

What a Tangled Web We Weave: Understanding the Interconnectedness of the Third Party Cookie Ecosystem

Xuehui Hu
King's College London
London, UK
xuehui.hu@kcl.ac.uk

Nishanth Sastry
King's College London, UK
University of Surrey, UK
nishanth.sastry@kcl.ac.uk

ABSTRACT

When users browse to a so-called “First Party” website, other third parties are able to place cookies on the users’ browsers. Although this practice can enable some important use cases, in practice, these third party cookies also allow trackers to identify that a user has visited two or more first parties which both share the second party. This simple feature has been used to bootstrap an extensive tracking ecosystem that can severely compromise user privacy.

In this paper, we develop a metric called “tangle factor” that measures how a set of first party websites may be interconnected or tangled with each other based on the common third parties used. Our insight is that the interconnectedness can be calculated as the chromatic number of a graph where the first party sites are the nodes, and edges are induced based on shared third parties.

We use this technique to measure the interconnectedness of the browsing patterns of over 100 users in 25 different countries, through a Chrome browser plugin which we have deployed. The users of our plugin consist of a small carefully selected set of 15 test users in UK and China, and 1000+ in-the-wild users, of whom 124 have shared data with us. We show that different countries have different levels of interconnectedness, for example China has a lower tangle factor than the UK. We also show that when visiting the same sets of websites from China, the tangle factor is smaller, due to blocking of major operators like Google and Facebook.

We show that selectively removing the largest trackers is a very effective way of decreasing the interconnectedness of third party websites. We then consider blocking practices employed by privacy-conscious users (such as ad blockers) as well as those enabled by default by Chrome and Firefox, and compare their effectiveness using the tangle factor metric we have defined. Our results help quantify for the first time the extent to which one ad blocker is more effective than others, and how Firefox defaults also greatly help decrease third party tracking compared to Chrome.

CCS CONCEPTS

• **Security and privacy** → **Web application security**; *Social network security and privacy*.

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

WebSci '20, July 6–10, 2020, Southampton, United Kingdom

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7989-2/20/07...\$15.00

<https://doi.org/10.1145/3394231.3397897>

KEYWORDS

Container Identity, Third Party Cookie

ACM Reference Format:

Xuehui Hu and Nishanth Sastry. 2020. What a Tangled Web We Weave: Understanding the Interconnectedness of the Third Party Cookie Ecosystem. In *12th ACM Conference on Web Science (WebSci '20), July 6–10, 2020, Southampton, United Kingdom*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3394231.3397897>

1 INTRODUCTION

It is common knowledge that cookies are placed in a user’s browser when they visit a website. The websites that the users visit are called first party websites (FPs). Although some of these cookies are placed by FP domains, many are placed by third party affiliates (TPs) of the first party sites, for reasons such as advertising, or analytics. Examples include advertising networks such as `adnxs.com`, `amazon-adsystem.com` and `doubleclick.net`, analytics platforms such as `google-analytics.com`, or social media trackers such as Facebook and Twitter.

Such TPs have proved to be an extremely powerful method for aggregating users’ browser histories – for example, a TP that appears on both `https://www.bbc.com`, and `https://www.nytimes.com` is able to infer that a user visited both sites, and therefore can infer that the user might be someone who is interested in news and current affairs. An important motivation for such detailed profiling of users is that it can lead to more “targeted” advertisements, which are much more profitable. For instance, recent research shows that right-leaning hyperpartisan websites in the US use more sophisticated and more intensive tracking techniques and are therefore able to command higher prices than left-leaning websites [3].

As advertisers’ demand for detailed profiles continues to grow, so does the sophistication of third party tracking technologies. Users and many browsers have responded with new technologies to safeguard user privacy. Indeed, this has led to an “arms race”, with users installing ad blockers such as Adblock Plus[14], uBlock Origin[18] and Ghostery[9], and certain websites (e.g., `memeburn.com`, `englishforum.ch`) responding with anti ad-blocking technology[31, 37, 46] that refuse to deliver content unless users unblock ad blockers whilst visiting their sites.

An important addition to the arsenal of technologies for preventing tracking of users is the notion of *multi-account containers*, often referred to simply as “containers”. Pioneered by Firefox [33], containers are a way to separate different sets of cookies from each other. Containers were initially intended for providing “contextual identities”¹, i.e., to create different user identities depending on

¹https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/Work_with_contextual_identities

the context of operation. For example, containers allow a user to cleanly separate their work and personal identities on the same browser. Different containers can also be used to login to the same sites with multiple user ids (for example, users having several email accounts from the same provider can be simultaneously logged in to each of them in different containers). In this sense, containers are similar to other browser extensions such as Multifox [28] or CookieSwap [13] (or the similar Swap my Cookies extension on Chrome [15]).

Firefox containers essentially create a different user profile within each container, providing a different database of cookies and storage for each². Thus, each container identity is kept separate from the others, and information such as third party cookies are not shared across containers. Firefox suggests [33] that this can be used to also achieve additional privacy, for instance by placing a user's shopping websites into a separate container from financial websites such as banks and credit cards.

Although containers offer a clean way to separate third (and first) parties from each other by creating different cookie stores, they are only effective if interfering parties are manually placed in separate containers. Recently, Firefox has come up with add-ons which *automatically, or by default*, provide protection for one of the most prevalent trackers – Facebook. This add-on creates a specialised container [36] for Facebook. Once installed, it opens Facebook in its own container, and the user is logged out of Facebook in other containers, thus automatically preventing Facebook Pixel [35] and other tracking by Facebook of users who are logged in.

While a solution for Facebook's tracking is indeed important as it is widely used on many websites, it is not sufficient, as it does not offer protection against other third parties. Furthermore, Facebook is blocked by the Government in countries such as China, and thus is not a major third party in that country [20]. Thus, solutions are needed that work for other important third parties, as well as for other country and cultural contexts.

Generalising from the Facebook example, any third party can be prevented from learning the (partial) browsing histories of users if two first parties sharing the same third party are placed in different containers. Fig. 1 illustrates this with an example. The websites in green and red share one or more third parties (e.g., Facebook, Google DoubleClick etc.), and therefore need to be placed in separate containers; otherwise the common third parties (e.g., Google's DoubleClick) will be able to infer that the same user visited both the green and red websites. However, the blue site does not share any third parties with either the green or red website and therefore can be placed in the same container with either of those two websites.

This paper asks the simple question: if the above policy is applied uniformly across different sets of websites, how many containers will be needed to separate different first party websites that share one or more third parties? This number provides us with a way to characterise and quantify the interconnectedness of the third party ecosystem for a given set of first party websites. We term this as the *Third Party Tangle Factor*, or simply the **Tangle Factor** of that set of websites. The higher the **Tangle Factor**, the more the interconnectedness of third party cookie ecosystem.




Figure 1: Overlapping first parties which share a third party tracker must be placed in separate containers, thus the red and green sites must be separated. The blue website does not share any third parties with either red or green and can be placed together with either of them in a container.

We calculate the tangle factor by modelling the set of FPs as nodes of a graph, drawing edges between two FPs when they share one or more TPs. We call this the *first party interconnection graph* (FPIG). Two FPs in the FPIG that share an edge (i.e., share one or more TPs) must be therefore placed in separate containers in order to prevent tracking by the shared TPs. If we assign a colour to each container, and label FPs in the FPIG with the colour of the container they are placed in, it is easy to see that the *vertex chromatic number* of the FPIG, i.e., the number of colours needed for nodes or vertices of the FPIG such that neighbouring vertices which share an edge are coloured differently, gives the minimum number of containers needed to effectively separate that set of first party websites. We call this the *tangle factor* of a given set of first party websites.

We apply the Tangle Factor metric to three different sets of first party websites. First, we look at the Alexa top- k websites for different values of k , and for two different countries, UK and China. Second, we leverage an ongoing user study³ which developed a Chrome plugin and collected anonymised browsing histories for a period of one year from a small cohort of users in the same two countries (UK and China). Our plugin has since been released in-the-wild, to help users to visualise their own browsing histories. We also provide these users an option to manually send us their histories, and contribute to our study. Our third and final set of websites is two months of browsing histories of 124 users from 25 countries who have decided to voluntarily contribute data to our user study.

We can use the tangle factor to understand the third party ecosystem from different vantage points. For example, we visit the top- k most popular websites from UK and China, and find that for the same set of websites (Global top 2K websites according to Alexa), visiting from the UK results in much higher interconnectivity than from China. In other words, the most popular sites have a higher tangle factor from UK locations, i.e., it requires many more separate containers to prevent third parties from tracking browsing of top- k websites in the UK, than in China.

A similar result carries over into actual browsing histories of real users in both countries, which are based on country-specific websites rather than synthetic “top- k ” websites: browsing histories of users in China have a lower tangle factor, i.e., are more easily separated, and with fewer containers, than browsing histories in

²In practice, a userContextID column is added to the cookie database, and only cookies matching the context ID of the container are sent to the website.

³This study has been approved by King's College London Research Ethics Committee (Approval no. MRS-1718-6539).

UK. We then expand to investigate the browsing histories of our in-the-wild users across 25 different countries, and show that the interconnectedness of websites has only a low (-0.0726) correlation with the actual numbers of third parties. Rather, it is the ubiquity of large third party trackers (e.g., Facebook and Google DoubleClick) which are present on a large proportion of websites, that increases the tangle factors.

Consequently, we explore a “what if” scenario where the most common third parties simply did not exist, or were prevented from operating (e.g., through ad blockers). We show that deleting the most common third party trackers, which corresponds to deleting the most common sources of edge creation in the FPIG, is a highly effective approach, and the tangle factor drops very quickly with the removal of the most common trackers. We then use this approach to measure the effectiveness of several ad blockers, and show that uBlock origin is more effective than Adblocker Plus, Ghostery and Ad Guard. We also show that *a)* that the largest containers of non-interfering first parties contain regional and computer-related websites, and *b)* using Adblocker Plus and uBlock origin results in UK interconnectivity dropping to levels similar to that of China. This lower interconnectivity is quantified in terms of the increase in the sizes of the largest set of non-interfering first parties (i.e., the size of the largest container when separating different first parties).

The rest of this paper is organised as follows: §2 defines the tangle factor, and provides details about our datasets. §3 quantifies the tangle factor by country, in terms of the numbers of containers needed to separate first party websites that share a third party; §4 uses the tangle factor to assess different methods to restrict the numbers of containers needed; §5 discusses related work and context, and §6 concludes.

2 DATA COLLECTION & METHODOLOGY

2.1 Browser plugin for data collection

The backbone of our data collection methodology is a plugin called “ThunderBeam”, which we have created for Google Chrome, by extending a Firefox extension called Lightbeam [32]⁴.

Thunderbeam not only allows users to log their browsing history, but it also provides support to track third-party networks *across* sites. The original Lightbeam for Firefox was based on an add-on called Collusion developed by Mozilla in 2012 [43]. Branching from Collusion, there is an add-on called Disconnect [11] for Google Chrome browsers. However, as opposed to Lightbeam, Disconnect can only log trackers in a uncontextualized manner (i.e., by looking at websites individually). Thus, Disconnect does not support tracking third parties across sites since there is no direct mechanism to capture and match the correspondence of first-party and third-party requests in Chrome. Our plugin develops several adaptations to make this work on Chrome, as detailed in a separate paper [20].

2.2 Third party disambiguation

In our analysis, we take the data provided by the Thunderbeam plugin, and identify different third parties. This requires some subtle disambiguation. First, we may have different domain names from the same third party entity: For instance, if there are both

`x.doubleclick.net` and `y.doubleclick.net` tracking `bbc.com`, we record them both as `doubleclick.net` and only count them once. Thus, we refer to third party services based on the 2^{nd} -level domain names [5].

Next, we may have different third party domains that all belong to the same parent entity. For example, DoubleClick and Google Analytics are both owned by Google. Following [23], We detect such cases by tracing the Authoritative DNS (ADNS) servers for the 2^{nd} -level domain names of TP entities, and merging TPs that share the same ADNS server.

Finally, some third parties use a cookie synchronisation mechanism to establish a “data sharing tunnel” between different third-party vendors. Following guidelines in [39], we detect cookie synchronisation by correlating shared unique userIDs embedded in cookies stored by different third parties.

2.3 First Party Interconnectivity Graphs (FPIG)

After merging third parties using the techniques mentioned above, we then consider all the third parties of a given set of first parties. We then model this as a graph, where the nodes are the first party websites, and edges are drawn between two nodes if the corresponding first party websites share a third party (after third parties are merged using the disambiguations discussed above).

2.4 Calculating Tangle Factor

The tangle factor of a given set of websites, i.e., a given FPIG, is calculated by computing an assignment of FPs to different containers, whilst respecting the restriction that two FPs that share a TP (i.e., two nodes connected by an edge in the FPIG) must be placed in different containers.




Figure 2: Restrictions in this first-party (FP) model example:

- ① FP₁ cannot be in the same container with FP₃;
- ② FP₂ cannot place with FP₄ and FP₆;
- ③ FP₆ cannot be with FP₅.

Effectively, this corresponds to a vertex graph colouring problem, where nodes with the same colour can be placed within the same container, and nodes which share an edge must be assigned different colours. Fig. 2 illustrates how a particular set of edges between different first parties would give rise to a container assignment. The minimum number of colours for the vertex colouring problem, i.e., the minimum number of containers needed in the FPIG, is the tangle factor of a given FPIG.

⁴Lightbeam is no longer a supported Firefox extension after Oct 2019

Table 1: Year-long (Jan 2018 to Jan 2019) data collection from 15 users in UK and China

User Group	1 st party sites	3 rd party cookies
UK Users	8416	113,003
CN Users	6144	74,313
Total	14,827	187,316

2.5 FPIG datasets

We apply the above methodology to obtain tangle factors for several different sets of first parties. First, we use an automated browsing system, built on top of Selenium [12] running in non-headless mode, and collect the cookies set by the Alexa⁵ top- k most popular websites. Specifically, we programmatically visit the country-specific Alexa top500 sites in UK and China, and also the global top2k websites from UK and China locations, in order to obtain a comparison between the number of containers required in both countries. The degree of demand for containers is taken as an indication of the degree of third-party risks in different countries. Further, we visit the UK top500 websites after installing different ad blockers (uBlock Origin, Adblock Plus, Ghostery and Adguard Adblocker) on two different browsers (Chrome and Firefox), to understand the privacy protection provided by different ad blockers.

To complement this dataset, we also collect data from real users who consented to support our work by providing anonymised browser histories⁶. Our users come in two cohorts. First, we have a small cohort of 9 users in the UK and 6 users in China, whose browsing activities were collected for a year-long period from Jan 2018–Jan 2019. Altogether, these users have visited around 15k first-party websites across one year, involving over 187k third-party domains (Table 1). We have also published the official version of our add-on in Chrome Web Store as “Thunderbeam-Lightbeam for Chrome”, and have seen over 1000+ installs of our plugin by Feb. 7th, 2020. This plugin offers users the option of submitting their data to our study. Through this mechanism, we have collected two-months of data from 124 users, who collectively provide us a picture of the third party ecosystem from 25 countries, as detailed in Table 2).


3 CONTAINER DEMAND BY COUNTRY

For any two sets of websites, the set with the lower tangle factor (i.e., the number of separate containers needed to ensure that a common third party does not learn about two first parties visited) is the one with the better privacy, and with less powerful third party tracking practices.

We begin by exploring the tangle factor of popular sets of websites from different countries. We then go on to show that tangle factor depends to a large extent on the country rather than actual numbers of third parties involved.

Table 2: Data collected from extension installers (ver.2) from Dec. 2019 to Feb. 2020: 124 users across 25 countries in total, Countries with fewer than 5 users are shown in gray.

Continent	Country	Num (Users)
Africa (8)	Tunisia	5
	Egypt	3
Asia (20)	China	8
	India	5
	Israel	3
	Singapore	2
	Thailand	2
Europe (70)	Germany	13
	United Kingdom	10
	France	10
	Netherlands	9
	Denmark	8
	Spain	6
	Italy	5
	Belgium	5
	Switzerland	3
	Luxembourg	2
	Norway	1
Sweden	1	
North America (15)	United States of America	8
	Canada	5
	Mexico	2
Oceania (3)	Austria	3
South America (5)	Chile	3
	Brazil	2

**Figure 3: Tangle factors for Alexa top2k (global) websites visited from UK and China vantage points.**

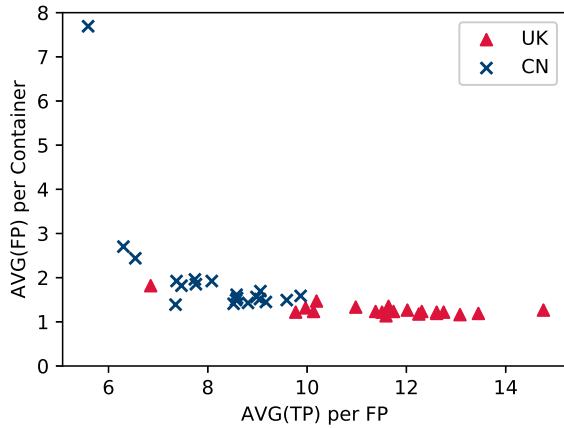



Figure 4: Linear correlation between TPs/FP and FPs/Container. Each point represents the average TPs/FP and FPs/container for one of the different ranked bins depicted in Figure 3. The Pearson correlation coefficient between TPs/FP and FPs/Container in China is -0.681 and UK is -0.728

3.1 Tangle factors of Alexa.com topsites visited from China and the UK


We begin by binning the Alexa top 2K websites (global ranks) into sets of 100 websites (i.e., the first bin has websites ranked 1–100, the second bin has websites ranked 101–200, and so on). We then ask whether users from different countries experience different levels of tracking, by accessing these websites from vantage points in the UK and China, using our automated browsing system (cf. Sec. 2.5). Fig. 3 shows the tangle factors of these sites for different rank bins. We find that overall, even when visiting the same sets of websites, users visiting from China are “less tangled” than users visiting from UK, i.e., need fewer containers. This is especially true for the most popular websites (ranks < 500). This is possibly because the Great Firewall of China blocks websites such as Facebook and Google, which are among of the most prevalent of trackers of Western countries [20]. The scatter plot of Figure 4 shows that indeed, the numbers of third parties per website is lower when visiting the Alexa global top2k websites from China than from the UK (blue points are largely to the left of red points with one exception), and that there is a strong (anti-)correlation between number of third parties loaded, and the number of containers needed for the trends among both countries.

3.2 Tangle factors of real users

The above result provides an indication of how blocking, especially at a countrywide level, can have an unintended positive side effect of decreasing the level of tracking by global giants for users from those countries. However, individual users’ browsing histories



(a) Regression between container numbers and the weekly FPs by each user with 95% confidence interval. ($r_{UK}^2 = 0.9688$; $r_{CN}^2 = 0.9162$)



(b) Residuals

Figure 5: Relation between numbers of first parties visited, and tangle factors (numbers of separate containers required) for weekly browsing histories of users from UK and China.

are unlikely to solely consist of the *global* top- k websites. Instead, country specific websites are likely to be hugely important.

We therefore next turn to understand the numbers of containers required for real users, by focusing on one year of browsing histories of our panel of users (Table 1). Figure 5(a) shows the number of containers needed (i.e., tangle factors) by individual users, as a function of how many first party websites they visit in each week. There is a clear correlation, with users who visit more FPs having higher tangle factors. However, in general, the Chinese users require fewer containers (have lower tangle factors) than UK users.

We can also compute a linear regression, with coefficients as:

$$n_{UKcontainer} = m_{UK} \times n_{firstpart} + c_{UK}$$

$$n_{CNcontainer} = m_{CN} \times n_{firstpart} + c_{CN}$$

where m_{UK}, m_{CN} represents the slope of the UK and CN growth lines respectively and c_{UK}, c_{CN} are constant terms. Figure 5(a)

³<https://alexa.com>

⁶This research is conducted under our university’s Ethics Approval no. MRS-1718-6539, and the data protection review by Chrome Web Store

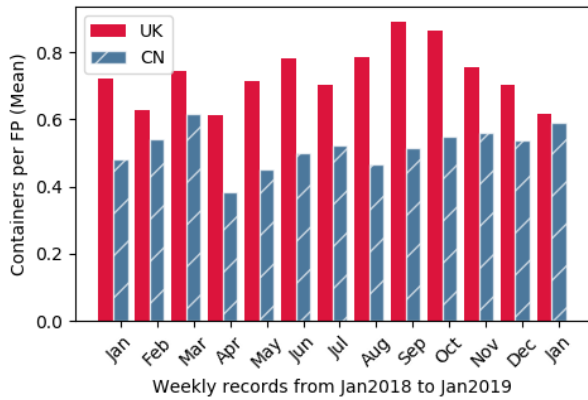


Figure 6: Average number of FPs stored in each container from Jan. 2018 to Jan. 2019 after ADNS disambiguation, based on weekly browsing records of UK and China participants.

shows the trends, and a good fit with high r^2 values. We also plot residuals in Figure 5(b), to check whether values are scattered around the y axis randomly, demonstrating that the linear regression model is appropriate for both UK and China datasets.

The regression lines of Figure 5(a) clearly show that the *slope* of the regression is lower in China than in the UK. Thus, even though heavier browsing leads to more tracking in both countries, the *growth* in tracking is slower in China than in the UK.

3.3 Container shareability in 25 countries

The lower tangle factor in China means that Chinese users need fewer containers to keep their browsing habits private from third parties. Figure 6 confirms this, showing that across the duration of the entire year of our study, Chinese users can pack more first parties into each container, without compromising privacy.

We expand on the above observation, and turn to the in-the-wild users of our Thunderbeam plugin (cf. Table 2). We ask how many first parties can be packed into containers (on average) for users in different countries around the world.

Figure 7 shows the results. In some countries such as China and Singapore (Figure 7(a)), it is possible to pack many more first parties into each container, whereas in others, the first parties tend to have common third parties, and therefore need to be placed in separate containers. We simultaneously plot the average numbers of third parties used by each first party, and the figures show visually that there is not necessarily a correlation. Note that this is also the case even when we discount countries with small numbers of users (countries with fewer than five users are grayed out). Figure 7(b) confirms the above visual observation with a scatter plot and corresponding Pearson correlation calculation, that shows there is no strong correlations between higher numbers of third parties used and the numbers of first parties that can be packed into a container. Rather, it is the *interconnectedness* of the first parties, i.e., the numbers of third parties shared, that determines whether or not two first parties can be placed into the same container.

4 RESTRICTING CONTAINER DEMAND

The results of the previous section lead us to consider ways to decrease the interconnectedness of the third party ecosystem in a given setting. We consider three options: First, we consider how effective are the default “content blocking” protections of different browsers, comparing Firefox and Chrome (§4.1). Then we consider what happens if browsers like Firefox generalised from the current special purpose “Facebook” containers, and instead just removed or blocked the top most prevalent third parties (§4.2). After this, we consider user interventions, such as installing ad blockers, which removes certain third parties. Each ad blocker removes different third parties based on their lists, and we use the tangle factor as a metric to understand which ones are the most effective (§4.3). Finally, in §4.4 we consider how well different categories of websites are able to co-exist with each other in the same container, without information leakage.

4.1 Chrome vs. Firefox

To begin with, we assess the default levels of privacy protection provided by two popular browsers - Firefox and Chrome. In particular, many casual or non privacy-conscious users may not install and deploy extensions such as ad blockers, which we consider later. However, even without installing extensions, there is a possibility of using mechanisms such as “Do Not Track” in most modern browsers, including Chrome⁷ and Firefox⁸, although it is known that Do Not Track is not very effective in practice [7, 19]. Compared with Chrome, Firefox provides users with a collection of additional privacy protection features known as content blocking⁹. Users could turn on strict content blocking in Firefox even without installing extensions, and prevent more harmful practices.

To test the effectiveness of these practices, we visit the Alexa top500 websites in an automated fashion, using Chrome and Firefox with the default settings. We then check the tangle factors, i.e., the number of containers required to separate third parties after the browsers have done their blocking (in the case of Firefox with content blocking), and taking into account the decrease in tracking due to the “Do Not Track” option.

Table 3: Number of containers required on Chrome and Firefox, when employing no extensions, but enabling privacy controls available by default.

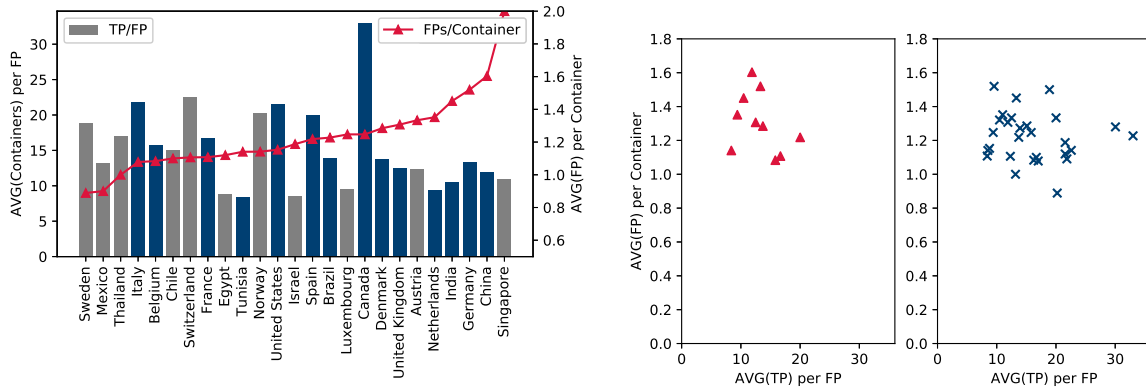
Num (Containers)	Chrome	Firefox
Original	408	410
Do Not Track	405	409
Strict Content Block	×	339

Table 3 shows the results. As expected, “Do Not Track” barely has any effect at all. It is merely a request to websites not to track, and if a site chooses not to respect it, there is no effect whatsoever on tracking. However, “content blocking” in Firefox decreases the

⁷Turn “Do Not Track” on or off in Chrome: <https://support.google.com/chrome/answer/2790761?co=GENIE.Platform%3DDesktop&hl=en>

⁸<https://support.mozilla.org/en-US/kb/how-do-i-turn-do-not-track-feature>

⁹Content blocking in Firefox: <https://support.mozilla.org/en-US/kb/content-blocking>



(a) Numbers of FPs in each container, and the corresponding numbers of in 25 countries (countries with fewer than 5 users grayed out) (b) Scatter plot of number of TPs per FP and number of FPs per container; in red for users in all 25 countries (Pearson Correlation Coefficient -0.301), and in blue for users in 13 countries with more than 5 users (correlation -0.360)

Figure 7: Comparison between the average number of TPs per FP and container per FP in countries.

number of containers required by about 17%, from 410 to 339 containers, because it executes additional blocking on the browser side. However, neither of these options is as effective as employing an ad blocker such as uBlock origin (which, as we will show in §4.3, results in a 40% reduction in number of containers needed).

4.2 Effectiveness of removing top third parties

Next step up in terms of user effort, Firefox has an interesting “suggested” add on that isolates Facebook logins from other websites. As Facebook is one of the most commonly used trackers, this is highly effective in decreasing overall levels of tracking. In this subsection, we ask, as a “what-if” scenario, what would happen if all the top-*k* trackers were removed or blocked by default.

This is visualised in Figure 8, which shows the First Party Interconnection Graphs (FPIG) of the Alexa Global top 500 websites, when these websites are visited from UK and China respectively (similar to the setup of Figure. 3). The graphs are drawn using a force-directed layout algorithm, with lays out the most connected nodes as the central core, and largely isolated nodes as a ring around the edges. The figure shows how even removing only the top 20 third parties can drastically decrease the tangledness of the FPIGs.

This informal but visually clear result is formalised in Figure 9, which shows how the tangle factor progressively decreases when visiting the global Alexa top 500 websites from UK and China, but with the top third parties removed. Initially, UK users need nearly 408 containers for the 500 websites, whereas CN users need 227. However, after removing just the top 50 third parties, the number of containers required drops to 9 and 8 respectively. Thus, *the third party ecosystem is highly interconnected mainly because of the predominance of a few large third parties. Protection against these large players can greatly decrease the extent of third party tracking in today’s web ecosystem.*

4.3 Assessment of ad blockers

Next, we use the same technique as above to determine the effectiveness of ad blockers. We visit the Alexa top 500 websites from a UK location, using Selenium in non-headless mode. We perform this experiment by installing four different popular ad blockers in turn: uBlock Origin, Adblock Plus, Ghostery and Adguard Ad-blocker. The selection of these four popular ad blockers is based on the recommendations in [4].

In Figure 10, we explore the performance of these ad blockers and show the percentage decrease in numbers of third parties as well as the tangle factor or numbers of containers required, when those ad blockers are deployed. This shows that uBlock origin performs the best, with nearly 60% reduction in raw numbers of third parties, and over 40% reduction in the number of containers required.

Both uBlock Origin and Adblocker Plus use Easy list for the list of third parties to filter. However, different default privacy settings could lead to different degrees of protection. In addition to Easy List, uBlock applies additional filters from Easy Privacy, Malware domain list and Peter Lowe’s tracker list[30]. These are enabled by default; thus, even if the user installs uBlock Origin without any custom settings, protection is provided by default. In contrast, Adblocker Plus takes a more moderate approach, and also allows some acceptable ads [2] (e.g., ads that comply with “Do Not Track” or those generated from the same origin as the first-party site), which results in the slight increase in the number of required containers. Adguard Adblocker’s poor performance in stopping third parties requests seems to be related to the fact that it hides ad elements after loading the entire site, rather than pre-blocking ad elements[1].

4.4 Interconnectedness across web categories

Finally, we ask how different categories of websites are able to co-exist with each other, given the above kinds of interventions. We use Alexa’s categorisation of the top 500 websites into 16 global categories, and examine the distribution of these categories in the




Figure 8: Force-directed layout of the FPIG of the Alexa Global Top 500 websites visited from a Chinese Location ((a) and (b)), and visited from the UK ((c) and (d)). The inner core is highly connected, and the outer ring is largely isolated nodes which can share a container. Nodes which can share a container are given the same colour. (b) and (d) show how the initial layouts in (a) and (c) for CN and UK respectively improve with many more isolated nodes after the top 20 third parties are removed.




Figure 9: Decrease in tangle factor as the top trackers are removed or blocked. After removing about 50 top third parties, UK and China respectively require only 9 and 8 containers, as opposed to initial numbers of 408 and 227 containers.




Figure 10: Percentage of decline in the number of containers and third parties, when users apply different ad blockers to visit Alexa top500 websites from a UK location. (Larger decline is better).

largest of the containers that may be formed after separating websites that share common third parties.

Fig. 11 shows the results. The first two columns show the category distribution of the contents of the largest container, when first parties from the Alexa top websites of China and UK are separated into containers based on shared third parties. The remaining four columns show the category distribution of the largest container for the the UK top 500 websites, when different ad blockers are applied in order from the least effective (Adguard) to the most effective (uBlock origin)¹⁰.

The CN500 column has a larger number of sites (94) than UK 500 (61), as tracking is less evolved in China. However, as more intrusive ad block extensions are introduced, the size of the largest container increases even for the UK500, and with both Ad block Plus (ABP) and uBlock origin, the largest container for UK is comparable to or larger than the largest container for China.

Looking across the categories, we find that the largest proportion of sites are those related to computers in both China and UK. Regional websites in the UK also track less and are therefore more easily incorporated into this large container.

5 RELATED WORK

Vyas *et al.* [44] propose to extend the same origin policy by adding a so-called origin attributes field, and separating cookies from different origin attributes. Our mechanism detects collision between two first parties automatically if third parties are shared, and can be used to create different origin attributes automatically. Thus, our method can feed into the origin attributes mechanism.

Origin attributes are used for contextual IDs in Mozilla multi-account containers[34]. Mozilla also introduced a special-purpose container targeted at the control of Facebook trackers [36]. Our work is inspired by these efforts, and generalises. Indeed, we show that removing the top 10-20 containers can have a hugely beneficial effect.

Related works such as [22, 29, 41] have been looking at better ways to detect online trackers, including anonymizing the *referrer* field in HTTP requests [25]. Although our work does not directly aim at blocking trackers, or attacking them in other ways (e.g.,[21]), identification of third parties is a paramount first step

¹⁰We do not apply Ad blockers to the CN 500 websites, as the ad blocker lists are not adapted for Chinese websites.

	CN500	UK500	Adguard	Ghostery	ABP	uBlock
<i>Num(max)</i>	94	61	73	88	97	112
adults	0.00%	2.86%	3.77%	3.51%	2.70%	2.44%
arts	0.00%	2.86%	1.89%	1.75%	0.00%	3.66%
business	10.00%	8.57%	9.43%	8.77%	12.16%	7.32%
computers	25.00%	17.14%	18.87%	21.05%	16.22%	23.17%
games	10.00%	2.86%	5.66%	3.51%	5.41%	6.10%
health	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
homes	0.00%	2.86%	1.89%	1.75%	2.70%	1.22%
kidsteen	0.00%	2.86%	1.89%	3.51%	2.70%	1.22%
news	0.00%	2.86%	1.89%	1.75%	2.70%	1.22%
recreation	5.00%	2.86%	3.77%	3.51%	4.05%	2.44%
reference	15.00%	11.43%	11.32%	10.53%	9.46%	8.54%
regional	10.00%	17.14%	15.09%	21.05%	22.97%	24.39%
science	5.00%	11.43%	7.55%	8.77%	6.76%	4.88%
shopping	0.00%	5.71%	7.55%	5.26%	5.41%	7.32%
society	10.00%	5.71%	5.66%	3.51%	4.05%	3.66%
sports	10.00%	2.86%	3.77%	1.75%	2.70%	2.44%

Figure 11: Proportion of the first-party websites from different Alexa web categories in the largest container

and a key concern for us. For this, we have referred to and used strategies, heuristics or third party lists from a number of efforts like ChromeDanger [6], Ghostery [12], Brave [16], AdReveal [27], Adblock [38], Plus [17], XRay [24], TrackAdvisor [25], and Disconnect [8]. Other works focus on advertising [25, 27] or on service media [42] alone as well as they do not consider country-specific trackers. Previous works like [23, 45], also mentioned the need for disambiguation of third parties based on authoritative DNS servers, which we use.

An overview of the evolution of the third-party tracking ecosystem is given in [29]. Authors show in 2012 that web measurements are an effective way to understand trackers. Later in 2015, authors in [26] provide an overview of the usage of cookies over 1 million sites. In another work from 2016, authors look at an advanced form of tracking that uses a cookie hijacking attack [40]. The setting proposed is adversarial and can therefore be considered less realistic than our study. More importantly, authors focus largely on DoubleClick, Google, and Amazon that hardly operate in China.

Our work looks at trackers today, but follows the design principles proposed in [29]. The size of our measurement is not as large as in [26]. However, we provide special attention to the most popular sites in different countries. We further consider the browsing habits of a user group that volunteer to our study.

In terms of the ad-block performance evaluation, [10] uses inspectors to evaluate the difference of blocking/capturing results caused by different strategies of PETs or browsers. For example, Ghostery and Disconnect only capture requests but do not modify ad attributions, while uBlock Origin uses filters to change the attributions of ad scripts to block the embedded advertisements. Our paper examines the performance of ad-blockers on the basis of their ability to restrict interconnections between first party websites.

6 CONCLUSION

In this paper, we considered the interconnectedness of the third party ecosystem, using vantage points from the UK and China to visit Alexa top 500 websites, and relying on real browsing histories

two cohorts of users: one, a carefully selected panel of 15 users, 9 from the UK and 6 from China, the other a set of 124 users from 25 different countries in the world who have chosen to donate data to our research project. We believe we are one of the first to use data from the browsing histories of real users. Note that we do not collect demographic information about our users due to privacy considerations.

We introduced a novel metric, which we call the “tangle factor”, to measure the interconnectedness of the third party ecosystem. The tangle factor is based on the insight that if we were to create a “first party interconnection graph” by drawing edges between first party websites which share a third party, the websites on either side of edge would share one or more third parties and therefore need to be placed in separate containers to prevent tracking. Thus, the minimum number of colours needed to colour this graph, i.e., its vertex chromatic number, also represents the number of containers needed. Using this metric, we showed that when visiting the same websites, users from Chinese locations are less tracked than users from UK locations, likely due to automatic blocking of major trackers like Google and Facebook from the Great Firewall of China. We also showed that this result carries over into the actual browsing histories of our panel of users, which are based on country-specific websites rather than global most popular sites considered in the synthetic evaluation.

We then used the tangle factor metric to assess the effectiveness of different methods of blocking trackers. We showed that blocking the top 20 trackers alone is sufficient to bring down the interconnectivity greatly, and only 9 containers, instead of 400 containers, are needed in the UK to separate out Alexa top 500 first parties that share a third party. We also used the tangle factor metric to compare ad blockers and showed that uBlock origin works better than others such as ghostery, Ad blocker plus and ad guard. We also showed that the default protection offered by Firefox is better than that offered by Google Chrome. These measurements are intended as proof-of-concept and our method can be expanded to compare other content blocking and protections apart from the ones we consider in this paper. These results provide quantitative evidence that the third party ecosystem is highly interconnected mainly because of a few large players, and protection against these can greatly decrease the extent and impact of tracking on the web.

REFERENCES

- [1] Adblock_Plus. 2018. Allowing acceptable ads in Adblock Plus. Available at <https://adblockplus.org/acceptable-ads>.
- [2] AdGuard. 2019. FAQ-What is the difference between AdGuard filtering methods? Available at <https://kb.adguard.com/en/android/faq#what-is-the-difference-between-adguard-filtering-methods>.
- [3] Pushkal Agarwal, Sagar Joglekar, Panagiotis Papadopoulos, Nishanth Sastry, and Nicolas Kourtellis. 2020. Stop tracking me Bro! Differential Tracking of User Demographics on Hyper-Partisan Websites. In *Proceedings of the The Web Conference (WWW 2020) (WWW '20)*. International World Wide Web Conferences Steering Committee, Taipei, Taiwan, 10.
- [4] Mshabab Alrizah, Sencun Zhu, Xinyu Xing, and Gang Wang. 2019. Errors, Misunderstandings, and Attacks: Analyzing the Crowdsourcing Process of Ad-blocking Systems. In *Proceedings of the Internet Measurement Conference*. ACM, 230–244.
- [5] Muhammad Ahmad Bashir and Christo Wilson. 2018. Diffusion of user tracking data in the online advertising ecosystem. *Proceedings on Privacy Enhancing Technologies* 2018, 4 (2018), 85–103.
- [6] Lujio Bauer, Shaoying Cai, Limin Jia, Timothy Passaro, and Yuan Tian. 2014. Analyzing the dangers posed by Chrome extensions. In *2014 IEEE Conference on Communications and Network Security*. IEEE, 184–192.

- [7] Frederik Zuiderveen Borgesius. 2013. Behavioral targeting: A European legal perspective. *IEEE security & privacy* 11, 1 (2013), 82–85.
- [8] Willem Boumans and Ir Erik Poll. 2017. Web Tracking And Current Countermeasures. (2017).
- [9] Cliqz. 2017. Ghostery. <https://github.com/ghostery>.
- [10] Amit Datta, Jianan Lu, and Michael Carl Tschantz. 2018. The Effectiveness of Privacy Enhancing Technologies against Fingerprinting. *arXiv preprint arXiv:1812.03920* (2018).
- [11] Disconnect. 2013. Take back your privacy. Available at <https://disconnect.me/>.
- [12] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. ACM, 1388–1401.
- [13] ExtUp. 2019. CookieSwap. Available at <https://addons.mozilla.org/en-GB/firefox/addon/cookieswap-q/> (last accessed on 28 Feb 2020).
- [14] Eyeo_GmbH. 2015. AdBlock Plus. <https://github.com/adblockplus>.
- [15] FDev. 2013. Swap my Cookie. Available at <https://chrome.google.com/webstore/detail/swap-my-cookies/dffhipnliikbklhjpapbecpmoilcama?hl=en> (last accessed on 28 Feb 2020).
- [16] Gertjan Franken, Tom Van Goethem, and Wouter I Joosen. 2019. Exposing Cookie Policy Flaws Through an Extensive Evaluation of Browsers and Their Extensions. *IEEE Security & Privacy* (2019).
- [17] Arthur Gervais, Alexandros Filios, Vincent Lenders, and Srdjan Capkun. 2017. Quantifying web adblocker privacy. In *European Symposium on Research in Computer Security*. Springer, 21–42.
- [18] Raymond Hill. 2016. uBlock Origin. <https://github.com/gorhill/uBlock>.
- [19] Chris Hoffman. 2019. RIP “Do Not Track,” the Privacy Standard Everyone Ignored. Available at <https://www.howtogeek.com/fyi/rip-do-not-track-the-privacy-standard-everyone-ignored/>.
- [20] Xuehui Hu, Guillermo Suarez-Tangil, and Nishanth Sastry. 2020. Multi-country Study of Third Party Trackers from Real Browser Histories. (Accepted). In *Proceedings of the 5th IEEE European Symposium on Security and Privacy (Euro S&P)*.
- [21] I Luk Kim, Weihang Wang, Yonghwi Kwon, Yunhui Zheng, Youstra Aafer, Weijie Meng, and Xiangyu Zhang. 2018. Adbudgetkiller: Online advertising budget draining attack. In *Proceedings of the 2018 World Wide Web Conference*. International World Wide Web Conferences Steering Committee, 297–307.
- [22] Balachander Krishnamurthy, Konstantin Naryshkin, and Craig Wills. 2011. Privacy leakage vs. protection measures: the growing disconnect. In *Proceedings of the Web*, Vol. 2. 1–10.
- [23] Balachander Krishnamurthy and Craig Wills. 2009. Privacy diffusion on the web: a longitudinal perspective. In *Proceedings of the 18th international conference on World wide web*. ACM, 541–550.
- [24] Mathias Lécuyer, Guillaume Ducoffe, Francis Lan, Andrei Papancea, Theofilos Petsios, Riley Spahn, Augustin Chaintreau, and Roxana Geambasu. 2014. Xray: Enhancing the web’s transparency with differential correlation. In *23rd {USENIX} Security Symposium ({USENIX} Security 14)*. 49–64.
- [25] Tai-Ching Li, Huy Hang, Michalis Faloutsos, and Petros Efstathopoulos. 2015. Trackadvisor: Taking back browsing privacy from third-party trackers. In *International Conference on Passive and Active Network Measurement*. Springer, 277–289.
- [26] Timothy Libert. 2015. Exposing the hidden web: An analysis of third-party HTTP requests on 1 million websites. *arXiv preprint arXiv:1511.00619* (2015).
- [27] Bin Liu, Anmol Sheth, Udi Weinsberg, Jaideep Chandrashekar, and Ramesh Govindan. 2013. AdReveal: improving transparency into online targeted advertising. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*. ACM, 12.
- [28] Dave Martorana. 2006. MultiFirefox. Available at <https://davemartorana.com/multifirefox/> (last accessed on 28 Feb 2020).
- [29] Jonathan R Mayer and John C Mitchell. 2012. Third-party web tracking: Policy and technology. In *2012 IEEE symposium on security and privacy*. IEEE, 413–427.
- [30] Johan Mazel, Richard Garnier, and Kensuke Fukuda. 2019. A comparison of web privacy protection techniques. *Computer Communications* 144 (2019), 162–174.
- [31] Daniele Moro, Filippo Benati, Michele Mangili, and Antonio Capone. 2018. Catching free-riders: in-network adblock detection with machine learning techniques. In *2018 IEEE 23rd International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*. IEEE, 1–6.
- [32] Mozilla. 2012. Firefox Lightbeam. Available at <https://github.com/mozilla/lightbeam-we>.
- [33] Mozilla. 2015. Multi-Account Containers. Available at https://support.mozilla.org/en-US/kb/containers#w_what-are-containers (last accessed on 28 Feb 2020).
- [34] Mozilla. 2017. Firefox Multi-Account Containers. Available at <https://addons.mozilla.org/en-GB/firefox/addon/multi-account-containers/versions/s/> (last accessed on 28 Feb 2020).
- [35] Mozilla. 2018. Facebook Container - Prevent Facebook from tracking you on other websites. Available at <https://support.mozilla.org/en-US/kb/facebook-container-prevent-facebook-tracking> (last accessed on 28 Feb 2020).
- [36] Mozilla. 2018. Facebook Container Extension: Take control of how you’re being tracked. Available at <https://blog.mozilla.org/firefox/facebook-container-extension/> (last accessed on 28 Feb 2020).
- [37] Muhammad Haris Mughees, Zhiyun Qian, and Zubair Shafiq. 2017. Detecting anti ad-blockers in the wild. *Proceedings on Privacy Enhancing Technologies* 2017, 3 (2017), 130–146.
- [38] Panagiotis Papadopoulos. [n.d.]. Analyzing the Impact of Digital Advertising on User Privacy. Available at http://users.ics.forth.gr/~panpap/thesis/panpap_phd_dissertation.pdf.
- [39] Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos Markatos. 2019. Cookie synchronization: Everything you always wanted to know but were afraid to ask. In *The World Wide Web Conference*. ACM, 1432–1442.
- [40] Suphannee Sivakorn, Iasonas Polakis, and Angelos D Keromytis. 2016. The cracked cookie jar: HTTP cookie hijacking and the exposure of private information. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 724–742.
- [41] Ashkan Soltani, Shannon Canty, Quentin Mayo, Lauren Thomas, and Chris Jay Hoofnagle. 2010. Flash cookies and privacy. In *2010 AAAI Spring Symposium Series*.
- [42] Jannick Kirk Sørensen and Hilde Van den Bulck. 2018. Public service media online, advertising and the third-party user data business: A trade versus trust dilemma? *Convergence* (2018), 1354856518790203.
- [43] Toolness. 2012. Collusion. Available at <http://www.toolness.com/wp/2011/07/collusion/>.
- [44] Tanvi Vyas, Andrea Marchesini, and Christoph Kerschbaumer. 2017. Extending the Same Origin Policy with Origin Attributes.
- [45] Craig E Wills and Doruk C Uzunoglu. 2016. What ad blockers are (and are not) doing. In *2016 Fourth IEEE Workshop on Hot Topics in Web Systems and Technologies (HotWeb)*. IEEE, 72–77.
- [46] Shitong Zhu, Xunchao Hu, Zhiyun Qian, Zubair Shafiq, and Heng Yin. 2018. Measuring and disrupting anti-adblockers using differential execution analysis. In *The Network and Distributed System Security Symposium (NDSS)*.