

On the Ubiquity of Web Tracking: Insights from a Billion-Page Web Crawl

Sebastian Schelter¹ and Jérôme Kunegis²

¹*Database Systems and Information Management Group, Technische Universität Berlin, Germany
sebastian.schelter@tu-berlin.de*

²*Namur Center for Complex Systems, University of Namur, Belgium
jerome.kunegis@unamur.be*

ABSTRACT

We perform a large-scale analysis of third-party trackers on the World Wide Web. We extract third-party embeddings from more than 3.5 billion web pages of the CommonCrawl 2012 corpus, and aggregate those to a dataset containing more than 140 million third-party embeddings in over 41 million domains. We study this data on several levels and provide the following contributions:

(1) Our work leverages the largest empirical web tracking dataset collected so far, and exceeds related studies by more than an order of magnitude in the number of domains and web pages analyzed. As our dataset also contains the link structure of the web, we are able to derive a ranking measure for tracker occurrences based on aggregated network centrality rather than simple domain counts. We make our extracted data and computed rankings available to the research community.

(2) On a global level, we give a precise figure for the extent of tracking, give insights into the structural properties of the ‘online tracking sphere’ and analyse which trackers (and subsequently, which companies) are used by how many websites, leveraging our ranking measure derived from the link structure of the web.

(3) On a country-specific level, we analyse which trackers are used by websites in different countries, and identify the countries in which websites choose significantly different trackers than in the rest of the world. In particular, the three tracking domains with the highest PageRank are all owned by Google. The only exception to this pattern are a handful of countries such as China and Russia. Our results suggest that this dominance is strongly associated with country-specific political factors such as freedom of the press.

(4) We investigate whether the content of websites influences the choice of trackers they use, leveraging more than ninety thousand categorized domains. In particular, we analyse whether highly privacy-critical websites about health and addiction make different choices of trackers than other websites. Our findings indicate that websites with highly privacy-critical content are less likely to contain trackers (60% vs 90% for other websites), even though the majority of them still do contain trackers.

Keywords: Web, Tracking, CommonCrawl

ISSN 2332-4031; DOI 10.1561/106.00000014
©2018 S. Schelter and J. Kunegis

1 Introduction

The ability of a website to track the pages read by its visitors has been present since the beginnings of the World Wide Web. In recent years however, another tracking mechanism has become widespread: that of third-party websites embedded into the visited site by mechanisms such as JavaScript and images.

In fact, the majority of websites contain third-party content, i.e., content from another domain that a visitor’s browser loads and renders upon displaying the website. Such an embedding of third-party content has always been possible, but was relatively rare, since most embedded images were located on the same server as the page itself, and in any case, the embedding of content was not intended for tracking. With the rise of social media and *Web 2.0*, websites increasingly began

to embed links (using various technologies) to third-party content, and thereby allowed the providers of such content to track users on a wide scale. The inclusion of third-party content occurs for a variety of reasons, of which not all require tracking, e.g., advertising, conversion tracking, acceleration of content loading, and provision of widgets. Social media and advertising companies primarily use the data collected via these tracking mechanisms to improve their capability to show personalized, tailored advertisements to their users, which represent their main source of income.

Regardless of their primary purpose, third-party components can (and in many cases do) track web users across many sites and record their browsing behavior. Therefore, these third-parties constitute a privacy hazard in many aspects, such as the collection of data about health conditions (Electronic

Frontier Foundation, 2015; Libert, 2015b), news consumption (Trackography, 2014), or their instrumentalization for mass surveillance by intelligence agencies (The Intercept, 2015). In order to understand and control these hazards, it is desirable to gain a deeper understanding of the ‘online tracking sphere’ as a whole. Previous research has studied small samples of this sphere, e.g., 1,200 English-language domains from Alexa’s popular sites (Krishnamurthy and Wills, 2009b), while researchers have only recently started to study larger datasets, e.g., tracking on the Alexa top 1 million domains (Libert, 2015a). The main reason for this is that, until lately, such a study would only have been possible for large companies possessing their own ‘copy’ of the web. Recent developments however allow us to study this online tracking sphere at a large scale: the availability of enormous web crawls comprised of hundreds of terabytes of web data (Spiegler, 2013), such as CommonCrawl 2012.¹

We process all 3.5 billion webpages from this corpus (which amounts to more than 200 terabytes of data), and extract the bipartite ‘tracking graph’, which describes the tracking of more than 41 million pay-level domains by 355 tracking services. A pay-level domain (PLD) is a sub-domain of a public top-level domain that users pay for individually. PLDs allow us to identify a realm, where a single organization is likely to be in control², e.g., the PLD for `www.example.com` would be `example.com`. Subsequently, we use this data to study three key research questions: (i) Which are the predominant tracking companies on the Web, and how many websites do they track? (ii) How does the distribution of web trackers vary from country to country? How is that variation related to different political, socio-cultural and economic factors? (iii) How are trackers distributed with respect to highly privacy-critical topics such as health and addiction? Do authors of such websites avoid web trackers?

To the best of our knowledge, the data extracted in the paper constitutes the largest empirical dataset of web trackers collected so far, containing an order of magnitude more domains than comparable studies. In this article, we extend our findings from previous work (Schelter and Kunegis, 2016) in several directions: (1) We conduct an analysis of tracking in the long tail of domains on the web, looking into 41 million pay-level domains (Section 5), and identify global structural properties of the online tracking sphere. (2) We analyze the tracking company distribution in top-level domains (TLDs) of non-English speaking countries, and find that a small set of US-based companies have a dominating role in the majority of countries; a correlation analysis suggests that this role is strongly associated with political factors such as freedom of the press (Section 6). Furthermore, we investigate two additional research directions: (3) We uncover country-specific as well as category-specific patterns in the relations between third-party trackers by clustering their co-occurrences (Section 5.5). (4) We compare the tracking capabilities on more than ninety thousand highly privacy-critical and less privacy-critical domains, finding that overall, websites with highly privacy-critical content do avoid trackers

as opposed to other websites, even though the majority of them still do contain trackers (Section 7). We also show that certain types of trackers are in fact present with higher probability on websites with privacy-critical content, indicating a lack of awareness of the tracking ability of their underlying services. (5) Finally, we provide our dataset to the scientific community³ and contribute it to the *Koblenz Network Collection* (Kunegis, 2013)⁴.

2 Online Tracking Fundamentals

In this section, we give a technical introduction to web tracking, and highlight resulting privacy implications.

2.1 Technical Foundations

In its basic form, online tracking involves *users* browsing the Web and the *website* that they intentionally visit. In addition to these two actors, one or more services may be present that record users’ browsing to the website; we call such services *third-parties*. We refer to third-parties as *trackers* if their main purpose or the business model of their owning company depends on collecting browsing data of users. Therefore, we categorize advertising services and social network plugins as trackers, because their main source of income comes from targeted, personalized advertisements which heavily depend on user data. On the other hand, we do not consider content delivery networks as trackers, because their main business is to accelerate website loading (see Krishnamurthy and Wills, 2009b; Englehardt and Narayanan, 2016).

Specifically, the user visits a website, which means that her browser issues an HTTP “GET” request to the web server hosting the website. The web server returns an HTTP response containing the HTML code of the website to render. This returned HTML code typically contains references to external resources, such as style sheets, JavaScript code and images that are required to render the page in the client browser. Next, the user’s browser automatically issues requests for these additional resources. In many cases, resources used by a web page reside on a different server than the one hosting the website. This possibility allows third-parties to track the user by recording and inspecting the external requests to their servers.

In the case of online tracking, the website’s HTML code embeds external resources from a third-party which aims to track the user. A typical example for such an external resource is a piece of JavaScript code, which the user’s browser will automatically load and execute from the third-party server. This external loading enables the third-party to record a wide variety of information about the user: (i) The third-party sees the current IP address of the user machine, which reveals the internet service provider as well as the approximate geolocation of the user. The browser typically announces its version, the underlying operating system, the user’s screen resolution and other information which allow the third-party to compute a ‘fingerprint’ of the browser. These fingerprints have been

¹<http://blog.commoncrawl.org/2012/07/2012-crawl-data-now-available/>

²<http://webdatacommons.org/structureddata/vocabulary-usage-analysis/>

³available at <https://ssc.io/trackingthetrackers>

⁴available at <http://konect.cc/networks/trackers-trackers/>

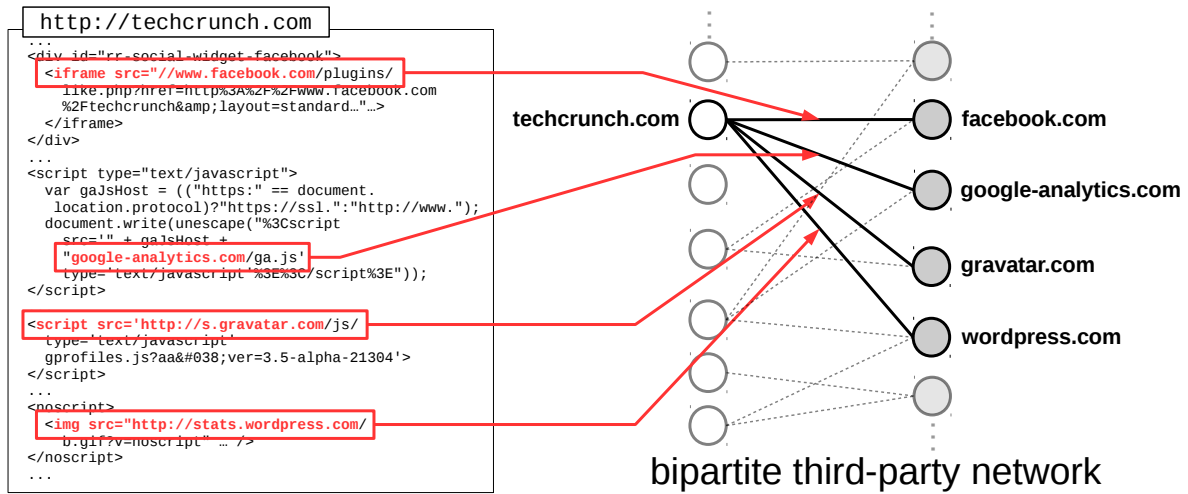


Figure 1: Example of our extraction process which results in the bipartite third-party network. We parse the HTML code of a website to extract the pay-level domains of all embedded third-parties. The captured embeddings form edges in the bipartite third-party network.

found to be sufficient to recognize individual users with a high precision (Eckersley, 2010). (ii) The third-party has access to the URI of the page which the user is currently visiting, and, by looking at the HTTP referrer, also learns the URI of the previous page the user has visited. (iii) The third-party is able to read existing cookies designated for its domain and can set new cookies. If the user chose to be persistently logged in to a service such as Google, Facebook or Twitter, these third-parties will recognize the user as they will receive the persistent login cookies. Third-parties are either embedded dynamically such as via JavaScript or the `iframe` element, or statically via `link` or `image` tags. In the latter case, they lose some of their tracking abilities such as reading the referrer and browser properties such as the screen resolution.

2.2 Privacy Implications

Third-parties vary considerably. They include advertisers, analytics services, social widgets, content delivery networks and image hosters. All of these have legitimate uses such as monetization, webpage optimization, A/B testing, conversion tracking, the provision of widgets, and fraud detection. However, the ability of many third-parties to record large portions of the browsing behavior of many users across a huge number of sites on the web poses a privacy risk, and is the subject of ongoing legal disputes (The Guardian, 2015a). The data recorded by this tracking infrastructure has been reported to contain large portions of online news consumption (Trackography, 2014), as well as intimate, health-related personal data (Electronic Frontier Foundation, 2015; Libert, 2015b), and personal data from social networks (Mayer and Mitchell, 2012; Krishnamurthy and Wills, 2009a). A large variety of browser-addons, (e.g., *Disconnect* 2016; *Ghostery* 2016; *Privacy Badger* 2016) have been developed to allow individuals to block online trackers from their computer. The ability to consume news and form a political opinion in an independent and unwatched manner, as well as the privacy of personal health-related data are vital for an

open society, and should not be subject to commercially motivated data collection. Furthermore, recent reports suggest that intelligence agencies piggyback on online tracking identifiers to build databases of the surfing behavior of millions of people (The Intercept, 2015). Frighteningly, this data collection seems to take place outside of the legal supervision of the governments and parliaments of many of the people affected.

3 Data Acquisition

3.1 Collection Methodology and Limitations

CommonCrawl 2012 is a very large publicly available crawl of the web, which consists of more than 3.5 billion HTML pages and amounts to approximately 210 terabytes in uncompressed form (Spiegler, 2013). While the CommonCrawl 2012 dataset is four years old as of 2016, we choose it over latter corpora, as it produces a more faithful representation of the underlying link structure of the web. This is due to the fact that it has been created using a breadth-first search crawling strategy (Lehmberg *et al.*, 2014), which leads to a more realistic representation of the connected network structure of the web compared to the crawling strategy employed in latter corpora. In these latter corpora, the crawling process fetched URLs from a predefined seed list and did not follow links, which results in networks with an unreasonably sized largest strongly connected component, which is not well suited for structural analysis.⁵

During our extraction process, we represent websites and third-parties by their pay-level domains, basing our data acquisition strategy on the fact that the URLs of third-party services have to be embedded in the HTML of the websites. Our extractor takes an HTML document as input and retrieves all embedded domains of third-party services, as shown in Figure 1. In a first pass through the document, we parse the HTML

⁵<http://webdatacommons.org/hyperlinkgraph/2014-04/topology.html>

and investigate the `src` attribute of `script`, `iframe`, `link` and `image` tags. In order to also find third-parties that are embedded via JavaScript code, we run a parser⁶ on all JavaScript code (but do not execute the code itself), and collect string variables that match a URI pattern. This investigation of the JavaScript parse tree allows us to also detect trackers that are dynamically added into the website’s HTML from Javascript code, e.g., `scorecardresearch.com` in the following example:

```
<script type="text/javascript">document.write(
unescape("%3Cscript src=' " +
(document.location.protocol == "https:" ?
"https://sb" : "http://b") +
".scorecardresearch.com/beacon.js'
%3E%3C/script%3E"));</script>
```

This approach allows for scaling to unprecedented data sizes, exceeding related work by more than an order of magnitude in the number of domains and pages analyzed. However, we note that it also imposes a set of limitations in comparison to data collection techniques that rely on instrumenting real browsers or on the collection of users’ HTTP traces (cf. Krishnamurthy and Wills, 2006; Roesner *et al.*, 2012; Libert, 2015b; Englehardt and Narayanan, 2016). One limitation of the use of the static CommonCrawl corpus is that transient trackers are excluded, i.e., trackers which are generated from JavaScript code that is fetched and executed dynamically at rendering time of the web page. Furthermore, we do not manually investigate the data recorded by each of the trackers, nor information that may influence tracking such as caching policies. We focus our analysis on the ability of third-parties to track (instead of the detailed tracking mechanisms) and on the high-level structural patterns of the tracking sphere.

We run a MapReduce (Dean and Ghemawat, 2010) implementation of our extractor in a massively parallel manner on the CommonCrawl 2012 web corpus, via the Amazon Elastic MapReduce service. In aggregate, the extraction takes about 60 hours using a cluster of 20 ‘m4.2xlarge’ instances (8 virtual CPUs and 32 GB of RAM per instance).

3.2 Dataset Statistics and Characteristics

Our dataset contains 3,536,611,510 web pages, which we aggregate by pay-level domain to create the *bipartite third-party network*. This network represents the embeddings of 12,756,244 third-parties (the right vertex set) into 41,192,060 pay-level domains (the left vertex set)⁷. We create an edge between a website domain and a third-party domain whenever we find the third-party embedded at least once in the web pages belonging to a website. For the identification of pay-level domains, we rely on the `InternetDomainName` class contained in the popular *Guava* core libraries for Java by Google.⁸

Table 1: Networks used in our study.

Count	Entities	Role
<i>Bipartite third-party network</i>		
41,192,060	Website PLDs	Left vertex set
12,756,244	Third-party PLDs	Right vertex set
140,613,762	Third-party embeddings	Edge set
<i>Bipartite tracking network</i>		
41,192,060	Website PLDs	Left vertex set
355	Tracking PLDs	Right vertex set
36,982,655	Third-party embeddings	Edge set
<i>Hyperlink network (obtained from webdata commons)</i>		
41,192,060	Website PLDs	vertex set
623,056,313	Hyperlinks	edge set

Next, we order the third-parties by the fraction of domains on which they occur and select the top third-parties for manual labeling. We leave unlabeled the long tail of third-parties that are only present on a few thousand out of the 41 million domains to keep the labeling effort manageable. We enrich the data for the resulting 1,375 third-parties as follows: We determine the registration countries and registering organizations for the third-party domains from WHOIS records.⁹ Then, we manually check the websites of the domains to determine their owning companies. Next, we label the third-parties according to their purpose and business model into categories such as ‘content delivery’, ‘analytics’, ‘advertising’, ‘website builder’, ‘image hoster’, and make the resulting labeled data available for the research community. Note that we did not have access to a data source denoting whether some of these trackers are being operated by governmental actors and therefore cannot provide insights into the question of access to tracking data by states.

From 1375 labeled services, we find 355 pay-level domains to belong to tracking services, as these match our definition of trackers, which we introduced and discussed in Section 2.1. (We refer to third-parties as trackers if their main purpose or the business model of their owning company depends on collecting browsing data of users.) Thereby, we construct the *bipartite tracking network*, which represents the 36,982,655 embeddings of the 355 tracking services in the 41,192,060 website pay-level domains of our corpus. Table 1 gives an overview over the datasets we use for this study. We acknowledge that it is difficult to come up with an objective, universally agreed upon labeling and selection of trackers; therefore, we provide our obtained and curated data to the scientific community to increase the transparency of our analysis.¹⁰

Additionally, we obtain the network of 623,056,313 hyperlinks between the pay-level domains in CommonCrawl (*Web Data Commons – Hyperlink Graphs* n.d.), for ranking websites according to the link structure of the web.

⁶<https://github.com/google/closure-compiler>

⁷Note that the extremely high number of embedded domains stems from the fact that all externally embedded resources (such as images) are counted here.

⁸<https://github.com/google/guava>.

⁹Unfortunately, we are constrained to current records as we do not have access to a historic database going back to 2012.

¹⁰available at <https://ssc.io/trackingthetrackers/>

4 Related Work

In this section, we give a brief overview over related work on web tracking and analysis of the CommonCrawl web crawl, and set this work in context to our findings. The analysis of the web tracking phenomenon has received much attention in the research community. In contrast to previous studies, our data acquisition technique allows us to look at an order of magnitude more domains, and several orders of magnitude more pages. With the exception of work by Englehardt and Narayanan, (2016)¹¹, related studies do not apply a ranking to tracker occurrences on domains. This can distort the results, as a tracker occurrence on an important, frequently visited domain is counted with the same weight as an occurrence on a less frequently visited page. Therefore, in Section 5.1, we introduce a measure for ranking trackers based on the accumulation of PageRank (Page *et al.*, 1999), a well-known approximation of the relevance of a website. We can compute these rankings in a replicable way, because PageRank is based on the underlying link structure of the webgraph, which we have available as part of our data.

4.1 Global Analysis of Web Tracking

Krishnamurthy and Wills, (2006) investigate web tracking and its privacy implications on about 1,000 domains by looking at the connections between visited domains and ‘hidden domains’ (similar to our bipartite third-party network). They also conduct a country-specific analysis by analyzing trackers on the top 100 domains for 68 countries, and find Google Analytics to be the dominating tracker. They extend this work in (Krishnamurthy and Wills, 2009b) where they investigate how the amount of tracking develops over time, both by single trackers as well as on a company level. They cover 1,200 English-language domains from Alexa top sites over four years between 2005 and 2008, and detect an increasing aggregation of private data by a decreasing number of trackers: the penetration of the top 10 trackers among popular sites has grown from 40% to 70% in the period of investigation. In contrast to our work, they do not report on tracking on privacy-critical categories of websites and do not investigate country-specific tracking patterns. Roesner *et al.*, (2012) develop a client-side method for detecting web trackers on 500 popular domains. They create a taxonomy of five different tracker types, based on how these manipulate browser state. Analogous to (Krishnamurthy and Wills, 2006), their data suggests a dominating role of Google Analytics. Furthermore, they estimate how much of users web search browsing sessions is visible to trackers by sampling from query logs, and find that more than 20% can be covered by several trackers. The authors rank trackers solely according to the number of appearances on a selection of domains and report surprisingly high ranks for unknown trackers such as `scorecardresearch.com`. We think that our ranking based on the accumulation of PageRank provides a much more realistic measure of the likelihood of being tracked, as it incorporates the link structure of the Web. In recent years, research has started to investigate tracking on a larger numbers

of domains. Libert, (2015a) presents a quantitative analysis of tracking mechanisms on the top 1 million sites from the Alexa ranking. He finds that nearly nine in ten websites leak user data to third-parties, and that a handful of American companies (including Google, Facebook and Twitter) receive the vast bulk of this user data. In contrast to our work, this study puts a strong focus on the potential of state surveillance, but only briefly differentiates its results between different countries or website categories. Additionally, Libert, (2015b) researches the privacy risks imposed by visiting health-related web pages on the web, based on the top 50 search results for about 2,000 common diseases. He finds that more than two thirds of these pages leak information about specific conditions, treatments and diseases to third-parties. Yu *et al.*, (2016) develop a novel defense approach based on k -anonymity. For that, they first process tracking data on 21 million pages from 350,000 domains, and collect traces of 200,000 internet users. Next, they show how to leverage this data collectively to dynamically identify unsafe elements and remove these elements from the requests. Englehardt and Narayanan, (2016) present ‘OpenWPM’, a comprehensive and scalable tracking measurement platform, which simulates real browsing. They use this platform to collect data on the top 1 million domains from Alexa. They also identify the heavy-tailed distribution of tracking capability, postulate the dominating role of Google Analytics and state that many of the top-occurring third-party domains belong to Google. This constitutes the only related study that also applies a weighting of occurrences of trackers on domains. They continue to investigate a wide variety of aspects of tracking, such as the low adoption of HTTPS encryption by trackers, they evaluate tracking protection techniques, and provide new insights into sophisticated tracking mechanisms, such as cookie syncing, ‘promiscuous’ cookies, and previously unknown fingerprinting techniques.

4.2 Web Tracking by Social Networks

A special focus has been given to actors that know about the real identities of internet users, such as online social networks. Krishnamurthy and Wills, (2009a) investigate the leakage of personally identifiable information from social networks to third-parties and suggest protection mechanisms. They extend this research to leakage from online social networks in (Krishnamurthy and Wills, 2010), and find that similar leakage happens there. Roosendaal, (2011) exemplarily investigates the privacy implications of the Facebook ‘like’ button. Chaabane *et al.*, (2012) look at the tracking capabilities of the three major social networks (Facebook, Twitter and Google+) on the top ten thousand domains. The authors also study tracking across different categories of websites, and find that tracking is encountered on all categories of sites. However, they do not perform an explicit analysis distinguishing between the amount of tracking on privacy-critical and less privacy-critical categories (where we encounter a reduction in tracking). The work exclusively covers trackers related to online social networks.

¹¹who use a measure that relies on externally provided, non-replicable domain rankings

4.3 Tracking Mechanisms and Detection

Eckersley, (2010) investigates ‘fingerprinting’ techniques, based on version and configuration settings of modern web browsers. They find that these techniques work surprisingly well at identifying individual users: for a random browser, only one in about three hundred thousand browsers will share its fingerprint. Furthermore, they developed a well-known test site¹² for investigating such fingerprints. Acar *et al.*, (2014) present an in-depth study of advanced tracking mechanisms such as browser fingerprinting via canvas images, re-spawning of HTTP cookies via the Adobe Flash plugin and cookie syncing between different trackers. They conclude that modern browsers – with the exception of the Tor Browser (Perry *et al.*, 2015) – fail at effectively protecting users from most of these techniques. Bau *et al.*, (2013) develop a machine learning-based approach for real-time tracker identification, which draws its features off a network derived from of script loading relationships in web pages. Kalavri *et al.*, (2016) inspect a graph similar to our bipartite third-party network, which they derived from user traces collected by a web proxy. Their aim is to automatically identify trackers among the third-parties, based on structural properties of their collected graph. They show that simple classifiers such as a nearest neighbor approach, as well as label propagation techniques perform surprisingly well on this identification task. Englehardt *et al.*, (2015) investigate the surveillance implications of passive eavesdroppers (e.g., intelligence agencies) piggybacking on existing tracking identifiers, and find that this allows them to reconstruct about two thirds of people’s browsing histories.

4.4 Structural Analysis of CommonCrawl

Spiegler, (2013) presents an exploratory analysis of the CommonCrawl 2012 corpus, computing elementary statistics such as the distribution of top-level domains, character encodings and media types. Lehmborg *et al.*, (2014), Meusel *et al.*, (2014), and Meusel *et al.*, (2015) study the fundamental graph structure of the web using the CommonCrawl 2012 dataset. Their findings suggest that the previously reported ‘bow-tie’ structure from (Broder *et al.*, 2000) is an artifact of the applied crawling process, and not a structural property of the web. Furthermore, their data shows that the distributions of in-degree, out-degree and sizes of strongly connected components of the page and host graph do not follow power laws.

5 Global Analysis of Tracking Services

In this section, we study the extracted third-party and tracking networks to gain insights about the distribution of online tracking services on the web. We want to answer the question to what proportion the browsing behavior of users on the web is visible to particular tracking services. Furthermore, we are interested in how strongly the tracking capabilities differ among various services. Unfortunately, there is a lack of openly available data sources that allow to quantify the number of visitors over time for the 41 million extracted pay-level domains in




Figure 2: Google owns the three tracking websites with the largest rank share. The twenty most common third-parties are shown by rank share. Tracking third-parties are highlighted in red, and domain share is shown as small black bars.

2012. Therefore, we resort to PageRank (Page *et al.*, 1999), a measure of the relevance of websites, as proxy for ranking them by the traffic which they attract. We compute the PageRank distribution of the domains in the hyperlink network to get an importance ranking for all the pay-level domains in our corpus. Next, we use this distribution to define our main ranking measure for the tracking capability of a third-party.

5.1 Ranking Tracking Services

We derive a ranking measure for third-parties from the PageRank value p of the pay-level domains in the hyperlink network as follows. Let D denote a set of domains to inspect, e.g., all pay-level domains belonging to a certain top-level domain, and let $t(D)$ denote the subset $\{d \mid d \in D \wedge t \text{ embedded in } d\}$ of domains contained in D having third-party t embedded. We define the *rank share* $\rho_{D,t}$ of a third-party t in domain set D as the sum of the PageRank values of domains from D that have t embedded, normalized by the overall sum of the PageRank values of domains in D :

$$\rho_{D,t} = \frac{\sum_{i \in t(D)} p_i}{\sum_{i \in D} p_i} \quad (\text{rank share})$$

We use the common parameter setting of 0.85 for the damping factor of the PageRank computation. In some cases, it is useful to look only at the number of pay-level domains visible to a third-party, regardless of their individual relevance. In these cases, we additionally report the *domain share* $d_{D,t}$ of third-party t in domain set D as the number of domains $t(D)$ in D having t embedded, normalized by the overall number of domains in D :

$$d_{D,t} = \frac{|t(D)|}{|D|} \quad (\text{domain share})$$

¹²<https://panopticklick.eff.org>




Figure 3: Cumulative distribution of the number of pay-level domains visible to tracking services in the bipartite tracking network.

5.2 Predominant Third-Parties

We first compute the third-parties with the highest rank share in our corpus. Figure 2 shows the top twenty third-parties by rank share (and also reports their domain share as small black bars). By far, the most common third-party is `googleanalytics.com` with a rank share of 0.507, which implies that the pay-level domains embedding `googleanalytics.com` amount to more than half of the mass of the PageRank distribution in our corpus. Its domain share is 0.248, implying that it is embedded on 24.8% of all pay-level domains. We find that five out of the ten most dominant third-parties belong to Google. The next dominant family are social media-related third-parties such as `facebook.com`, `twitter.com` and `addthis.com`. On the lower end, we find content delivery services such as `twimg.com`, the image hosting platform of Twitter, the cloud platform Amazon Webservices (`amazonaws.com`) and Facebook’s content delivery platform `fbcdn.net`. We highlight third-parties that match our definition of tracking services from Section 2.1 in Figure 2, which includes the dominant third-party `googleanalytics.com` as well as eight additional out of the twenty predominant third-parties.

5.3 Distribution of Large and Small Trackers

Web trackers as analysed in this article differ widely in size: A few trackers are used on millions of sites, while many more are used only on a few thousand. Such skewed distributions are often modeled by *power laws*, and can be explained by positing preferential attachment, i.e., a situation in which the most popular trackers are added to more new sites than small trackers, increasing the disparity over time, and leading to a monopoly or oligopoly of trackers dominating a large majority of all websites. In order to study this discrepancy, we investigate the differences in tracking capability between individual tracking services. Therefore we restrict the following analysis to the bipartite tracking network, which only contains the third-parties that match our definition of tracking services. We investigate the distribution of the number of pay-level domains visible to an individual tracking service. This distribution corresponds to the degree distribution of the right vertex set in the bipartite tracking network. Figure 3 shows the corresponding cumulative distribution function (on logarithmic axes) of this distribution. We encounter a highly disproportionate distribu-




Figure 4: Heatmap representing the conditional probabilities $P(t_1 | t_2)$ of encountering tracking third-party t_1 (Y axis) on a website pay-level domain, given that this website has tracking third-party t_2 (X axis) embedded.

tion: 50% of the tracking services are embedded on less than ten thousand pay-level domains, while tracking services in the top 1% of the distribution are integrated into more than a million pay-level domains. To assess whether this distribution follows a power law, we apply the methodology by Clauset and colleagues (2009), which uses a goodness-of-fit test to find a minimum value of the degree from which a power law is valid, as well as estimates the power law exponent. We find that starting from degree 6,848, the distribution indeed follows a power law with exponent 1.725, which is at the upper end of the range of exponents observed in many other hyperlink networks¹³, and which is compatible with a phenomenon of preferential attachment, i.e., trackers already deployed to many websites attracting proportionally more new sites than others.

Another related question concerns the use of multiple trackers: Do sites using few (or many) trackers prefer big (or small) trackers? This question can be answered by investigating the assortativity of the tracking network. In fact, we find that low degree domains (websites with few trackers) tend to connect to high-degree tracking services. This amounts to a dissortative mixing with a negative correlation coefficient of -0.1863 . The found dissortativity is statistically significant to a p -value under 0.001. This encountered dissortativity is again consistent with findings for other hyperlink networks¹⁴, and implies that sites using few (including a single) trackers use large trackers more than an average site, in turn increasing the predominance of a small number of very widely deployed tracking services.

¹³<http://konect.cc/statistics/power/>

¹⁴<http://konect.cc/statistics/assortativity/>

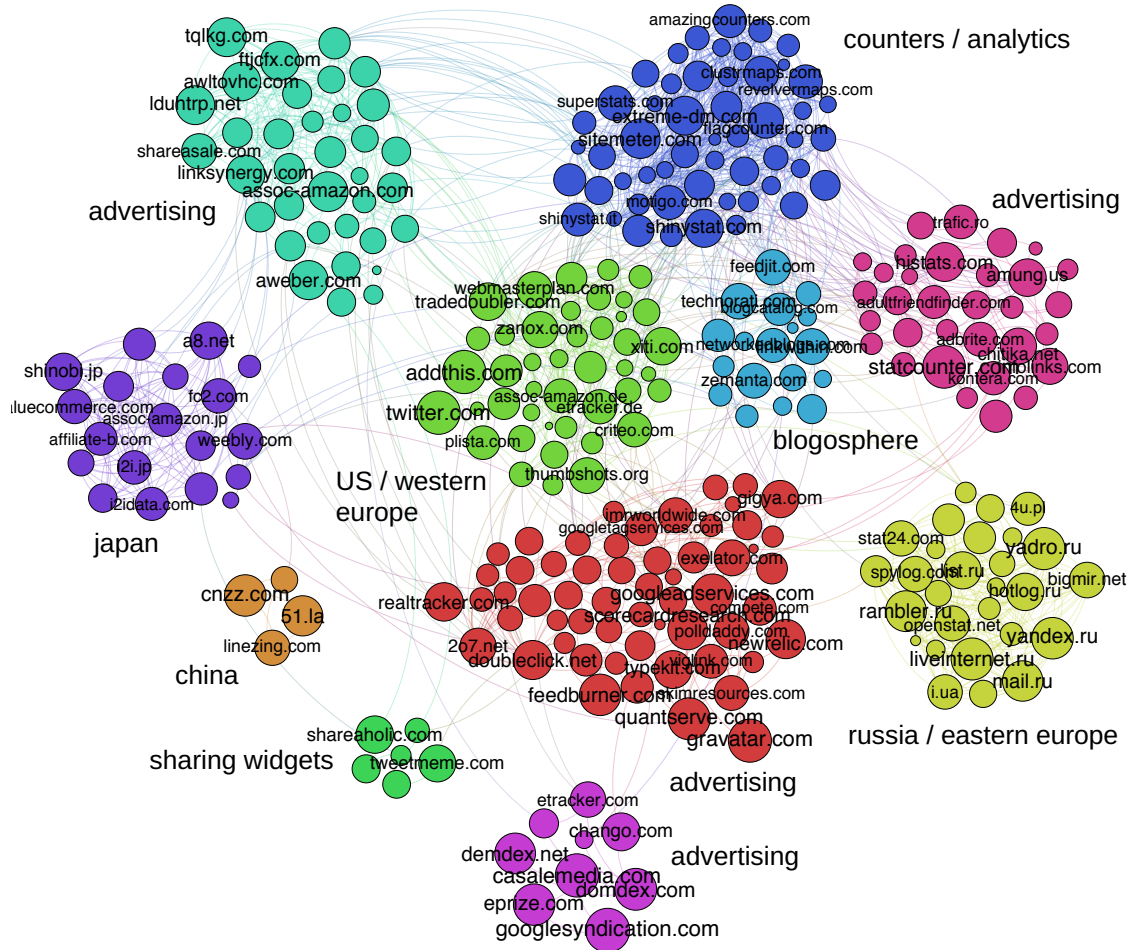


Figure 5: Cluster analysis of the co-occurrence network of tracking third-parties using modularity maximization. We manually derive the labels by investigating the distribution of third-party categories and registration countries inside the clusters. Vertex sizes are logarithmically proportional to the overall number of website PLDs on which a tracking service is present.

5.4 Tracker Co-Occurrences

We now focus on the relationships between individual pairs of tracking third-parties, in order to answer questions of the form *What proportion of the website domains that embed facebook.com also embed twitter.com?* We therefore compute the conditional probability $P(t_1 | t_2)$ of encountering a tracker t_1 on a website pay-level domain, given that it already has another tracker t_2 embedded.

Figure 4 illustrates the results for all pairs within the fifteen most predominant tracking services using a heat map. From this experiment, we make multiple observations. First, we observe the predominant role of `google-analytics.com`, having by far the highest probability of co-occurring with any other service, consistent with its wide spread over the web. Furthermore, we find a high agreement in the probability patterns for social media-related services such as `facebook.com`, `twitter.com` and `gravatar.com`.

Additionally, the heatmap hints at country-specific patterns, hence the isolation of Russian tracking third-parties such as `yadro.ru`, `liveinternet.ru` and `yandex.ru`.

5.5 Cluster Analysis of Tracker Co-Occurrences

We focus our analysis on the country-specific and media-specific patterns that the previous analysis pointed to in the data. We proceed in two steps to uncover these patterns: First, we compute the co-occurrence network of tracking services to find pairs of trackers that occur on the same sites very often. Second, we cluster this network using the Louvain method for community detection (Blondel *et al.*, 2008), and inspect the resulting clusters.

We compute the adjacency matrix of the co-occurrence network as the one-mode projection $A = B^T B$, with B being the bi-adjacency matrix of the tracking network (Kunegis, 2015) pointing from website PLDs to tracking PLDs. This means that we process each website PLD (corresponding to a row in B) and increment A_{ij} for every pair of tracking PLDs i and j present on the particular website. Therefore, an entry A_{ij} in the resulting co-occurrence matrix denotes the number of website PLDs on which the trackers i and j have both been present. We use the signed root-log-likelihood ratio (LLR) derived from log-likelihood ratio-based G^2 tests (Dunning, 1993) to prune the co-occurrence matrix. We choose a pruning thresh-

Netherlands. We attribute this to the fact that many Russian internet companies are legally based in the Netherlands for reasons of taxation (Vsevolod, 2014). Next, we encounter a cluster containing `shinobi.jp` which comprises mostly Japanese third-parties. Finally there is a small cluster which contains a majority of Chinese third-parties, such as `cnzz.com`.

Not all clusters lend themselves to a geographic interpretation however. The remaining clusters exhibit similarities in the type of trackers embedded or in the type of website tracked: We find a very large cluster of mostly web counter and web analytics services, with representatives such as `shinystat.com` and `amazingcounters.com`. Additionally, we find a small cluster of sharing widgets (e.g., `shareaholic.com` and `tweetmeme.com`). A third cluster, represented by `technorati.com` and `networkedblogs.com`, consists mostly of domains related to blogging. The remaining clusters mostly consist of advertising services.

6 Country-Specific Analysis of Tracking

In the previous sections, we have found that many clusters in the co-occurrence network exhibit a country-specific pattern. In order to better understand the relationships between individual countries and their distribution of web tracking, we conduct a set of deeper analyses on the country level. Every country is represented by a top-level domain such as `.de` for Germany and `.fr` for France. We omit the United States which in addition to its country code-based top-level domain `.us` also uses `.com`, which in turn is also widely used by non-US websites. We look at the 50 predominant top-level domains in our corpus and their corresponding countries belonging to many continents, ranging from Germany (`.de`) hosting 2.87 million pay-level domains to Thailand (`.th`) with 15,858 pay-level domains.

6.1 Predominant Tracking Companies per Country

We compute the rank share of tracking services in the subset of pay-level domains belonging to a specific country, represented by its top-level domain. We aggregate the results on company level in order to match a tracker to the country in which its owning company is based. Then, we compare the distribution of tracking companies in a specific country to the distribution of companies in the corpus as a whole, and to the distribution in the `.com` top-level domain. Figure 6 shows the results of this analysis for a selection of countries: Ireland, Germany, France, Poland, Israel, Turkey, Korea, China, Russia and Iran. We highlight the bars for the three globally most dominant companies: Google, Facebook and Twitter. These three companies have a special role, as we encounter them in the majority of top ten lists, in many cases accumulating the largest amount of rank share. Even among these three, Google has a preeminent position: we find it in a dominating role in the majority of countries, often with an amount of rank share that is more than double of what the second-placed company accumulates. We find that in many cases, the top ten companies consist of the three dominant US companies, Google, Facebook and Twitter, accompanied by a set of companies resident in the country under observation. Examples are Zanox (affiliate marketing) and

Table 2: Correlation of various country-level indicators with the tracking dominance of Google, Facebook and Twitter.

Variable	Correlation
Democracy index	0.662***
Freedom of the press	0.612***
Proportion of English speakers in the population	0.343*
Online ad spending per capita	0.333
US trade volume	0.167
Online ad spending ratio	0.062

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

INFOnline (digital audience measurement) in Germany, Criteo (advertising) in France, as well as Yandex and LiveInternet in Russia. These country-resident companies hardly ever appear in the top ten list of another country. The pattern of dominance of Google, followed by Facebook and Twitter is present in the overall corpus as well as in the vast majority of countries; however there are a few notable outliers, e.g., China and Russia where country-resident companies such as Yandex and CNZZ outrank Google.

6.2 Correlation Analysis of the Dominance of Google, Facebook and Twitter

We further investigate the country specifics of the three companies with a dominating role: Google, Facebook and Twitter. We therefore define a simple, dichotomous measure of dominance per country. We say that these three companies have a dominating role in a country if they accumulate more than half of the mass of the rank share distribution in the top ten companies. We compute this measure for the 50 countries under investigation in our corpus, and encounter the dominance pattern for a vast majority of 46 countries. Only four countries do not exhibit this pattern: China, Russia, Iran and Ukraine.

We would like to gain insights into factors which explain these findings. Intuitively, there could be influential factors of different character, e.g., political (e.g., states purposefully discriminate foreign companies to uphold control over media), socio-cultural (e.g., a language is dominant which is very different from English, making it harder for foreign companies to enter the market) or economical (e.g., certain countries might be more attractive to advertisers due to higher ad spending). We leverage a set of additional data sources for this analysis: an index describing the democratic character of a country (The Economist Intelligence Unit, 2012), another index measuring the degree of freedom of the press in a particular country (Freedom House, 2012), data about the ratio of English speakers in a country's population (Wikipedia, 2015), as well as population counts (The World Bank, 2015). Furthermore, we include economic data about US foreign trade (U.S. Census Bureau, 2015) and statistics about worldwide ad spending (CatchaDigital, 2013).

We compute several country-specific indicator variables from these additional datasets: Our political indicators consist of the *democracy index* and the *freedom of the press index*. For the

latter, we simply revert the existing scale to make high values indicate high freedom of the press. We use the *percentage of the population which speaks English* as socio-cultural indicator. Finally, we derive several economic indicators. We compute the *online ad spending per capita* as the sum of digital and mobile ad spending per country normalized by its population. The *online ad spending ratio* is the ratio of the sum of digital and mobile and spending to the overall media ad spending. Lastly, the *US trade volume* is the sum of imports and exports of the US with the given country, normalized by the size of the population of the country.

We calculate the point-biserial correlation coefficient of the indicator variables to our dichotomous dominance variable, in order to measure their strength of association. Table 2 shows the results. We find a very strong and at the same time statistically significant correlation with the political indicators, democracy index and freedom of the press. The socio-cultural indicator, amount of English speakers, is only moderately correlated and only statistically significant at the 0.05 level. The economic indicators, online ad spending per capita, US trade volume and online ad spending ratio show low to moderate correlation, which is not statistically significant. These findings are surprising as they indicate that a positive characteristic such as freedom of the press is accompanied by a potentially very negative characteristic: the recording of people’s browsing behavior by companies outside of the legal control of the countries institutions.

7 Content-Specific Analysis of Tracking

In this section, we investigate how the presence and amount of trackers in a website is related to its content. We focus on comparing domains with highly privacy-critical, intimate content to domains with less critical content. We want to know whether there is a reduction in tracking due to the intimate nature of the content. We leverage the DMOZ database (2016), a large, human-edited directory of the web which provides an extensive labeling of websites. We define four highly privacy-critical categories of websites, as well as four less privacy-critical categories in order to compare the distribution of trackers among them. We associate every such defined category with a set of DMOZ labels and extract the domains of all websites with that label. We repeat this for all available translations of the label in up to 41 different languages. Next, we intersect these domains with the domains in our corpus. We use the resulting set of labeled domains as a representative for the category.

This process results in 24,026 *health*-related domains, originating from DMOZ labels such as `Health/Mental_Health` and `Health/Medicine/Surgery` and 2,436 *addiction*-related domains, which have labels such as `Health/Addictions`, `Society/Sexuality/Sexual_Addiction`. Furthermore we find 3,421 *sexual-ity*-related domains, originating from labels such as `Society/Gay,_Lesbian,_and_Bisexual` and 611 domains related to *gender identity*, which we get by looking at labels such as `Society/Transgendered`. We chose these categories, as the association of a person with the topics discussed in these websites could potentially result in discrimination against the person (e.g., less chances during job search due to an illness), or even physical

Table 3: Number of highly privacy-critical and less privacy-critical websites with and without trackers. Overall, less privacy-critical websites are more likely to contain trackers than highly privacy-critical websites.

	Highly critical	Less critical
Without trackers	11,001	7,014
With trackers	16,697	57,850
Fraction with trackers	60.3%	89.2%

harm (e.g., prosecution of homosexual persons with prison or even the death penalty in some countries). As less privacy-critical categories, we choose *cooking & food* (40,698 domains), *soccer* (11,677 domains), *television* (5,793 domains) and *video games* (12,016 domains). Overall, this analysis includes 91,813 distinct labeled domains tagged with 450 labels (We provide a list of the labels used to the research community¹⁵).

7.1 Overall Avoidance of Trackers on Highly Privacy-Critical Websites

We perform a statistical test to find out whether in general, privacy-critical websites are more likely to contain trackers than privacy-noncritical websites. Table 3 summarises the total number of analysed websites in both categories with and without trackers. We observe that across all types of trackers combined, privacy-critical websites are less likely to contain trackers than privacy-noncritical websites. Indeed, about 90% of websites with less privacy-critical content contain trackers, while only about 60% of websites with highly privacy-critical content do. This result is significant to a p -value of $p < 0.001$. While this confirms that the choice of including trackers on a website seems to correlate with the topic of the website, indicating that trackers are avoided on possibly privacy-incriminating topics, the numbers still show that a majority of websites with highly privacy-critical content *do* contain trackers. This result also raises the question of *which* trackers are employed on which type of website, and whether some type of trackers are preponderant on any type of website, highly privacy-critical or less so. Therefore, the next experiment investigates each tracker separately.

7.2 Trackers Specifically Present on Highly Privacy-Critical Websites

In this experiment, we want to find out for each tracker whether it is represented more on privacy-critical or less-critical sites. We perform a statistical test based on the null hypothesis that the trackers are distributed randomly among all websites, and compare each tracker’s distribution to that. The results are shown in Table 4. We observe that the big three trackers (Google, Facebook, and Twitter) are less prevalent on highly privacy-critical websites. This is consistent with the overall results that websites containing highly privacy-critical

¹⁵<https://ssc.io/trackingthetrackers/dmoz.tsv>




Figure 7: Top ten companies embedded on highly privacy-critical (top) and less privacy-critical domains (bottom).

content avoid trackers. However, other trackers are more prevalent on highly privacy-critical websites than on less privacy-critical ones. These include AddThis, StatCounter, Amazon, Sharethis, Site Meter, Adobe, and many more. Many of these trackers are owned by companies much smaller than the three dominating ones (Google, Facebook, Twitter) and share a common characteristic: they are likely to not be perceived as *tracking* services, but as additional functionalities on a website, such as visitor counters or sharing buttons. We conjecture that the fact that these trackers are perceived as less of a threat to privacy leads them to not be avoided by hosters of highly privacy-critical websites. Nonetheless, it remains the case that even among websites with highly privacy-critical content, the most common trackers are Google, Facebook, AddThis and Twitter as shown in Figure 7. As largest three overall trackers are present in the top four spots (with only AddThis having a larger share of websites than Twitter), we specifically analyse the role of these top three trackers in the following experiments.

7.3 Distribution of Top Trackers in Privacy-Critical Websites

We compute the domain share and rank share of the three top trackers Google, Facebook and Twitter, as determined in Section 6.1, for all the domains in a particular category, and list the results in Table 5. Google has by far the highest domain-share in all categories, indicating that it is present on exceedingly more domains than any other company (cf. Table 5a). On average, we encounter Google services on 51% of domains per category, in contrast to services from Facebook and Twitter, which we appear on 26% and 14% of domains only. Surprisingly, the picture changes once we look at rank share instead of domain share, as shown in Table 5b: Google has less rank share than Facebook. That suggests that Facebook tracks less but more prominent domains (in terms of PageRank) for our defined categories. To further investigate this finding, we plot the cumulative distribution of ranks corresponding to tracked domains per company of the highly privacy-critical categories in log-log scale in Figure 8. Investigating this distribution confirms our hypothesis that on average, Google occurs on lower-ranking domains than Facebook and Twitter.

8 Conclusion and Discussion

The scope of our analysis allows us to make several novel observations about online tracking. We found that 9 out of the

20 predominant third-party domains belong to trackers, and confirmed the extraordinary tracking capability of Google Analytics (analogous to Roesner *et al.*, 2012; Krishnamurthy and Wills, 2006). Furthermore, we found that the distribution of the number of website domains tracked follows a power law, and that the overall tracking network is of a dissortative character. While there are many small trackers which are country-specific (e.g., to Germany and Japan), this is not true for the largest tracking services. These belong to Google, Facebook and Twitter, all US companies acting on a global scale, and representing the largest trackers in almost all countries. The exception to this pattern are a small number of countries such as China, Russia and Iran and which represent outliers in terms of political factors such as democracy and freedom of the press. In particular, we could not determine a statistically significant correlation with economic factors such as amount of foreign trade between a country and the US, or with indicators related to ad spending. Based on these findings, we suppose that the choice of tracking software made by website owners is largely independent of the website’s topic, and mostly depends on political factors, e.g., freedom of the press. In economical terms, this confirms that social media companies in US-friendly countries such as Germany and Japan have a hard time getting large market share due to the dominance of the US-based companies Google, Facebook and Twitter, while social media companies in countries such as Russia and China have better prospects due to the (voluntary or legislated) avoidance of US companies in those countries. Additionally, our findings support a previous study (The Guardian, 2015b) which concluded that Google still operates tracking services on Chinese websites, despite its proclaimed retreat from the Chinese market in 2010.

Our results indicate that the fact that a website covers highly privacy-critical topics does not imply a lack of tracking. Even though the rate of tracked websites among those with highly privacy-critical content is lower than for other websites (60% versus 90%), the majority of such websites does still contain trackers. One aspect of this high number is the apparent high number of trackers which are seemingly not perceived to be as dangerous as Google, Facebook and Twitter, namely those trackers which implement services benefitting the website itself, such as visitor counters. For such websites, our results indicate that they are even more prevalent on privacy-critical websites. For end users, we can conclude that tracking is to be expected on all types of websites, regardless of the topic.




Figure 8: Per-company cumulative distribution of the PageRank of tracked domains with highly privacy-critical content.

Table 4: Classification of trackers by the type of website (highly privacy-critical or less privacy-critical) they are more prevalent on. The table contains all trackers among the 20 most used trackers (by number of websites) which are significantly present for each type of website.

Prevalent on highly privacy-critical sites	Prevalent on less privacy-critical sites
AddThis ^{***} StatCounter ^{***} Amazon ^{***} Sharethis Inc ^{***}	Google ^{***} Facebook ^{***} Twitter ^{***} LiveInternet ^{***}
Site Meter ^{***} Adobe ^{***} extreme digital ^{***} AddToAny ^{***}	FullCircleStudies ^{***} Rakuten [*] Yandex ^{***} Rambler ^{***}
Disqus ^{**} ClustrMaps ^{***} ShinyStat ^{***} Yahoo [*] Boardhost ^{***}	Nielsen ^{***} AOL ^{***} Tradedoubler ^{***} FC2 ^{***} Samurai ^{***}
applied technologies ^{**} CommissionJunction [*] Microsoft ^{***}	Factory ^{***} mail.ru ^{***} New Relic ^{***} A8 ^{***} Chartbeat ^{***}
Histats [*] Technorati ^{***} Motigo ^{***} fav.or.it. [*]	ValueCommerce ^{**} VibrantMedia ^{***} INFOnline ^{***}

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Tracking capabilities of Google, Facebook and Twitter on different categories of domains.

(a) Domain share on highly privacy-critical categories and less privacy-critical categories of domains.

	Google	Facebook	Twitter
Health	0.473	0.157	0.081
Addiction	0.427	0.152	0.092
Sexuality	0.521	0.299	0.158
Gender Identity	0.462	0.295	0.201
Cooking	0.482	0.176	0.070
Soccer	0.519	0.301	0.160
Television	0.617	0.394	0.239
Video Games	0.590	0.268	0.144

(b) Rank share on highly privacy-critical categories and less privacy-critical categories of domains.

	Google	Facebook	Twitter
Health	0.903	0.912	0.855
Addiction	0.798	0.959	0.953
Sexuality	0.888	0.989	0.834
Gender Identity	0.802	0.990	0.914
Cooking	0.890	0.925	0.788
Soccer	0.904	0.981	0.840
Television	0.909	0.976	0.839
Video Games	0.916	0.965	0.834

Acknowledgments

We thank Felix Neutatz, Hung Chang, Sergey Dudoladov and Stephen Merity for their help in exploring the Common Crawl corpus. This work has been supported by an AWS in Education Research Grant award and by the EC FP7 grant n° 610928, REVEAL.

References

- Acar, G., C. Eubank, S. Englehardt, M. Juarez, A. Narayanan, and C. Diaz. 2014. “The Web never forgets: Persistent tracking mechanisms in the wild”. In: *ACM CCS*. 674–689.
- Bau, J., J. Mayer, H. Paskov, and J. C. Mitchell. 2013. “A promising direction for web tracking countermeasures”. In: *Proc. Web 2.0 Security and Privacy*.
- Blondel, V. D., J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. 2008. “Fast unfolding of communities in large networks”. *Journal of Statistical Mechanics: Theory and Experiment*. 2008(10): P10008.
- Broder, A., R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. 2000. “Graph structure in the Web”. *Computer networks*. 33(1): 309–320.
- CatchaDigital. 2013. “Worldwide ad spending forecast: Emerging markets, mobile provide opportunities for growth”. URL: <http://www.slideshare.net/catchadigital/emarketer-worldwideadspendingforecast>.
- U.S. Census Bureau. 2015. “U.S. trade in goods by country”. URL: <https://www.census.gov/foreign-trade/balance/country.xlsx>.
- Chaabane, A., M. A. Kaafar, and R. Boreli. 2012. “Big friend is watching you: Analyzing online social networks tracking capabilities”. In: *2012 ACM Workshop on Online Social Networks*. 7–12.
- Clauset, A., C. R. Shalizi, and M. E. Newman. 2009. “Power-law distributions in empirical data”. *SIAM*. 51(4): 661–703.
- Dean, J. and S. Ghemawat. 2010. “MapReduce: A flexible data processing tool”. *Communications of the ACM*. 53(1): 72–77.
- “Disconnect”. 2016. <https://disconnect.me/>.
- DMOZ. 2016. “Directory of the Web”. <https://www.dmoz.org/>.
- Dunning, T. 1993. “Accurate methods for the statistics of surprise and coincidence”. *Comput. Linguist.* 19(1): 61–74.
- Eckersley, P. 2010. “How unique is your web browser?” In: *Privacy Enhancing Technologies*. Springer. 1–18.
- The Economist Intelligence Unit. 2012. “DemocracyIndex”. URL: http://www.eiu.com/public/topical_report.aspx?campaignid=DemocracyIndex12.
- Electronic Frontier Foundation. 2015. “HealthCare.gov sends personal data to dozens of tracking websites”. URL: <https://www.eff.org/de/deeplinks/2015/01/healthcare.gov-sends-personal-data>.
- Englehardt, S. and A. Narayanan. 2016. “Online tracking: A 1-million-site measurement and analysis”. URL: http://randomwalker.info/publications/OpenWPM_1_million_site_tracking_measurement.pdf.

- Englehardt, S., D. Reisman, C. Eubank, P. Zimmerman, J. Mayer, A. Narayanan, and E. W. Felten. 2015. "Cookies that give you away: The surveillance implications of web tracking". In: *WWW*. 289–299.
- Freedom House. 2012. "Freedom of the press". URL: <https://freedomhouse.org/sites/default/files/FOTP%5C%202012%5C%20Final%5C%20Full%5C%20Report.pdf>.
- "Ghostery". 2016. <https://www.ghostery.com/>.
- Kalavri, V., J. Blackburn, M. Varvello, and K. Papaginannaki. 2016. "Like a pack of wolves: Community structure of web trackers". In: *PAMS*.
- Krishnamurthy, B. and C. E. Wills. 2006. "Generating a privacy footprint on the Internet". In: *ACM SIGCOMM*. 65–70.
- Krishnamurthy, B. and C. E. Wills. 2009a. "On the leakage of personally identifiable information via online social networks". In: *ACM workshop on Online Social Networks*. 7–12.
- Krishnamurthy, B. and C. E. Wills. 2009b. "Privacy diffusion on the Web: A longitudinal perspective". In: *WWW*. 541–550.
- Krishnamurthy, B. and C. E. Wills. 2010. "Privacy leakage in mobile online social networks". In: *USENIX Conference on Online social networks*. 4–4.
- Kunegis, J. 2013. "KONECT: The Koblenz Network Collection". In: *Proceedings of the 22nd International Conference on World Wide Web Companion*. International World Wide Web Conferences Steering Committee. 1343–1350.
- Kunegis, J. 2015. "Exploiting the structure of bipartite graphs for algebraic and spectral graph theory applications". *Internet Math*. 11(3): 201–321. DOI: 10.1080/15427951.2014.958250.
- Lehmberg, O., R. Meusel, and C. Bizer. 2014. "Graph structure in the Web: Aggregated by pay-level domain". In: *ACM Web Science*. 119–128.
- Libert, T. 2015a. "Exposing the hidden Web: An analysis of third-party HTTP requests on 1 million websites". *International Journal of Communication*.
- Libert, T. 2015b. "Privacy implications of health information seeking on the Web". *Communications of the ACM*. 58(3): 68–77.
- Mayer, J. R. and J. C. Mitchell. 2012. "Third-party web tracking: Policy and technology". In: *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE. 413–427.
- Meusel, R., S. Vigna, O. Lehmberg, and C. Bizer. 2014. "Graph structure in the Web—revisited: A trick of the heavy tail". In: *WWW*. 427–432.
- Meusel, R., S. Vigna, O. Lehmberg, and C. Bizer. 2015. "The graph structure in the Web—Analyzed on different aggregation levels". *The Journal of Web Science*. 1(1).
- Page, L., S. Brin, R. Motwani, and T. Winograd. 1999. "The PageRank citation ranking: Bringing order to the Web". *Tech. rep.*
- Perry, M., E. Clark, and S. Murdoch. 2015. "The design and implementation of the Tor browser [draft]". URL: <https://www.torproject.org/projects/torbrowser/design>.
- "Privacy Badger". 2016. <https://www.eff.org/de/node/73969>.
- Roesner, F., T. Kohno, and D. Wetherall. 2012. "Detecting and defending against third-party tracking on the Web". In: *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association. 12–12.
- Roosendaal, A. 2011. "Facebook tracks and traces everyone: Like this!" *Tilburg Law School Legal Studies Research Paper Series*. (03).
- Schelter, S. and J. Kunegis. 2016. "Tracking the trackers: A large-scale analysis of embedded web trackers". *ICWSM*.
- Spiegler, S. 2013. "Statistics of the Common Crawl Corpus 2012". *Tech. rep.* Technical report, SwiftKey.
- The Guardian. 2015a. "Belgian court orders Facebook to stop tracking non-members". URL: <http://www.theguardian.com/technology/2015/nov/10/belgian-court-orders-facebook-to-stop-tracking-non-members>.
- The Guardian. 2015b. "Google is returning to China? It never really left". URL: <https://www.theguardian.com/technology/2015/sep/21/google-is-returning-to-china-it-never-really-left>.
- The Intercept. 2015. "From radio to porn, British spies track web users' online identities". URL: <https://theintercept.com/2015/09/25/gchq-radio-porn-spies-track-web-users-online-identities/>.
- Trackography. 2014. "Meet the trackers; me and my shadow". URL: <https://myshadow.org/trackography-meet-the-trackers>.
- Vsevolod, P. 2014. "Russia beyond the headlines: Yandex reacts to Putin comments about foreign influence as share price falls". URL: https://rbth.com/business/2014/04/28/yandex_reacts_to_putin_comments_about_foreign_influence_as_share_pri_36283.html.
- "Web Data Commons – Hyperlink Graphs". <http://webdatacommons.org/hyperlinkgraph/>.
- Wikipedia. 2015. "List of countries by English-speaking population". URL: https://en.wikipedia.org/wiki/List_of_countries_by_English-speaking_population.
- The World Bank. 2015. "Population, total". URL: <http://data.worldbank.org/indicator/SP.POP.TOTL>.
- Yu, Z., S. Macbeth, K. Modi, and J. M. Pujol. 2016. "Tracking the trackers". In: *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee. 121–132.