvpn.at	openvpn.xihuan	savetibet.nl
loudtwitter.com	openvpn.belau.pw	openvpn.xin	savetibet.org
lparchive.org	openvpn.bjarkoy.no	openvpn.xyz	savetibetstore.org
lyrics-youtube.com	openvpn.blog	openvpn.yun	scholarlyexchange.org
maarip.org	openvpn.ca	operaunite.com	se.net
maarip.org.ph	openvpn.cal	org.ru	searchman.com
maarip.org.ws	openvpn.cc	outsystemscloud.com	secretunnel.com
mangyshlak.su	openvpn.cf	overplay.net	securitykiss.com
marine.ru	openvpn.chrome	pagefrontapp.com	seed4.me
masyoutube.com	openvpn.cn	pagespeedmobilizer.com	selfip.com
medium.com	openvpn.co	paltalk.com	selfip.info
mems-exchange.org	openvpn.co.cm	panoramio.com	sendtodropbox.com
meteorapp.com	openvpn.co.pw	peacehall.com	serveftp.net
mex.com	openvpn.com	peliculasyoutube.com	servegame.com
micrologger.com	openvpn.com.au	penza.su	shadowsocks.biz
mine.nu	openvpn.com.ba	phayul.com	shadowsocks.com
minghui.ca	openvpn.com.br	pinnacle.com	shadowsocks.com.au
minghui.cc	openvpn.com.cn	pixolino.com	shadowsocks.com.hk
minghui.de	openvpn.com.ph	pkarchive.org	shadowsocks.company
minghui.org	openvpn.com.ua	poetryarchive.org	shadowsocks.gift
minghui.tv	openvpn.com.ws	pokerstars.com	shadowsocks.im
minghuischool.us	openvpn.computerhistory.museum	pokrovsk.su	shadowsocks.info
mingpao.com	openvpn.cz	popularyoutube.com	shadowsocks.me
mitbs.ca	openvpn.dclik	privateinternetaccess.com	shadowsocks.network
mitbs.com	openvpn.de	privatetunnel.com	shadowsocks.org
mitbs.org	openvpn.dev	privatevpn.com	shadowsocks.software
mooo.com	openvpn.drive	privatevpn.org	shadowsocks.space
mordovia.ru	openvpn.ed.pw	probblogger.com	shadowsocks.team
mordovia.su	openvpn.edu.ws	proxfree.com	shadowsocks.work
morphism.info	openvpn.es	proxies.asia	shenyun.com
movements.org	openvpn.etisalat	proxies.pw	shenyunperformingarts.org
mp3-juices.com	openvpn.eu	proxies.sx	simbolostwitter.com
msexchange.org	openvpn.fm	proxpn.biz	sinaapp.com
msk.ru	openvpn.fr	proxpn.com	siteintelgroup.com
msk.su	openvpn.ga	proxy-list.org	skyoutube.com
murmask.su	openvpn.gle	proxy-site.us	smartblogger.com
my-private-network.co.uk	openvpn.go.pw	proxy.business	smartdnsproxy.com
my-router.de	openvpn.gov.ve	proxy.cab	smhric.org
mybluemix.net	openvpn.gov.vc	proxy.cm	sochi.su
mydropbox.com	openvpn.guge	proxy.college	softonic.aip.ee
myfusion.cloud	openvpn.hangout	proxy.fit	softonic.android
myhome-server.de	openvpn.hu	proxy.forsale	softonic.anquan
myshopblocks.com	openvpn.in	proxy.fyi	softonic.anquan
mytis.ru	openvpn.io	proxy.mba	softonic.arab
myvnc.com	openvpn.ir	proxy.nu	softonic.belau.pw
nalchik.ru	openvpn.jp	proxy.org	softonic.bjarkoy.no
nalchik.su	openvpn.kr	proxy.poker	softonic.bz
nationofchange.org	openvpn.la	proxy.porn	softonic.ca
navoi.su	openvpn.lt	proxy.sx	softonic.cal
netlify.com	openvpn.map	proxy.web.id	softonic.cc
ngrok.io	openvpn.md.ci	proxy.wiki	softonic.cf
nicotwitter.com	openvpn.me	proxy.world	softonic.chrome
no-ip.biz	openvpn.mil.ph	proxy.yt	softonic.cm
no-ip.net	openvpn.mobi	proxygerman.com	softonic.cn
no-ip.org	openvpn.mosvik.no	proxygerman.com.ba	softonic.co
nodeart.io	openvpn.ne.pw	proxygerman.com.ph	softonic.co.cm
nodum.co	openvpn.net	proxygerman.com.ws	softonic.co.in
nodum.io	openvpn.net.cm	proxygerman.computerhistory.museum	softonic.co.pw
noip.me	openvpn.net.ph	proxynova.com	softonic.com
nom.al	openvpn.net.ws	proxyserver.com	softonic.com.au
nom.im	openvpn.network	proxysite.club	softonic.com.ba
nom.si	openvpn.nexus	proxysite.com	softonic.com.br
nonviolent-conflict.org	openvpn.ngo.ph	proxysite.fr	softonic.com.cn
nordvpn.biz	openvpn.nl	proxysite.it	softonic.com.ph
nordvpn.com	openvpn.no	proxysite.nl	softonic.com.us
north-kazakhstan.su	openvpn.nom.za	proxysite.org	softonic.computerhistory.museum
nov.ru	openvpn.nowruz	proxysite.pw	softonic.dclik
nov.su	openvpn.online	proxysite.tk	softonic.de
nov.sh	openvpn.or.pw	proxyturbo.com	softonic.dev
ntdtv.ca	openvpn.org	ps3youtube.com	softonic.dk
ntdtv.com	openvpn.org.cn	ptplus.fit	softonic.drive
ntdtv.com.ba	openvpn.org.ph	ptt.cc	softonic.ed.pw
ntdtv.com.mx	openvpn.org.pl	publicartarchive.org	softonic.edu.ws
ntdtv.com.ph	openvpn.org.ws	puffinbrowser.com	softonic.es
ntdtv.com.tw	openvpn.ovh	pureinsight.org	softonic.etisalat
ntdtv.com.ws	openvpn.pars	pyatigorsk.ru	softonic.eu
ntdtv.computerhistory.museum	openvpn.ph	rackmaze.com	softonic.fm
nuclearweaponarchive.org	openvpn.phd	rael.org	softonic.fr
nym.by	openvpn.pl	realchange.org	softonic.gdn
nym.me	openvpn.play	reklamyoutube.com	softonic.gle
nym.tw	openvpn.politie	reuters.com	softonic.go.pw
nytimes.com	openvpn.presse.ml	rfi.fr	softonic.gov.ve
obninsk.su	openvpn.pw	rimmer.su	softonic.gov.vc
octanevpn.com	openvpn.red	rohingyablogger.com	softonic.gq
octanevpn.com.ba	openvpn.ren	rothschildarchive.org	softonic.guge
octanevpn.com.ph	openvpn.ru	rsf-chinese.org	softonic.hangout
octanevpn.com.ws	openvpn.sch.lk	rsf.org	softonic.in
octanevpn.computerhistory.museum	openvpn.se	rsf.org.au	softonic.it
okayfreedom.com	openvpn.search	rsf.org.cn	softonic.jp
omnitalk.com	openvpn.shia	rsf.org.in	softonic.la
on-web.fr	openvpn.shouji	rsf.org.ph	softonic.map
onthewifi.com	openvpn.site	rsf.org.pk	softonic.md.ci

softonic.mil.ph
softonic.mosvik.no
softonic.mx
softonic.name
softonic.ne.pw
softonic.net
softonic.net.cm
softonic.net.ph
softonic.net.ws
softonic.nexus
softonic.ngo.ph
softonic.nl
softonic.no
softonic.nom.za
softonic.nowruz
softonic.or.pw
softonic.org
softonic.org.ph
softonic.org.us
softonic.pars
softonic.phd
softonic.pl
softonic.play
softonic.politie
softonic.presse.ml
softonic.pw
softonic.sch.lk
softonic.se
softonic.search
softonic.shia
softonic.shouji
softonic.store
softonic.tci
softonic.top
softonic.tv
softonic.us
softonic.vg
softonic.vip
softonic.wis
softonic.xihuan
softonic.xxx
softonic.xyz
softonic.yun
solucionfacebook.com
soundcloud.com
southfront.org
spb.ru
spb.su
spdns.de
spdns.org
speakerdeck.com
spideroak.com
spotflux.com
square7.ch
square7.de
square7.net
sraffaarchive.org
sslsecureproxy.com
sslunlocked.com
start-vpn.com
startpage.com
static.land
status-for-facebook.com
stepchange.org
storify.com
strongvpn.com
stuartxchange.org
studentsforafreetibet.org
sumrando.com
sunporno.com
sunvpn.net
superyoutube.com
suprememastertv.com
svn-repos.de
tapatalk.com
tashkent.su
telegram.org
termez.su
thebobs.com
thebrickblogger.com
theepochtimes.com
thenewslens.com
thepiratebay.org
thestandnews.com
thetoryexchange.org
thingdust.io
thyroidchange.org
tibet.net
tibet.org.tw
tibetaction.net
tibethouse.org
tibethouse.us
tibetnetwork.org
tibetsun.com
tibetruth.com
togliatti.su
toolur.com
top-proxies.co.uk
toptwitter.com
torguard.net
torguard.tv
torproject.org
torrentz.eu
torvpn.com
tosaveyoutube.com
totalvpn.com
townnews-staging.com
tr-youtube.com
trafficplex.cloud
troitsk.su
trucchifacebook.com
tselinohrad.su
tubewolf.com
tula.su
tumblr.com
tunnelbear.com
tunnelguru.com
tunnelr.com
turkiyefacebook.com
tuva.su
tuxfamily.org
tv-shows-youtube.com
tweetdeck-twitter.com
twelve.today
twiends.com
twister.net.co
twitpic.com
twitter.com
twittercounter.com
uhrp.org
uk.com
uk.net
unlock-everything.com
unlock-everything.com
unlockfreeproxy.com
unlockvideos.com
unlockvideos.com.ba
unlockvideos.com.ph
unlockvideos.com.ws
unlockvideos.computerhistory.museum
unlockweb.co
unlockweb.gq
unlockweb.net
unlockyoutube.co
unlockyoutube.com
unlockyoutube.me
unlockyoutube.us
us.com
us.org
ushahidi.com
uyghur.kim
uyghuramerican.org
uyghurcongress.org
vapor.cloud
vaporcloud.io
vdyoutube.com
venetianmacao.com
veooz.com
viber.com
vid.me
videogamesblogger.com
vimeo.com
vipsinaapp.com
viyoutube.com
vladikavkaz.ru
vladikavkaz.su
vladimir.ru
vladimir.su
voachinese.com
volodga.su
voxer.com
vpnbook.com
vpndada.com
vpnmag.fr
vpnme.me
vpnmentor.com
vpnreactor.com
vpnsecure.me
vpntraffic.aip.ee
vpntraffic.android
vpntraffic.anquan
vpntraffic.aquila.it
vpntraffic.arab
vpntraffic.belau.pw
vpntraffic.bjarkoy.no
vpntraffic.cal
vpntraffic.chrome
vpntraffic.co.cm
vpntraffic.co.pw
vpntraffic.com
vpntraffic.com.ba
vpntraffic.com.ph
vpntraffic.com.ws
vpntraffic.computerhistory.museum
vpntraffic.dclk
vpntraffic.dev
vpntraffic.drive
vpntraffic.ed.pw
vpntraffic.edu.us
vpntraffic.etisalat
vpntraffic.fu
vpntraffic.gle
vpntraffic.go.pw
vpntraffic.gob.ve
vpntraffic.gov.vc
vpntraffic.guge
vpntraffic.hangout
vpntraffic.la
vpntraffic.map
vpntraffic.md.ci
vpntraffic.mil.ph
vpntraffic.mosvik.no
vpntraffic.ne.pw
vpntraffic.net.cm
vpntraffic.net.ph
vpntraffic.net.ws
vpntraffic.nexus
vpntraffic.ngo.ph
vpntraffic.nom.za
vpntraffic.nowruz
vpntraffic.or.pw
vpntraffic.org.ph
vpntraffic.org.us
vpntraffic.pars
vpntraffic.ph
vpntraffic.phd
vpntraffic.play
vpntraffic.politie
vpntraffic.presse.ml
vpntraffic.pw
vpntraffic.sch.lk
vpntraffic.search
vpntraffic.shia
vpntraffic.shouji
vpntraffic.tci
vpntraffic.vg
vpntraffic.ws
vpntraffic.xihuan
vpntraffic.yun
vpntunnel.com
wantyoutube.com
warayblogger.com
wattpad.com
wearchange.org
webproxy.ca
webproxy.com
webproxy.id
webproxy.li
webproxy.online
webproxy.pw
webproxy.ru
webproxy.to
webproxy.us
webproxy.yt
webspace.rocks
webtunnel.com
webtunnel.org
wedeploy.io
wedeploy.me
wedeploy.sh
wellbeingzone.co.uk
wellbeingzone.eu
wenxuecity.com
wikileaks.com
wikileaks.org
wikimapia.org
wikimapia.org.in
wikimapia.org.ph
wikimapia.org.ua
wikimapia.org.ws
wikipedia.org
wikipedia.org.cn
wikipedia.org.il
wikipedia.org.in
wikipedia.org.ph
wikipedia.org.pl
wikipedia.org.ws
windscribe.com
withgoogle.com
withyoutube.com
wordpress.com
wsj.com
xltovens.com
xyzyoutube.com
yahoo.com
ybo.faith
ybo.party
ybo.review
ybo.science
ybo.trade
yolasite.com
yombo.me
yourprivatevpn.com
yourprivatevpn.net
youtube.com
youtube2mp3.cc
youtube2mp3.to
za.com
zacebook.com
zacebookpk.com
zend2.com
zenmate.com
zenmate.io
zerocensorship.com
zpn.in
zvuk.me

A.2 Indonesia

Adult domains (45,863) not shown

10bet.com
141jav.com
141jav.xyz
19.com
1date.nl
1staab.com
21x.com
21x.org
2link.be
2much.com
2much.net
2much.tv
2plus2.fr
2plus2.net
3design3.com
3xnews.com
40best.com
4hen.com
4hen.net
4realcash.com
4ucash.nl
5000fotos.com
50plus-treff.de
775533.com
8x8.be
9news.com.au
9volttaco.com
a-pic.net
a14k.com
aaaaaaa.com
acanthus.be
acceleratedweb.com
acefree.com
actual-host.com
actualcash.com
adameveshops.com
adbucks.com
adcycle.com
addict-to.com
adorableaudrey.com
adressex.com
adrianswebpage.com
adventuredating.com
aerobicise.com
affaire.com
affairmatch.com
afunnystuff.com
agenbola.club
agenbola.com
agenbola.io
agenbola.net
agreatsite.com
ahoo.com
airwoodmedia.com
akt.de
alexiscapriblog.net
alexistexasclub.com
alexz-traffic.com
alexzandra.com
alfamina.com
allariagiovanni.com
allbestclips.com
allosponsor.com
allshiny.com
allurecash.net
almacendefamosas.com
alsbikinis.com
alsscanfan.com
alternative-footwear.co.uk
alternativeconnections.com
alterskontrolle.de
amanojyaku.net
amazons.com
amazingaila.com
amazingjokes.com
amber-michaels.org
amber-michaels.ws
amor.at
amor.com
amorsi.com
amouret.nl
amsterdammarijuanaseeds.com
andrevanamstel.nl
aneros.com
angelsweb.com
angelsweb.nl
angryduck.com
anna-angel.com
annamills.com
anonymouse.com
anonymouse.org
anonymouse.ws
anotherite.co.uk
answering-islam.org
antweb.biz
antweb.cz
antweb.info
any-time.de
aoisola.net
apa-gbi.org
apcsites.com
apexglamour.com
aphrodite-travel.co.uk
aphrodite-travel.com
aphrodite-travel.nl
aprosupport.com
aquafinity.com
aquarium.tv
archenemys.com
archives.com
archivodefamosas.com
archivodefamosas.net
art-forum.org
asacp.com
asacp.org
asanava.com
ashleymadison.com
askzoe.com
athena2.net
atkcash.com
atkol.com
atkol.tv
atomicbooks.com
atsukokudo.com
attuworid.com
audreylive.com
autorank.nl
avenueblue.com
avenueblue.net
avnads.com
avnawards.com
avnlive.com
avnonline.com
awm-help.com
azz4ever.com
bacabacaquran.com
backpage.com
badoink.com
badoink.net
badtales.com
badtales.nl
balkanstgp.com
ballblasters.com
baraskit.se
barbaranitke.com
barbie-boy.com
bastardly.com
baxterstgp.com
bbcd.de
bbs-tv.com
beach4fun.com
beam.to
beer.com
beertraffic.biz
bellezasnordicas.com
bentvoices.org
bestofamsterdam.com
bet365.com
bet365.es
bet365.it
bet365affiliates.com
betfair.com
bethemask.com
bgafd.co.uk
bi.org
biancabeauchamp.com
biancaenrob.nl
biertijd.com
biertijd.xxx
bikini-collection.com
bikiniholiday.com
bikinilist.com
bindme.nl
bizar.com
bizar.dk
bizar.hu
bizar.net
bizar.nu
bizar.ws
blinkbits.com
blogspot.com
blogspot.org
blueangel.nl
blueblood.net
blueoot.com
blunts.com
bolinga.com
bonesteel.net
bonkwire.com
boogo.nl
boozetime.com
boredtown.com
bosstelenet.com
bovada.lv
boycherries.com
boytemper.com
braincash.com
bravedigger.net
breast.com
breastfiles.com
breasts.com
breasts.me.uk
breasts.org
breasts.ws
briantarsis.com
bride.ru
brizis.com
brizis.nl
brobible.com
buktidansaksi.com
bullwhip.co.uk
bullwhip.org
busyx.com
buzzlink.com
buzzsession.com
c.la
caesarspalace.de
atsukokudo.com
captain-outrageous.com
captiveculture.com
caramec.com
caramec.net
cartoonland.de
cartoonland.it
cartoonmodern.com
casey-parker.com
cash.com
cashcore.com
cashdorado.de
cashforge.com
cashmaniacs.com
castorcash.com
catalina-cruz.ws
catchycash.com
catfightpages.com
cathouseclothing.com
ccomg.com
cecash.com
censorwatch.co.uk
cfnm-movies.com
cfnm-pic.com
cfnm-story.com
cfnm-pic.com
cgi-works.net
chameleonsubmitter.com
charlielaine.com
charlielaine.org
charlielaine.us
chatbe.nl
chatten.co.at
chatten.nl
chatten.us
cheating-housewives.co.uk
cheating-housewives.com
cheating-housewives.net
cheatinghousewife.com
chez-asma.net
chicas.be
chicas.cc
chicas.ch
chicas.com
chicas.in
chicas.nu
chicas.tk
chicas.ws
chippendales.co.uk
chippendales.com
chippendales.com.pl
chippendales.de
chisel.com
chloel.com
chokedchicken.com
chokinchicken.com
chromeonline.com
citebeur.com
citebeur.fr
citebeur.net
citebeur.org
cleanseries.com
cleispress.com
clinched.net
clipsgrabber.com
clockworkcash.com
close2heaven.nl
club-absolut.de
club-de-sade.de
club-nikki.com
club-oase.de
club-secrets.com
club-secrets.de
club-sttrophez.de
clubaphroditesneek.com
clubcarmellabing.com
clubcenterfolds.com
clubdrenteboerderij.nl
clubelectricblue.com
clubmadonna.com
clubmadonna.nl
clubwideworld.com
clubxlive.com
cmd368.com
cmd368.net
cnv.com
cocoboyz.com
coed.com
coedcentral.com
collarme.com
collinstud.com
coltstudiogroup.com
com.ru
comalternativeconnections.net
comdotgame.com
comics-art.com
comics-house.com
comicsxd.com
comikzone.com
comrudevirtual.com
comsthumbs.com
condor-traffic.com
contactbox.nl
convertfamily.com
cookclips.com
coolbeans.com
cooxa.com
coquines.ch
coquines.com
coquines.info
coquines.org
coquines.ws
coral.co.uk
corpun.com
covet.ws
crazycasey.net
crimson-moon-ltd.com
crissycrankscars.com
crissymoran.com
crissymoran.ws
crissymoranblog.com
crushgiantess.com
cuentameya.com
daddy.com
dailymovies.biz
dailymovies.info
dailymovies.nl
dailymovies.org
dailymovies.tv
dailyniner.com
dailyrotten.com
dailytreats.com
dallasconnection.net
damesontvangen.nl
darenzia.com
darenzia.net
dasandyman.com
date-online.nl
date.com
date.info
date.nl
dateclub.be
datematch.com
datepage.nl
dating-portal.nl
dating-site.us
dating-site.ws
dating-websites.net
dating-websites.us
datingdirect.com
datinggold.com
datinghaus.com
datingnederland.com
datingservice.nl
datingservice.org
datingstart.com
datingstart.nl
dawnedire.com

dayom.com
dbasixx.com
ddgxxtreme.com
deezteez.com
degeilewebsite.nl
delights.com
delights.us
deltadivener.com
demicomix.com
denisemilani.com
der-kerker.de
descuidosdefamosas.com
desiderya.it
deviantart.com
devonsorlrd.com
devonsorlrd.net
devajoker.com
devaliga.com
die-landsauna.de
dieeule.de
discretos.net
dita.net
disite.com
dmm.co.jp
doggielist.com
dogging-site.co.uk
doing.my
dollars-paradise.com
dollhouse-agogo.com
donkeresletten.nl
donmai.us
doodmovies.com
dorcel.com
dorcel.fr
dorcel.tv
douche.com
dohou.com
downloadenaar.nl
downloadmaar.nl
downloads.nl
downloadsmovies.net
dragondungeon.com
dragonthumbz.com
drakaina.com
drbizaro.net
drbizaro.org
dreamstore.ch
dressed2play.com
drubskin.com
dubberley.com
dutchria.nl
dutchthumbs.net
dymantic.com
dynamix.net
eaglesmovies.com
easy-free.com
ebinacash.com
ecchiart.com
echangisme.biz
echangisme.net
edenwells.com
egafd.com
egotastic.com
ehowa.com
eigenstart.nl
ejhs.org
el-patio.be
elegant-sophisticated.net
elke-jeinsen.com
elles-se-mettent-nues-pour-nous.fr
emeliapapaige.com
emilydream.com
eminism.org
emptyclosets.com
enginecash.com
enjoy.be
enjoy.com
enormepikken.nl
erection-zone.com
erodating.com
erodating.nl
erohosting.de
erohosting.nl
eropodium.nl
erosportal.net
erostar.ch
erostar.net
erostar.nl
erouniforms.com
ethnicdarlings.com
eu.com
evilchilli.com
excitingillusions.com
excom.be
excretor.com
exgo.com
exoticgold.com
exotiqapparel.com
extape.com
extreme-anne.com
extreme-drawings.com
extremefunnyhumor.com
extremeoldies.com
facebook.com
facegoo.com
facesitting-discipline.com
fairsuchen.com
familydiscipline.com
famousboard.com
fasttimesatnau.com
fasttimesatnau.net
fasttimesatnaupage.com
favorietje.nl
favouritecash.com
fazed.net
fazed.org
fbi.tv
femalecompanions.co.uk
femalecompanions.com
femalecompanions.us
festinhasvip.com
fetatjejer.com
fhgking.com
fhm.co.uk
fhm.com
fhm.com.ph
fhm.fr
fhm.hu
fhm.nl
fhp-inc.com
filespace.com
findadeath.com
finelineinternet.com
finethumbs.com
fissionbyte.com
fitzmulti.com
fkpublications.com
flabber.nl
flavors.me
fling.com
floridawetdreams.com
follandotiasdretreinta.com
foobies.com
foogie.com
footprincess.jp
forcedwitness.com
forumtravesti.com.br
fotostrip.dk
foulmouthshirts.com
fpctraffic.com
fpctraffic2.com
free-be.com
free-comics.com
free-comics.net
freechickz.com
freeexclusivemovies.com
freemoviez.com
freewestpapua.org
friends.com
frieskoppel.nl
frivoli.at
frivoli.net
fuk.ca
fuk.org
fun-games.nl
fundorado.ch
fundorado.com
fundorado.de
funnydownloads.de
funnydownloads.nl
funnyinside.com
funnyville.com
funpic.hu
funpic.nl
funplek.nl
funwithamber.com
g-cash.biz
gabrio.com
garageglamour.com
gcruise.com
gebyarbola.co
gebyarbola.com
gebyarbola.info
gebyarbola.net
geile-rijpevrouwen.nl
geile.nl
geilefotos.be
geilefotos.nl
geileslet.net
geillive.nl
georgiajonesonline.com
get-em-first.com
getcybercash.com
gfy.com
gibbleguts.com
gigacash.com
gilfpot.com
glamourboutique.com
glamourbuckz.com
globill-systems.com
gloriabrame.com
gmw.cn
goaloo.com
goddessolga.com
goedbegin.com
goldmom.org
gonzomovieclub.com
gor.net
gorillamask.net
gotgauge.com
gothic.com
gothic.info
gothic.net
gothichorrorales.com
gratiscontacten.nl
gratisrukken.net
gratisrukken.nl
gratispelletjes.com
gratispelletjes.nl
gratisweb.com
greatincest.com
greelibie.com
greendollars.com
gregboone.com
grinchoo.com
grosnibards.net
gsport.com
gsport.nl
guba.com
gunzblazing.com
halhal.net
hanneblank.com
harry-red.com
hastriki.com
hcube.biz
hdvaccess.com
hdzog.com
heather-summers.net
heaven666.org
heavy.com
heavy.cz
heavy.us
hebus.com
hee.cn
heers.nl
hellsgoddess.co.uk
henta.biz
henta.org
hentime.com
herbalaffiliateprogram.com
herfirstdv.com
herfirstdv.ws
herwoodshed.com
hetbuurmeisje.be
heteschup.com
hightimes.com
himemix.com
hisclub.com
hitbooster.nl
hitslap.com
hivmme.com
hogehakkenspecialist.nl
hohsupport.com
hollywoodsleazy.com
holmaatjes.nl
homemadaactions.com
homepex.net
hostforx.com
howoldisshe.net
hpic.com
hqfisting.com
humorpages.net
hyperfree.com
iambuck.com
iamtrouble.com
iamverified.com
idleriot.com
ikvleenneukdate.nl
ilhadoprazer.com.br
ilhadoprazer.net
imagesystem.it
imcbill.com
imco.nl
imperidefamosas.com
indobolajalan.com
inet-cash.de
info-fiend.com
innulge.com
inrealife.net
insidetwists.com
intensecash.com
interfriendship.de
internationalbikini.com
interview.es
intiem-contact.nl
intimshop.hu
intimshop.info
intimshop.no
intimshop.ru
inzestfamily.com
islandhousekeywest.com
isna.org
iusw.org
ivysummer.com
j-webdesigns.com
jacquieetmichel.fr
jacquieetmichel.info
jacquieetmichel.net
jacquieetmichel.org
jaggle.net
jaggle.nl
janeduvall.com
japanese-action.com
jasperemerald.com
jastusa.com
javbucks.com
javhd.com
javhd.in
javhd.mobi
javhd.pro
javhd.sex
javhd.tokyo
jaxtravstudios.com
jayaliga.com
jayaliga.net
jennaclub.be
jennahaze-blog.com
jennajameson-blog.com
jerkingsresources.com
jessecapelli.com
jessecapelli.net
jessecapelli.ws
jessythumbnails.com
jeunes-cochannes.biz
jeunes-cochannes.eu
jezebelclair.com
jezebelclair.net
jheat.com
jimycanon.com
jimycanon.net
jiriruzek.net
jlist.com
jmbsoft.com
jockphysical.com
joephillips.com
johnandjohn.nl
joinrightnov.com
jonathantart.com
jongetienersletjes.nl
jongy.com
jonrhus.com
joshtucker.com
jousaunagids.nl
juicyclips.com
jurbi.com
justtuds.com
justuseme.com
k-cool.com
kamasutra.at
kamasutra.com
kamasutra.com.mx
kamasutra.es
kamasutra.fr
kamasutra.hu
kamasutra.it
kamasutra.lt
kamasutra.name
kamasutra.nl
kamasutra.nu
kamasutra.org
kamasutra.ru
kamasutrabears.nl
karumz.com
kates-playground.com
kates-playground.info
kavamediaaigroup.com
kawaii.cc
kawaii.com
kcash.biz
kcash.info
keepitnice.com
kelli.net
kenstwistedmind.com
kicken.com
killerbikinis.net
kinghost.com
kladblog.com
koi-de-neuf.fr
kongodongo.com
kontraband.com
koquin.com
kostenlose-toplisten.de
krystalwaters.com
kutlinkjes.nl
kvmediaproductions.com
kwikmed.com
kwikmed.net
kwinkies.com
kyravonkropp.com
l-a-tex.com
lablue.de
lagitane.com
largefriends.com
lastgasp.com
laurentem.nl
lauxanh.org
lauxanh.us
lavalife.com

lay.be
leatherroses.com
lecoq.de
ledix.com
legcash.com
leisuretown.com
lekkerepoesjes.nl
lenalist.com
leschattes.org
leukehumor.nl
leudhost.com
lexa.nl
liberties.be
libertin-online.com
libidopil.com
libidopil.nl
ligasbobet.com
like-em-straight.com
link4all.net
linkdumper.com
linkdumper.org
linkweb.nl
littleyellowdifferent.com
livem8.com
lizvicious.com
loaded.co.uk
lonelycheatingwives.com
longbucks.com
longnail.com
lonnie-waters.com
lotzadollars.com
lui.fr
lynnpaularusell.com
m88.com
m88.com.hk
m88id.com
ma-salope.com
machomedia.hu
madieanne.com
mafiaserver.com
maidenheads.com
male.com
mallcom.com
mallcom.de
mania-manga.com
maniacdiaries.com
maniacdiaries.nl
mantanmuslim.com
marcopolo.be
market4you.com
marquis.de
marriedsecrets.com
match.nl
maxim.com
maxim.it
mccoysguide.com
mea-culpa.com
meet2cheat.de
meetav.com
meetingcafe.nl
megaculos.com
megapic.net
megapic.ru
megaportal.ru
mellonland.com
melonfarmers.co.uk
memberdoorway.com
members.pl
membersupportcenter.com
miabuelaesunaputa.com
mibrujula.com
midget.de
midget.tv
midget.ws
mikesouth.com
milehighclub.com
milkandcookies.com
mintlist.com
missx.com
missx.org
misterwolfe.com
mjtop.com
modemac.com
moedertjes.nl
moistpixels.com
mokkels.be
mokkels.eu
mokkels.nl
monica-mayhem.com
moniquesweb.com
monstermanga.com
moodyz.com
moodyz.tv
moonangel.com
moono.com
morbositas.com
mota.ru
moulinrougeamsterdam.nl
moustiq.com
movie-hut.com
moviefort.com
mozzel.nl
mpegstation.com
mpegstation.nl
mr-bert.nl
mr-s-leather.com
mrcash.nl
mrngood.com
mschristine.com
msmargaretdavis.com
mufftop.com
multiserv.com
murrayandvern.com
muscle.tv
musclelegods.com
mutske.com
mutterundtochter.cc
mutterundtochter.tv
my-vasectomy.com
mygood.biz
mylinea.com
myreadingmanga.info
mysecurewallet.nl
mystique-magazine.com
mywife.cc
mywife.com
mywife.jp
nachtclubvenus.nl
nacktmuschi.de
nad-iksodas.com
nagabaru.com
namiolive.com
napped.com
narcane.com
nattekutje.nl
nattesputkutjes.nl
navashibari.com
nbmpub.com
ndfreehost.com
nearlygood.com
nederlanddating.com
nedsnookie.com
nemo-glamour.com
nenablue.com
nerve.com
nerve.de
netachtien.nl
netmanik.com
netpart.com
neuk-plaatjes.nl
neukafspreekje.nl
neukendoejezo.com
neukendoejezo.net
neukme.nl
neuknu.nl
neukplezier.nl
nevest.net
newgrounds.com
newsdump.com
newsdump.org
nglcc.org
nicerations.com
nichebucks.com
nichelist.com
nichepartners.com
nicheltoplist.com
nicheltopsite.com
nietvoordepoes.nl
nifty.org
nightoforgies.com
nightshiftpatrol.com
nightstation.com
nlgja.org
nlounge.com
nnarchive.org
nopubestgp.com
nordfx.com
nostalgicglamour.com
nowgoal.cc
nowgoal.com
nowgoal.net
nudaten.nl
nulive.nl
oanda.com
oasedresden.de
oast.com
ocioso.com.br
oddsportal.com
odesa.gov.vc
odesa.nl
oilreg.com
oldandcrazy.com
oldandcrazy.net
olderpic.com
oldwishes.com
oliversmoney.com
olympine.com
omegle.com
omimovie.com
on-line-customer-service.com
ondeugendstel.com
ondeugendstel.nl
one-and-lonely.com
only4you.be
onlyfans.com
ontheropes.com
openload.co
opwindend.net
orchiddesigns.com
orientbeach.com
orion-grosshandel.com
orthodykes.org
oska.com
ourtime.com
ouwezoen.nl
over30list.com
ovguide.com
oxcash.com
oy6.net
pacopacomama.com
palace.com
palcomix.com
papi.com
parchis.com
parenclub-aphrodite.nl
parishiltonpage.com
parisplayground.com
parkingcrew.net
partnercash.com
partnercash.de
partnerresource.com
partnerruil.com
partnerruil.net
partnerruil.nl
partnerruil.us
pauljohnballard.com
paycounter.com
payglad.com
paysite-cash.com
paysitereviews.biz
paysitereviews.com
paysitereviews.us
pcrdist.com
people.com.cn
people.nl
pepsaga.com
perl-princess.net
personshavers.com
perufornication.com
picalink.com
pigdog.org
pillsmoney.com
pinstripecash.com
pipax.com
pipedreamproducts.com
pixotna.com
pizorn.com
pk.com
planetelive.com
pluginaccess.com
pluginfeeds.com
plumper.co.uk
plumper.net
plumper.tv
plumper.ws
plumper.xxx
plumperz.com
plumpys.com
plynn.com
podryvacze.pl
poker88.asia
poker88.biz
poker88.in
poker88.org
poker88.plus
poker88.tv
poker88.win
pokermaya.net
pokermaya.org
pokernews.com
pokerstars.bg
pokerstars.com
polenwg.com
polenwg.de
popbytes.com
popmycheri.com
poppen.co.uk
poppen.com
poppen.de
poppen.tv
poppers-shop.com
poppers-shop.nl
porkyhost.com
pornhub.com
posteriorpenetrator.com
postmoderncourtesan.com
potofmoney.com
powernetx.com
prankplace.com
preggo.com
preggo.nl
pregnantlactation.info
pregnantschool.com
pridedreams.com
primerizas.com
princesskali.com
princesslissa.com
pro6.org
prohibidas.com
project-z.com
promocionesweb.com
prophetofdoom.net
protectionparentale.com
pwp-club.de
punishedbrats.com
punternet.com
purewebpower.com
pursedlips.com
quality-control-centre.com
quatero.nl
queen8.com
queenofheartany.com
queensandkings.biz
queerparents.org
queerworld.com
quebec.com
rachel-aziani.net
rachel-aziani.us
rachel-aziani.ws
rachelazianiblog.com
rachelazianiblog.net
rachelkramerbusse1.com
rafik.com
ranchophucko.com
rargb.to
ratemybits.com
ratemyknockers.com
ravinriley.com
rawrods.com
raymondibrahim.com
reddit.com
rencontre-region.com
retecool.com
retecool.nl
retroskank.com
ricefever.com
ringsurf.com
risque.com
risque.com.au
risque.net.au
rollanotherjoint.com
roodharigesletjes.nl
roomservice2000.com
root-top.com
ropeaffairs.com
ropemarks.com
ropemarks.nl
rosie.de
rosie.nl
rotci.com
royalle.com
royalpic.com
rudevirtual.com
runka.com
russische-vrouwen.nl
rvisions.com
sisis1.com
sabayland.nl
sadism.com
sadism.nu
sadism.ws
sadistic.nl
saishuu.com
sandrachang.net
sandyssecrets.com
sankakucomplex.com
sapphicmovies.com
sarofreuve.com
saucy4u.com
saudek.com
sawnpic.com
sbobet.co.uk
sbobet.com
sbobet.com.hk
sbobet.help
sbobet.mobi
sbobet.org
sbobet.sg
sbobet.sx
sbobet.top
sbobet.tv
sbobet.tv
sbobet.us
sbobet.website
sc-2257.com
sc-venus.de
scarletroom.com
scatmovie.com
scatmovie.nl
schih.com
schmacht.org
scoreslive.com
scorpionsoumis.net
sdcmmedia.com
seamountainranch.com
secretleather.com
securitysoft.com
seduction.ca

seduction.com
 seduction.cz
 seduction.net
 seduction.nl
 seiren.com.br
 sepulchritude.com
 serious-coin.com
 seska.com
 sfc.org.uk
 sfsi.org
 shadowlands.com.au
 shavenomore.com
 she-international.com
 shermshack.com
 shesilver.com
 shibaricon.com
 shockingcash.com
 showboat.nl
 siccash.com
 siecus.org
 silentbucks.com
 simonbolz.com
 sirkovski.com
 site90.com
 sizepro.com
 skaters.ws
 skins.be
 skins.org
 skoften.net
 skoop.com
 skybet.com
 slackernetwork.com
 slamhost.com
 slashdong.com
 slashdong.org
 sleeklegs.com
 slonopotam.com
 sm-bomber.com
 small-giant.com
 smaag.com
 smokeroom.com
 smokin-slappers.com
 smokingarchive.com
 smokingpole.com
 smokingworld.net
 snaz75.com
 sneak-a-peak.com
 soccernews.com
 soccerpunter.com
 soccervista.com
 soft-core.us
 softlinkers.org
 soj.org
 solarcash.com
 solidoak.com
 sologratias.biz
 sologratias.it
 sonja-adams.com
 sophiarossi-blog.com
 southcn.com
 southernbrooke.com
 spacash.com
 spacebling.com
 spamfreeforums.com
 spannertrust.org
 spicyblogs.com
 spiderthumbs.com
 spiece.com
 spiritworks-art.com
 splashsmoke.com
 splut.com
 springbreakbeach.com
 spyaddicts.com
 spydrive.net
 squirrel-nuts.com
 sqx.nl
 starfool.com
 starletz.com
 stars-tgp.com
 starsx.fr
 start4all.com
 startplezier.nl
 startspot.nl
 startstek.nl
 steadybucks.com
 stephm.com
 stevedietgoedde.com
 strandslet.nl
 streakerama.com
 stripgamecentral.com
 stripperweb.com
 strontslet.nl
 studio-esme.com
 studio-mystique.nl
 studio66tv.com
 stuffchannel.com
 suarapapua.com
 submissive.ws
 submitool.com
 sultryserver.com
 summerbunnies.com
 sundiary.com
 sunnydollars.com
 sunnydollars.net
 sunnypotd.com
 sunnytgp.com
 superhq.net
 supertgps.com
 supportservice24h.com
 sureflix.com
 surfplaza.be
 susanstgp.com
 suzannevinters.com
 sybian.cn
 sybian.com
 symtoys.com
 takeninhand.com
 talesofdeath.com
 tanlinesclub.com
 tantra.at
 tantra.co.nz
 tantra.com
 tantra.org
 tantraworks.com
 tantrumtrainers.com
 tapism.com
 tastelikepizza.com
 tatet.net
 tchatche.us
 tech-chick.com
 teddibarrett.com
 teddibarrett.net
 temptation.be
 temptation.de
 temptation.nu
 tendrebulle.fr
 terri-summers.com
 terri-summers.net
 terrisummers.com
 terrisummers.ws
 terryvision.net
 tes.org
 tetasgrandes.info
 tetasgrandes.org
 tetten.be
 tgphaven.com
 tgplink.com
 tgpteam.com
 tgptraffic.biz
 tgptraffic.info
 tgptraffic.net
 the-boss.nl
 the-iron-gate.com
 thebaddest.com
 thedatefinders.com
 thedatefinders.net
 thedatefinders.org
 theexiles.org
 thehouseofx.com
 thelondonanners.com
 thenichetraffic.com
 thenuproject.com
 thepillowbook.com
 thepinupfiles.com
 thepiratebay.org
 thepiratebay.se
 thequran.com
 thereligionofpeace.com
 thesultan.com
 thesultan.nl
 thevalkyrie.com
 theybf.com
 thongz.net
 threshold.org
 thrills.com
 thscore.cc
 thuisontvangen.nl
 thumblogger.com
 ticklingforum.com
 tienermokkels.be
 tienermokkels.nl
 tiffanypreston.com
 tiffanypreston.info
 tnacash.com
 tombianchi.com
 tombianchi.org
 tommyedwards.com
 tomspicpost.com
 tonysmovies.com
 tonyvard.com
 tooole.com
 top30.com.br
 topbucks.com
 topdatingsites.info
 topdatingsites.nl
 tophumor.nl
 toplingerie.fr
 toplingerie.ws
 toplist24.de
 toplista.pl
 toplistdirectory.com
 topsiteguide.com
 topsites.com.br
 topsites.it
 torsky.org
 tosurfy.com
 tosurfy.nl
 totally-briana.com
 totallycrap.com
 touchmeup.com
 tracilords.com
 traffic-avenue.com
 traffic-bomb.net
 trafficholder.com
 tramparamsimspons.com
 trqbdqahr.com
 truetere.com
 truewiveswhocheat.com
 trulyz.com
 truthkings.com
 tsclips.com
 tsday.com
 tuk.be
 tuk.nl
 twistedblogs.com
 twistedcash.com
 twistedfiles.com
 twistedmonk.com
 twistyscash.com
 uk.com
 uk.net
 ultraflix.com
 ultrapegs.com
 ultrasparky.org
 ultrastart.nl
 uncovered.net
 uncovered.ws
 uncutdvds.co.nz
 uncutdvds.com
 uncutdvds.com.au
 uni-cash.com
 uniform-thumbs.com
 untrue.com
 urbx.com
 uselessjunk.com
 uselessjunk.net
 userrules.com
 uthervers.com
 vanjas-world.com
 vcomm.net
 vd.com
 velvetmag.com
 venus-berlin.com
 venus-paireclub.de
 venuscash.com
 victorspage.com
 vigorelle.com
 vigrx.com
 vimeo.com
 violent-comix.com
 vipluxuria.com
 virtuguy2.com
 vkini.com
 volume-pills.com
 voomed.com
 voor-iedereen.nl
 vp7.com
 vpscash.be
 vpscash.com
 vpscash.nl
 vs.com
 w3open.com
 w88.com
 waarzo.nl
 wasteland.com
 wasteland.nl
 watcmvfg.com
 waytoomany.com
 wdating.com
 web-date.co.uk
 web-log.nl
 webair.com
 webcashmaker.com
 webpark.pl
 webpark.ru
 webpark.sk
 webtjigertje.nl
 webxfrance.com
 webxfrance.org
 wet-t-shirts.com
 weya.com
 whackin.com
 whitelinefirm.nl
 wildcomics.com
 wildhookups.com
 wildseries.com
 williamhiggins.com
 williamhill.com
 wilprive.nl
 winnifred.nl
 witchtorture.com
 wjwebdesigns.com
 wolume.com
 womaninprison.com
 wordoyster.com
 world-collections.com
 wrongsideoftown.com
 wsacp.org
 wutdd.com
 wyldebsites.com
 x2x.de
 x2x.nl
 xahead.com
 xamo.com
 xamo.net
 xaviersite.com
 xdamez.com
 xeromag.com
 xfreehosting.com
 xmovie.com
 xmovie.it
 xoriental.com
 xrevenge.org
 xstarsworld.com
 xyzcomics.com
 yakcash.com
 yankmycrank.com
 yankmycrank.ws
 yee-e.com
 yetisblog.com
 yhao.com
 ynot.com
 yonkis.com
 yonkis.net
 yourchoice.co.uk
 yourchoice.nl
 yourdreamdate.nl
 yumi.nl
 yummyly.com
 zapto.org
 zetacash.com
 zlata.de
 zonacaliente.es
 zone-archive.com
 zone-coquine.com
 zonemall.com
 zoaccess.com
 zootoday.com
 zwijnerij.nl

A.3 Iran

12bet.com
ix.com
2016bestnine.com
2forms.gov.il
4chan.org
4everproxy.com
500px.com
500px.org
8tracks.com
aasoo.org
adobe.com
adultfriendfinder.com
aebn.net
agacystore.com
alarabiya.net
alertpay.com
alkasir.com
alphacoders.com
amadnews.com
amadnews.org
amazonaws.com
amiclubwear.com
amontazeri.com
aniscartujo.com
annsummers.com
annunci69.it
anonymizer.com
anonymous-proxy-servers.net
anonymoussurf.us
apparelnews.net
aqeedeh.com
archive.gov.il
archive.is
artstation.com
ashemaltube.com
ashleymadison.com
asiasociety.org
ask.fm
asl19.org
asos.com
asranarshism.com
astrill.com
aswat.com
audioboom.com
avaaz.org
avn.com
azadegi.com
azadiandishsh.com
babes.com
babyblog.ru
bahai.fi
bahai.nl
bahai.org.au
bahai.us
balatarin.com
bandcamp.com
barebackrt.com
battlecreekenquir.com
bayatzanzani.net
bbc.co.uk
bbc.com
bbcpersian.com
bdsmcafe.com
beatport.com
beeg.club
beeg.com
behance.net
bellabellaboutique.com
bestvpn.com
betternet.co
bgboyfriendtv.com
bia2.com
bible.com
biblegateway.com
bic.org
bigcartel.com
bikini.com
bikini.com.gr
billabong.com
blacked.com
blackmilkclothing.com
blogger.com
bloglovin.com
blogspot.bg
blogspot.co.at
blogspot.co.id
blogspot.com
blogspot.com.au
blogspot.com.ng
bonyadhomayoun.com
boyfriendtv.com
bradbare.com
brightside.me
browse007.com
btl.gov.il
buffered.com
buysub.com
buzzfeed.com
bypassy.com
caa.gov.il
cam4.com
cam4bucks.com
causes.com
cbs.gov.il
chatrandom.com
chaturbate.com
chaturbate.global
checkedproxylists.com
christiandatingforfree.com
cia.gov
citizenlab.org
clarionproject.org
codyapp.com
cognitiforms.com
connectpal.com
contentabc.com
convio.net
corbinfisher.com
couchsurfing.com
court.gov.il
criterion.com
cyberhostvpn.com
daftarche.com
dailymotion.com
darkalley.com
deepikaghai.com
definebabe.com
derafsh-kaviyani.com
discogs.com
dnamagazine.com.au
donmai.us
douban.com
dramafever.com
dredown.com
dressin.com
dropbox.com
duckduckgo.com
e621.net
ekhartyoga.com
eliteproxyswitcher.com
ello.co
emerproxy.xyz
enghelabe-eslami.com
epicbrowser.com
eroshare.com
espacelibido.com
eurolive.com
expatdatingfrance.com
export.gov.il
eyeem.com
facebook.com
fanart.tv
fandor.com
farda.us
favoritewords.com
fb.com
feedburner.com
fetishcon.com
fetishpapa.com
fetlife.com
fidh.org
figurerealm.com
filmaffinity.com
filmon.com
filmon.tv
fitnessblender.com
flickr.com
flipboard.com
flyproxy.com
forbiddenplanet.com
forms.gov.il
fotostrana.ru
fozoolemahaleh.com
free-proxy-list.net
freedating.co.uk
freedomessenger.com
freemoviez.biz
freepeople.com
freevpn.ninja
freevpn.org
freevpn.pw
french-twinks.com
fstoppers.com
fuskator.com
galoremag.com
garow.me
gatherproxy.com
gay.com
gaydemon.com
gaypornlovers.com
gaysexpositionsguide.com
gb.com
geeksaresexy.net
getlantern.org
gfyca.com
gggt.com
gilgamishaan-books.org
giphy.com
glamour.it
glamour.nl
glwiz.com
glymp.com
gmfa.org.uk
goldenfrog.com
google.com
googleusercontent.com
googlevideo.com
gooya.com
gopetition.com
gotinder.com
gotquestions.org
gr.com
groupsexinthecity.com
gstatic.com
gunroad.com
gunzblazing.com
gunzblazingpromo.com
halazonmag.com
hambastegimeli.com
handami.com
happysocks.com
hbonov.com
hearstmags.com
helixstudios.com
helixstudios.net
hide-ip.us
hideip.co
hideipproxy.com
hideipvpn.com
hidemyass.com
hidemyip.info
hiderealip.net
hiveminer.com
hola.org
hornydesigns.com
hotmiamistyles.com
hotmovies.com
hotmovies.xxx
hotspotshield.com
howcool.com
hra-news.org
hrv.org
hsselite.com
humanism.org.uk
humanrightsiniran.com
humanrightsiniran.org
hunkemoller.com
hunkemoller.de
hwcdn.net
icfj.org
ifttt.com
imagecomics.com
imo.im
ims.gov.il
incloak.com
indexxx.com
info.al
info.gov.il
ink361.com
imagine.com
instagr.in
instagy.com
internetcloak.com
internetproxy.eu
ipcload.us
ipmask.us
ipower.com
iran-livetv.com
iranapi.net
iranhumanrights.org
iraniancanada.ca
iranianuk.com
iranrights.org
iransong.com
iranwire.com
irsfe.in
irsfe.ir
isohunt.to
israel-mfa.gov.il
issuu.com
itv.com
jackpot247.com
jadi.net
jameda.com
jcrew.com
jensiat.io
jlist.com
joopea.com
joopea.news
juicyads.com
justjared.com
justpaste.it
juxtapoz.com
jw.org
kalame.org
kamanagir.net
katch.me
keep.com
kendallj.com
khabareg.com
khabarnet.info
khodnevis.org
khoondi.com
kiiroo.com
kimkardashianwest.com
kink.com
kk.no
knesset.gov.il
kobo.com
kobobooks.com
lambiek.net
landmilkhoney.com
latinboyz.com
lds.org
lelo.com
lenceria-sexy.net
lesgrandsducs.com
likekhor.com
lintas.me
littlemonsters.com
livechatinc.com
livejasmin.com
livingsocial.com
locari.jp
lovehoney.com
lovehoney.eu
lubed.com
lucasentertainment.com
lullabellz.com
lulus.com
lulus.tw
lunss.com
magshop.com.au
mail.gov.il
mail.ru
majzooban.org
manjam.co
manjam.com
manjam.eu
manoto.news
manoto1.com
manototv.com
manyvids.com
marde-rooz.com
mardomak.org
mardomreport.net
marxists.org
maryam-rajavi.com
masculinities101.com
maskip.info
massageenvy.com
meccabingo.com
media-imdb.com
medium.com
mejalehhafeh.com
melimazhabi.com
melonbooks.co.jp
melonbooks.com
mfa.gov.il
mixcloud.com
mixpanel.com
mixx.com
mod.gov.il
models.com
mohabatnews.com
moital.gov.il
mormonchannel.org
mossad.gov.il
most.gov.il
mot.gov.il
movieguide.org
mr-s-leather.com
mubi.com
mulpix.com
my-proxy.com
myanimelist.net
mycomicshop.com
myipproxylist.com
myspace.com
naakojaa.com
namehbeanha.com
nebraska.gov
net-a-porter.com

netflix.com
 nevproxylist.net
 nevproxysites.com
 nextmagazine.com
 noisetrade.com
 notforprint.co
 nowness.com
 okcupid.com
 omiddana.net
 onlineproxy.eu
 onlinevideoconverter.com
 oovoo.com
 openload.co
 openvpn.net
 openvpn.org
 opera-mini.com
 opera-mini.net
 paniraniist.org
 panyuxin.com
 papodehomen.com.br
 paskoocheh.com
 passinternet.com
 patreon.com
 peacewithgod.net
 perfectweddingguide.com
 periscope.tv
 persecution.com
 persianhub.org
 pexels.com
 photofunia.com
 php-proxy.com
 piba.gov.il
 picclick.co.uk
 pinalove.com
 pining.com
 pinterest.com
 pinterest.de
 pinterest.ie
 pinupgirlclothing.com
 piratebayproxy.be
 piratebayproxy.co
 piratebayproxy.tf
 pixabay.com
 pixiv.net
 planetromeo.com
 playboy.com
 pleasuresments.com
 plurk.com
 pmo.gov.il
 pof.com
 pokemon.com
 pond5.com
 popjustice.com
 popsugar.com
 porn.com
 porn.es
 porndeals.com
 pornhub.com
 pornotube.com
 pornoxo.com
 poshmark.com
 posting.org
 prettyslittleting.com
 previewsworld.com
 privateinternetaccess.com
 promodj.com
 prosolutionpills.com
 proxifier.com
 proxy-free.de
 proxy-site.net
 proxy-site.us
 proxy.cm
 proxy.lu
 proxy.nu
 proxy.org
 proxy.porn
 proxy.rocks
 proxy.sx
 proxy62.com
 proxybrowsing.com
 proxylists.me
 proxynova.com
 proxyterminal.com
 proxyweb.net
 proxywiky.com
 proxz.com
 psiphon.ca
 psiphon.me
 psiphon3.com
 psiphon3.net
 publicdesire.co.uk
 publicdesire.com
 pulpcoverters.com
 pyknet.net
 qombol.com
 qpic.cn
 qq.com
 queerty.com
 quora.com
 rabb.it
 radiofarda.com
 radiojavan.biz
 radiojavan.com
 radiojavan.me
 radiojavan.tv
 radiojavan.ws
 radiozamanah.com
 radis.org
 ragingstallion.com
 rapfa.com
 realdoll.com
 rebelmouse.com
 recon.com
 reddit.com
 redtube.com
 reelrundown.com
 religinfoserv.gov.il
 renderosity.com
 rentmen.ae
 rentmen.at
 rentmen.ch
 rentmen.cl
 rentmen.cn
 rentmen.co
 rentmen.com
 rentmen.com.au
 rentmen.com.br
 rentmen.cz
 rentmen.dk
 rentmen.es
 rentmen.eu
 rentmen.fr
 rentmen.in
 rentmen.it
 rentmen.jp
 rentmen.kr
 rentmen.my
 rentmen.pe
 rentmen.ph
 rentmen.se
 rentmen.sg
 reverbnation.com
 revolutiondance.com
 rfi.fr
 riffsy.com
 risheha.com
 rocknrollbride.com
 roshanfekr.org
 rottentomatoes.com
 rowzane.com
 rutube.ru
 saanei.org
 sabzlink.com
 sahamnews.org
 savedeo.com
 scribd.com
 secondlife.com
 secrethallway.com
 securityinbox.org
 seekingarrangement.com
 sexfactor.com
 sexlikereal.com
 shahrvand.com
 shahvani.com
 shockblast.net
 shopbop.com
 shopcade.com
 shopwiki.com
 showpo.com
 shutterstock.com
 sjmaylee.com
 skinymint.com
 slaveregister.com
 smartandsexy.com
 smashwords.com
 snapchat.com
 snapwidget.com
 socks-proxy.net
 socksproxychecker.com
 solarmovie.click
 solarmovie.fm
 solarmovie.is
 solarmovie.ph
 solarmovie.site
 solarmovie.st
 soulsingles.com
 soundcloud.com
 southbeachswimsuits.com
 spark.com
 spencersonline.com
 rabb.it
 spotify.com
 sslproxies.org
 sslsecureproxy.com
 staticflickr.com
 stepbible.org
 stockroom.com
 storenvy.com
 strikingly.com
 stripchat.com
 strongvpn.com
 stylewe.com
 subscene.com
 sunni-news.net
 sunporno.com
 svpply.com
 swimmingtiger.com
 swimsport.com
 tableamag.com
 tagboard.com
 talktome.com
 tarnama.org
 tarnama.us
 tavaana.org
 tehila.gov.il
 telegra.ph
 tenor.co
 tenor.com
 terijon.com
 thebestvideocontentever.com
 thechive.com
 thehunt.com
 theknot.com
 themoviedb.org
 thenude.eu
 theoutnet.com
 thepiratebay.al
 thepiratebay.bid
 thepiratebay.cd
 thepiratebay.co.in
 thepiratebay.expert
 thepiratebay.gd
 thepiratebay.gg
 thepiratebay.id
 thepiratebay.lv
 thepiratebay.me.uk
 thepiratebay.mk
 thepiratebay.org
 thepiratebay.plus
 thepiratebay.rs
 thepiratebay.se
 thepiratebay.vg
 thesims.com
 theteenboy.com
 theuntappedsource.com
 thewrap.com
 thirdlove.com
 tictail.com
 tigerproxy.net
 timeout.com
 timtales.com
 tiny4k.com
 tjournal.ru
 tnaboard.com
 todorelatos.com
 tofo.me
 tolovehonorandvacuum.com
 tondar.ca
 tondar.org
 toonistan.com
 toosheh.org
 toovia.com
 top-proxies.co.uk
 toptopic.com
 torguard.net
 torlock.com
 torproject.org
 trakt.tv
 treasureislandmedia.com
 tribecafilm.com
 tribunezamanah.com
 tumblr.com
 tunnelbear.com
 twimg.com
 twitter.com
 ucweb.com
 uk.net
 ukipvpn.com
 ukswimwear.com
 ultrasurf.us
 unblocked.bid
 unblocked.la
 unblocked.lol
 unblocked.onl
 unblocked.pw
 unblocked.red
 unblocked.rocks
 unblocked.srl
 unblocked.video
 unblocker.me
 unblocker.us
 unblocker.yt
 unblockmyweb.com
 unblockthesites.com
 unblockvideos.com
 unlockyoutube.co
 united4iran.org
 uplust.com
 uproxy.org
 us-proxy.org
 usmagazine.com
 valuemags.com
 veoh.com
 vice.com
 victoriabeckham.com
 victoriasssecret.com
 vid.me
 vietnameselove.com
 viki.com
 vine.co
 vingle.net
 vixen.com
 vk.com
 vk.me
 voanews.com
 voat.co
 vogue.nl
 vpnbook.com
 vpnritic.com
 vpmme.me
 wanelo.com
 wapmia.com
 watsappad.com
 we-change.org
 websta.me
 wechat.com
 weheartit.com
 whisper.sh
 whispersystems.org
 whoer.net
 whyweprotest.net
 wikiislam.net
 wikinews.org
 wish.com
 womensecret.com
 wordpress.com
 wp.com
 wsop.com
 www.gov.il
 xhamster.com
 xhamster.one
 xhamsterlive.com
 xnxx.com
 xnxx.video
 xroxy.com
 xvideos.com
 xvideos.works
 yandex.com
 yekray.com
 youfreeproxytube.com
 younow.com
 youporn.com
 youtube-nocookie.com
 youtube.com
 yting.com
 yuki.la
 zahedi.com
 zalando.fr
 zamaaneh.com
 zamanalwsl.net
 zandiq.com
 zeitoons.com
 zendesk.com
 zenmate.com
 zoxy.net

A.4 Turkey

Adult domains (39,212) not shown

100ppi.com
itubehd.com
24video.net
360kad.com
3d-comics-tgp.com
3dbdmdungeon.com
3design3.com
3dtgp.com
3dchat.com
3wasonnet.com
3xnews.com
4flirt.nl
4shared.com
69.com
6monkey.com
777.com
888.com
888.it
888casino.com
888casino.it
888games.com
888poker.com
888poker.es
8ch.net
abspritz-bilder.com
actualcash.com
adam4adamlive.com
adinfinity.com.au
adorableaudrey.com
adsvinging.co.uk
aepartnership.com
affairalert.com
africantube.net
agents69.com
aka.ms
aladd.net
alexz-traffic.com
alluc.com
alluc.ee
alluc.to
amateurtaboo.net
amber-michaels.org
anfenglish.com
ankurdi.com
anfturkce.net
angel-live.com
aniloscash.com
animecontent.com
anonymouse.org
anycash.com
aphrodisiacsofnature.com
artmunkgames.com
artofzoo.com
ashleydoll.com
ashleymadison.com
ashleymadison.com
asiancam.eu
asiantubevideo.com
atkcash.com
audreybitoni.com
audreybitonitube.com
avmariaozawa.com
azianigold.com
backpage.com
badoink.com
balmtube.com
bankofbeijing.com.cn
bazoocam.org
bdsmbboard.org
bdsmsculture.com
bdsmkino.com
bearfilms.com
behindthetowel.com
bendibao.com
bestofblowjobs.com
bestoflegsandfeet.com
bet365.com
bethemask.com
betsafe.com
betterfap.com
biancabeauchamp.com
bigbootysingles.com
bigtftpatriol.com
bizarre100.com
black-women-tube.com
blackcuke.com
blackmama.com
blackplushwhite.com
blip.tv
blogspot.ro
blurayclips.com
bonetoob.com
bonkwire.com
bootyfix.com
braincash.com
bravedigger.net
brianaonline.com
brookbradford.com
brookelina.com
bucksmatrix.com
bullslinks.com
burbujasdeseo.com
c0930.com
caiyunapp.com
camelstyle.net
camvide.org
candydollchan.net
caribbeancom.com
cashdorado.de
cashmaniacs.com
catchycash.com
ccb.com
cduniverse.com
chanchan.com
cheeseepin.net
chemnet.com
cherrykiss.org
cherryspot.com
chocomelons.com
chubbytubby.com
clipsblacktubes.com
clubdomcash.com
cnkang.com
coltstudiostore.com
cmdotgame.com
comicsband.com
comicsxd.com
coolhandpoker.com
covers.com
craziescash.com
creative3dcash.com
crissymoran.net
crookedhalos.com
crunchydollars.com
crystalindian.com
cyonix.to
czechcasting.com
dablacktube.com
dailyniner.com
dangniao.com
daredormmobile.com
dawfilms.com
ddfcash.com
dirtylilly.com
dirtystain.com
dlxer.com
dmm.co.jp
dofantasy.com
dofiga.net
doing.my
domai-art.com
donita-dunes.net
doodmovies.com
drivenbyboredom.com
dudetubeonline.com
egotastic.com
elitedollars.com
elles-se-mettent-nues-pour-nous.fr
emilybound.com
enginecash.com
english-stockings.com
enjoybucks.com
epidemz.net
erotilink.com
eurolive.com
europacasino.com
evaangelinaonline.com
evennode.com
evilangelvideo.com
excluzive.net
exoticgold.com
extreme-forum.net
facebook.com
famedollars.com
famousboard.com
fappic.com
favebmasters.com
fckya.com
femdomstory.org
femlatex.com
flavors.me
fling.com
forcedmen.com
forumtravesti.com.br
fotoblow.com
fotolaski.ru
fotos-videos-gratis.com
fpctrffic2.com
freakjunction.com
freeblacktubes.com
freejabcomix.com
freevideobase.com
freewebs.com
frontarmy.com
fundorado.de
funny-base.com
funny-games.biz
funpic.hu
futanarisplash.com
g-cash.biz
geocities.com
gfy.com
gkiss.com
gmw.cn
gndbank.com
godsofmen.com
goldmom.org
goldtalen.com
gorillamask.net
greatincest.com
grindr.com
growlapp.com
gunblazing.com
guyswithiphones.com
h0930.com
hairyardraw.com
hajeez.com
hanime.tv
happygos.com
haremovies.com
hdmovieclub.com
heartbreakersforum.com
heaven666.org
hentine.com
himemix.com
hithomemovies.com
hmsites.com
honestvids.com
hostgator.com
hgfsting.com
hunterscash.com
hussytube.com
ichan.org
idealgasm.com
idolbucks.com
im9.eu
imgchili.net
imperiodefamosas.com
indecisivecaptions.com
industrytweeet.com
informe.com
instagram.com
interrspace.com
intimshop.ru
ip138.com
itc.nu
itr2010.org
jackpotcity.com
jackpotcitycasino.com
jacquieetmichelstour.com
jamestraffic.com
javgbucks.com
jvideo.com
jerkhour.com
jinha.com.tr
jlist.com
joinforjoy.com
juicyads.com
jurbi.com
jynxmaze.info
kaixin001.com
kaotic.com
keepitnice.com
kennyspennies.com
laftube.com
latotona.com
lauxanh.us
laydapipe.com
lenovo.com
lestaigames.com
lightspeedcash.com
link4all.net
lirio.us
littlesiangfs.com
livescore.com
lollipoptube.com
lookimonline.com
lotzadollars.com
lovelyirene.com
mafia-linkz.to
making-love.tv
malerevenue.com
manbukake.com
manicamoney.com
manjam.com
manplay.com
mashable.com
maverickmen.com
meat-rack.com
meetav.com
men.com
mendyou.com
milano.com
mistyanderson.com
mjtop.com
mom-next-door.com
moonangel.com
morexxx.com
moustiq.com
moviedollars.com
mr-s-leather.com
mrsnake.com
menbc.com
mundomais.com.br
mykokam.com
myreadingmanga.info
myreferer.com
mytrannies.com
myvidster.com
neatmovies.com
nemo-glamour.com
netflix.com
newgrounds.com
niceratiots.com
nifty.org
nightlife141.com
nightmovesonline.com
ninevids.com
nscash.com
nylons100.com
offbeatr.com
oh-myygoth.com
okokoras.gr
olympine.com
omegle.com
omgblog.com
oognip.com
ovguide.com
ownvideo.com
pantyhosenextdoor.com
pastebin.com
pc899.com
people.com.cn
pepsaga.com
peterfever.com
pharmnet.com.cn
phonemates.com
plugerr.com
pluginfeeds.com
pocitadlo.sk
poppen.de
poppers-shop.com
profitsdeluxe.com
pubacash.com
publicfeet.com
puffyfash.com
pumasvetetube.com
qqtube.com
rargb.com
rarg.to
realdoll.com.cn
realdolluk.com
realitydudes.com
realllyusefulcash.com
rebootcash.com
redtube.su
rent.men
rentboy.com
rentmen.com
retromoviestube.com
rewardsaffiliates.com
riskified.com
rktv.com
roksa.pl
ronharris.com
ronharrisgalleries.com
roughanalvideos.com
royal-cash.com
rudevirtual.com
rugerbugger.com
samrosee.net
sankakucomplex.com
savitabhahimovie.com
sensualarousalblog.com
sensualarousalgalleries.com
serious-films.com

seska.com
shockingcash.com
showcasecash.com
shuhai.com
simplydevon.com
sitetag.us
skins.be
sm-bomber.com
smaq.com
smarttubepro.com
smotrix.com
somaslots.com
spacash.com
spankingden.com
spiderdollars.com
studmuffinblog.com
suaveswing.com.br
sugarmpeg.com
sunmaker.com
supermen.com
sweetandraw.com
symboltube.com
takebucks.com
tanlinesclub.com
tastelikepizza.com
taylortrue.com
testrust.com
texnet.com.cn
tgpdreams.com

tgptraffic.biz
the-clitoris.com
theanalpics.com
thebiggestforums.com
thefloatingworld.com
thenupject.com
thepiratebay.se
thesword.com
thickcash.com
thugvideos.com
thumblogger.com
tiffanybrookeslove.com
titancasino.com
titanpoker.com
titsbouncers.com
titusblog.com
toastdontheinside.com
todayszaman.com
toocle.com
topbucksmobile.com
topsiteguide.com
toptobottom.com
torjackan.info
torskys.org
toydemon.com
trishauptown.com
tsdreams.com
tsflix.com
tube43.com

tube77.com
tubealliance.com
tubeinvite.com
tubemales.com
tubemike.com
tubemm.com
tubeowl.com
tubeufo.com
tubexe.com
twistycash.com
ucgalleries.com
ultramegabit.com
unkrossed.com
untrue.com
uplay-istrip.com
uploadhouse.com
uploadia.com
vamateur.com
venus-berlin.com
vid.blog.br
vid.me
videoclipa.com
videolovesyou.com
videomega.tv
videoslots.com
virgins-tgp.com
voyeur-house.tv
vp7.com
vpscash.com

waav.tv
wallpaperdome.com
warttube.com
wepayyoulongtime.com
wikipedia.org
wodrun.com
worldwidetopsites.com
wudage.com
www.edu.cn
wyldesites.com
xdating.com
xfliq.com
xfreehosting.com
xoriental.com
xossip.com
xtitstube.com
xyzcomics.com
yakcash.com
yaol-games.com
ydss.cn
ymkikaku.com
yonkis.com
yorozukoubou.jp
youku.com
young-legal.com
zakulisi.cz
zonacaliente.es
zone-archive.com

FIN