Truth in Web Mining: Measuring the Profitability and the Imposed Overheads of Cryptojacking

Panagiotis Papadopoulos[1], Panagiotis Ilia[2], and Evangelos Markatos[1]

[1] University of Crete / FORTH, Greece, {panpap,markatos}@ics.forth.gr
[2] University of Illinois at Chicago, USA, pilia@uic.edu

Abstract. In recent years, we have been observing a new paradigm of attacks, the so-called cryptojacking attacks. Given the lower-risk/lower-effort nature of cryptojacking, the number of such incidents in 2018 were nearly double of those of ransomware attacks. Apart from the cryptojackers, web-cryptomining library providers also enabled benign publishers to use this mechanism as an alternative monetization schema for web in the era of declined ad revenues. In spite of the buzz raised around web-cryptomining, it is not yet known *what is the profitability of web-cryptomining and what is the actual cost it imposes on the user side.* In this paper, we respond to this exact question by measuring the overhead imposed to the user with regards to power consumption, resources utilization, network traffic, device temperature and user experience. We compare those overheads along with the profitability of web-cryptomining to the ones imposed by advertising to examine if web-cryptomining can become a viable alternative revenue stream for websites. Our results show that web-cryptomining can reach the profitability of advertising under specific circumstances, but users need to sustain a significant cost on their devices.

Keywords: Cryptomining, Cost of In-Browser Mining, Digital Advertising, Cryptojacking

1 Introduction

The last 15 months, we observe on the web the uncommon case of a new conceptual type of attack cropping up in cybersecurity. This new type of attack is the well-known cryptojacking, which had a 35% share of all web threats [19] last year. Specifically, when ransomware attacks declined by 45% in fourth quarter of 2018 in comparison to the first quarter of the same year, cryptojacking attack incidents quadrupled by 450% in the same time-frame [25].

Despite of the reduction in cryptocurrency prices, cryptojacking continues to be prevalent on the web [42] due to the minimal effort it requires from the attackers. The concept behind this category is pretty simple: When a user visits a website, their system resources get abused by the attacker to mine cryptocurrency for as long user has the tab open (and in some cases even for longer [44]). Malicious miners have shown up in mobile devices, cloud infrastructure, IoT gadgets and game consoles [36], or even critical infrastructure [37].

Of course, in-browser mining is not a new idea. The compatibility of JavaScript miners with all modern browsers gave motivation for web coin mining attempts since the very early days of Bitcoin, back in 2011 [35]. However, the increased mining difficulty of Bitcoin was the primary factor that led such approaches to failure. Yet, the emergence of altcoins with features like transaction speed, mining speed and distributed mining, became the growth factor for some coins (e.g., Monero [7] grew from 13$ to 300$ within 2017 [20]) and was the catalyst for the incarnation of JavaScript-based coin mining [29].

Coinhive [6], was the first JavaScript-based in-browser mining library (September 2017), which promoted web-cryptomining as an alternative revenue stream for publishers. And indeed, we observed a significant [12,1,15] number of benign content providers deploying mining libraries in their websites (e.g., The PirateBay [11]). Of course the increased growth of cryptomining did not create opportunities only for benign publishers, but cyber-attackers as well. Several incidents have been reported the last couple of years with popular and prestigious websites [30,34,18,32,22] being infected, thus forcing their visitors to mine cryptocoins. Although there are many existing works studying the prevalence or proposing detection mechanisms and countermeasures on cryptojacking [27,55,28], yet some interesting questions still remain unaddressed: *What is the actual cost of web-cryptomining on the user side? What is the profitability for the attacker or the benign publisher? Can it become an alternative web monetization scheme for benign publishers?*

In this study, we aim to tackle these exact questions; we conduct the first full-scale analysis of the profitability and costs of web-cryptominingon the user-side, in an attempt to shed light in this newly emerged technology, and explore if/to what degree it can replace ads on the web. Specifically, we estimate the possible revenues based on advertising and web-cryptomining, aiming to determine under what circumstances a miner-supported website can reach the profits from digital advertising. To achieve that, we collect a large dataset of ad- and miner-supported websites, and we develop WebTestbench: a sensor-based testbed to measure the resource utilization of both monetization models and compare the imposed user-side costs. In particular, with WebTestbench we measure (i) the utilization of system resources such as CPU and main memory, (ii) the degradation of the user experience due to the mining workload, (iii) the energy consumption and how this affects battery-operated devices (e.g., laptops, tablets, smartphones), (iv) system temperature and how overheating affects the user's device and (v) network and how this can affect a possible Internet data-plan.

To summarize, the main contributions of this paper are the following:

(i) We study the profitability of web-based cryptocurrency mining, while questioning its ability to become a reliable monetization method for web services. Our results show that for the average duration of a website visit, ads are at least 5.5× more profitable than web-cryptomining. However, a miner-supported website can produce higher revenues if the visitor keeps their tab open for longer than 5.53 minutes, when there are no more ads served. Based

on these findings, a hybrid approach that leverages both ad and cryptomining would allow publishers to receive the maximum possible profit.

(ii) We design a methodology to assess the resource utilization patterns of in-browser cryptominers on the visitor's device. We implement our approach into the WebTestbench framework[3] and we investigate what costs these utilization patterns impose on the visitor's side with regards to the user experience, and energy consumption and battery autonomy.

(iii) We collect a large dataset of around 200K ad- and miner- supported websites that include different web-cryptomining libraries. We use this dataset as input for the WebTestbench framework and we compare the resource utilization and costs of the two monetization models. Our results show that while browsing a miner-supported website, the visitor's CPU gets utilized 59× more than while visiting an ad-supported website, thus increasing the temperature (52.8%) and power consumption (2×) of their device.

2 Background

2.1 In-Browser Mining

Web-based mining is a method of cryptocurrency mining that happens inside a browser, using a script delivered from a website. The first attempts of in-browser Bitcoin (or Ethereum) mining failed due to the increased mining difficulty. However, the rise of alternative crypto-coins (altcoins) that provide distributed mining, increased mining speed and ASIC (Application-Specific Integrated Circuit) resistance, made distributed CPU (i.e, x86, x86-64, ARM) based mining very effective, even when using commodity hardware, and opened new potential funding avenues on the web. The large growth of web-cryptomining started with the release of Coinhive's JavaScript-based Monero miner [6] in September 2017[4]. After Coinhive, many more companies launched their own web mining services[56]. Such miners compute hashes as a Proof-of-Work, and can be easily included in any website, thus enabling publishers to utilize visiting users' CPUs as an alternative monetization mechanism.

Upon visiting a miner-supported website, the user receives a mining library along with the website's content. Usually these libraries are provided by third parties or mining service providers (MSP), who are responsible for (i) maintaining the source code, (ii) controlling the synchronization of computations, (iii) collecting the computed hashes and (iv) pay the publishers. Upon rendering, a miner establishes a persistent connection with the remote MSP (e.g., coinhive.com) to communicate with the service/mining pool. Through this channel, the miner receives periodically PoW tasks and reports back the computed hashes.

[3] Open-source: https://github.com/panpap/webTestbench

[4] Months after this paper's submission, Coinhive announced that it shuts down its operations after 2 years due to the hard fork and algorithm update of the Monero network [5].

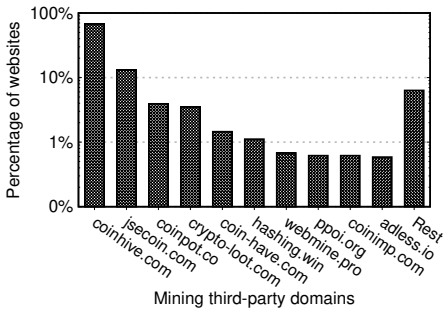

Fig. 1: Cryptomining market share per third party library in our dataset. Coinhive owns the dominant share (69%) when JSEcoin follows with 13%.

2.2 Cryptojacking

The so-called Drive-by Mining, or *cryptojacking*, takes place either (i) directly: when the publisher of a website performs web-cryptomining without the consent of their visitors[5], or (ii) indirectly: by compromising embedded third party libraries or by delivering malicious mining code through the ad ecosystem [21]. For example, the compromisation of a single screen reader third party (i.e., Browsealoud [43]) resulted in infecting more than 4000 websites that were using it. Victims of cryptojacking have been several popular and prestigious websites [30,34,18,32,22].

3 Data collection and analysis

To gather the necessary data for our study, we collect several miner blacklists [59] including the ones used by the 5 most popular mine-blocking browser extensions[6]. By merging these blacklists we compose a list of 3610 unique entries of mining libraries and keywords. Then, we use these entries to query PublicWWW archive [51], and we collect a total of 107511 mine-including domains. It should be noted that the domains we collected are ranked in the range from 1353 to 960540 in the Alexa rank of popular websites, and that the majority of them are based in the USA, Russia and Brazil.

The mining websites we collected, include more than 27 different third party miners, such as Coinhive, CryptoLoot, JSEcoin and CoinHave. In Fig. 1, we present the portion of websites in our dataset that use each one of these libraries. As seen, besides the large variety of mining libraries, there is a monopolistic phenomenon in the market of cryptominers, with Coinhive owning the dominant share (69%). From the rest of the libraries only JSEcoin miner surpasses 10%.

[5] Contrary to digital advertising where the visitors can discern ad-impressions, in web-cryptomining it is not easy to perceive the existence of a running miner.

[6] Coin-Blocker, No Mining, MinerBlock, noMiner and CoinBlock.

Table 1. Summary of our dataset

Type	Amount
Blacklist entries	3610
Miner-supported websites	107511
Ad-supported websites	100000
Unique third-party miners	27

Apart from miner-supported websites, we also collected an equal number of ad-supported ones, which are among the same popularity ranking range. Then, by using the blacklist of Ghostery adblocker, we enumerated all the ad-slots in these websites. We found the average number of ad-slots per website to be 3.4. Finally, Table 1 summarizes the contents of our dataset.

3.1 WebTestbench framework for utilization analysis

To measure the costs each domain in our dataset imposes on the user, we designed and developed WebTestbench: a web measuring testbed. A high-level overview of the architecture of WebTestbench is presented in Fig. 2. The WebTestbench framework follows an extensible modular design, and consists of several measuring components that work in a plug-and-play manner. Each such component is able to monitor utilization patterns across different system resources (e.g., memory, CPU, etc.). The main components of our platform currently include:

A. **crawler component**, which runs the browser (i.e., Google Chrome) in a headless mode. The crawling is responsible of stopping and purging *any* state after a website probe (e.g., set cookies, cache, local storage, registered service workers, etc.), and listening to the commands of the main controller (i.e., next visiting website, time interval, etc.).
B. **main controller**, which takes as input a list of domains and the visiting time per website. It is responsible for scheduling the execution of the monitoring components.
C. **monitoring platform**, which is responsible for the per time interval execution of the monitoring modules. This platform was build in order to be easily expandable in case of future additional modules.

For the scope of this analysis, we developed 6 different modules to measure the utilization that miners perform in 6 different system resources:

1. memory activity (physical and virtual), by using the `psrecord` utility [53] and attaching to the crawling browser tab's pid.
2. CPU utilization per core, by using the linux process status tool (`ps`).
3. system temperature (overall and per core), by leveraging the Linux monitoring sensors (`lm_sensors`[54]).
4. network traffic, by capturing (i) the network packets through `tcpdump` and (ii) the HTTP requests in the application layer along with their metadata (e.g., timing, initiator, transferred bytes, type, status code).

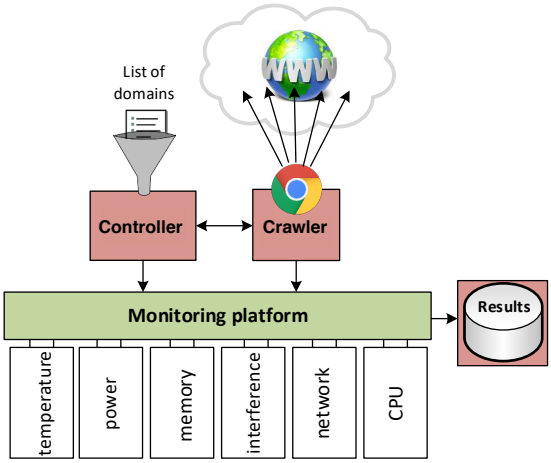

Fig. 2: High level overview of our measurement testbed. A Chrome-based platform fetches each website for a specific time, and its different components measure the resources.

5. process interference, to infer the degradation of user experience caused by the heavy CPU utilization of mining processes. This module consists of a CPU intensive benchmarking that includes multi-threaded MD5 hash calculations.
6. energy consumption, by utilizing an external Phidget21 power sensing package [50]. Phidget enable us to accurately measure the energy consumption of the 3 ATX power-supply lines (+12.0a, +12.0b +5.0, +3.3 Volts)[7]. The 12.0 Va line powers the processor, the 5.0V line powers the memory, and the 3.3V line powers the rest of the peripherals on the motherboard.

3.2 Methodology

In order to explore the different resource utilization patterns for miner- and ad-supported websites, we load our domain dataset in WebTestbench and we fetch each landing page for a certain amount of time. During this period the network monitoring module captures all outgoing HTTP(S) requests of the analyzed website. Additionally, the modules responsible for measuring the energy consumption, the CPU and memory utilization and the temperature, report the sensors' values in a per second interval. By the end of this first phase, WebTestbench erases any existing browser state and re-fetches the same website. This time, the only simultaneously running process is the interference measuring module which reports its progress at the end of the second phase.

[7] Instrumented in a similar way as in other studies [49].

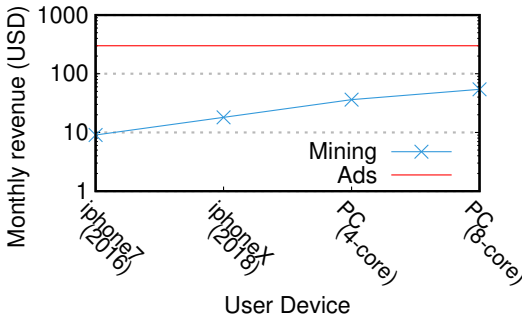

Fig. 3: Estimation of monthly profit for a website with 100K visitors and average visit duration of 1 minute. Visitors with mid-rage 8-core PCs (300H/s) provide 5.5× less revenue to the publisher than when having 3 ads in the website.

4 Profitability and Cost Analysis

In this section, we explore the profitability of cryptomining for the owner of the mining module (i.e., publisher or cryptojacker) and the cost this imposes on the users. For the following experiments, we use a Linux desktop equipped with a Hyper-Threading Quad-core Intel I7-3770 operating at 3.90 GHz, with 8 MB SmartCache, 8 GB RAM and an Intel 82567 1GbE network interface.

4.1 Profitability of In-Browser Mining

In the first set of experiments, we explore the profitability of in-browser miners and compare it to the current digital advertising model. Thereby, in the first experiment we simulate the monthly profit of the two strategies for a website of moderate popularity: 100,000 visitors per month. Studies have measured the average duration of a website visit being around 1 minute [38].

For this experiment, we use the popular Monero mining library of Coinhive and in order to measure the highest possible revenues for the miner's owner we assume the maximum rate this library provided ever: 0.0001468 XMR/1M hashes. This means that the miner's owner would get 0.0001468 Monero (XMR) per 1 million calculated hashes. Given the volatility of crypto-coin values against fiat money, in our simulations, we use the mean value of Monero across last year (2018): 1 Monero=205 USD. Apart from the visit duration, the amount of total calculated hashes of a website depends on the computation power of the visitors' devices. Thus, in this experiment, in order to cover a wider range of CPU hashrate capabilities, we use 4 different levels of computation rates based on [10]: the rate of 50 Hash/sec(e.g., iPhone7), 90 Hash/sec (e.g., iPhoneX), 200 Hash/sec (e.g., 4-core PC) and 300 Hash/sec (e.g., 8-core PC).

In the same experiment, we also compute the monthly revenue in case of a benign publisher when leverages instead of web-cryptomining,the advertising

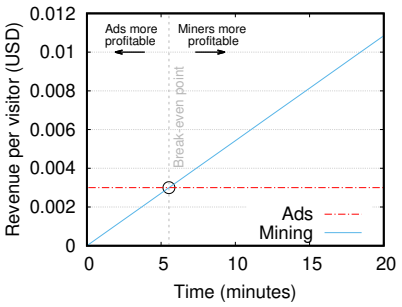

Fig. 4: Revenue per visitor. A publisher makes higher profit from mining than using ads (3 ad-slots) when the visitor (mid-rage 8-core PC) has his tab open for duration > 5.53 minutes.

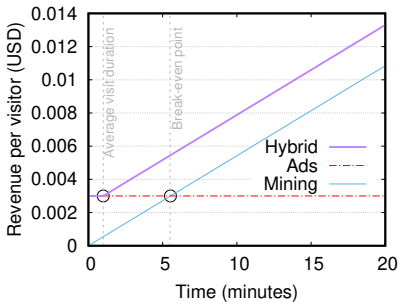

Fig. 5: Revenue per visitor according to a hybrid approach. The revenue is bounded either before or after the break-even point, to be always higher or equal to both ads and web-cryptomining

model to monetize their content. The most popular medium for personalized ad-buying nowadays [14] is the programmatic instantaneous auctions. The average number of ad-slots in an ad-supported website is 3 and the median charge price per ad impression as measured in previous studies [47,46] is 1 CPM[8] (Cost per thousand of impressions).

As can be seen in Fig. 3, for the average duration of a user's visit, the miner achieves an average computation rate from visitors of as high as 300Hash/sec. We see however that the website produces $5.5\times$ more revenue when using ads instead of web-cryptomining[9]. In addition, we see that as the visitor's hardware improves, the distance between these two monetization methods becomes smaller. This means that in the future web mining can be capable of providing comparable profits for the publishers.

It is apparent, that for a miner-supported website, *time matters*. Indeed, recent studies [1] show that the majority of miner-supported websites provide content that can keep the visitor on the website for a long time. Such content includes video or movie streaming, flash games, etc. Of course, in web-cryptomining, the user does not need to interact with the website's content per se. As such, numerous deceiving methods (e.g., via service workers [44,13] or pop-unders [57]) are currently used, aiming to allow the embedded miner to work in the background for as long as possible.

In the next experiment, we set out to identify the minimum time the website needs to remain open inside a tab of a visitor's browser in order to make higher profits than when using ads. In Fig. 4 we simulate the revenue per visitor for a website running in the background, and we use the same hash-rate levels as above. In order to produce revenues higher than when ads are delivered, the

[8] Advertiser pays 1 US dollar every 1000 successfully rendered impressions.
[9] Our simulation results have been also verified by real world experiments [8].

Table 2. Distribution of the average CPU Utilization for the different monetization methods. The median miner-supported website utilizes 59× more the user's CPU than the median ad-supported website.

Type	10^{th} **Perc.**	**Median**	90^{th} **Perc.**
Advertising	3.33%	9.71%	17.19%
Mining	560.11%	574.01%	580.71%

website must remain open in the user's browser for a duration longer than 5.53 minutes. When on background, the website does not receive fresh ads, since no new ad-auctions take place. In Fig. 4 we see that it is more profitable to use ads on the left of the break-even point, but when moving on the right of break-even point, web-cryptomining generate higher income.

To mitigate that, one could use a hybrid model to combine both. Specifically, as shown in Fig. 5, websites can utilize ads to generate a basic revenue from the visitor and move to web-cryptomining when she switches to different browser tab (e.g., after 1 minute). This way, publishers can continue making profit when their websites become idle. So a publisher's revenue when using ads is given by R_A (1), when using web-cryptomining by R_M (2), and when using the hybrid model by R_H (3), where t_0 is the average duration of a visit.

$$for\ t \in (0, t_0):\ R_A(t) = C_1 \tag{1}$$

$$for\ t \in [t_0, \infty):\ R_M(t) = C_2 * t \tag{2}$$

$$R_H(t) = C_1 + C_2 * (t - t_0) \tag{3}$$

As can be seen in Fig. 5, the revenue produced by the hybrid approach is always higher or equal to both ads and web-cryptomining.

4.2 Costs on the User Experience

After estimating the possible revenues of cryptomining, it is time to measure the costs this method imposes on the user and see how do these compare to the costs a website imposes when it uses ads to monetize its content.

CPU and Memory Utilization: In this set of experiments, we explore the average CPU and memory utilization by mining-supported websites. Note at this point, that the intense of mining is tunable. The majority of mining libraries enable their controller to fine tune the number of threads and the throttling of the miner. In this experiment we fetch each website in our two subsets for 3 minutes using WebTestbench and we extract the distribution of its CPU utilization through time. In Table 2 we report the average values for the median, the 10th and 90th percentiles. As we see, the median miner-supported website utilizes the visitor's CPU up to 59× more than an ad-supported website. We also measure the utilization of the visitor's main memory. In Fig. 6 we plot the

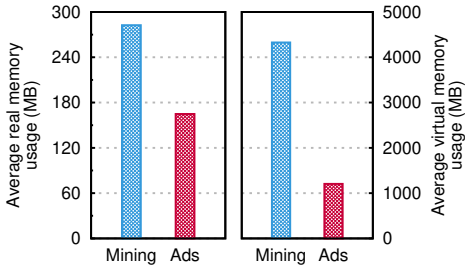

Fig. 6: Distribution of average real and virtual memory utilization through time. Miner-supported websites although reserve (3.59×) larger chunks of virtual memory, require 1.7× more MBytes of real memory than ad-supported websites.

average values for both real and virtual memory activity. As expected, miners do not utilize memory as heavy as CPU. On average the miner-supported websites require 1.7× more space in real memory than the ad-supported websites.

Network Activity: Next, we measure the network utilization of the average mining-supported website. As discussed in Section 2, a mining library needs to periodically communicate with a remote third party server (i.e., the MSP's server) in order to report the calculated hashes but also to obtain the next PoW. This communication in the vast majority of the libraries in our dataset takes place through a special persistent channel that allows bidirectional communication. To assess the network activity of web miners, we use the network capturing module of WebTestbench and we monitor the traffic of each (ad- and miner-supported) website for 3 minutes.

Based on the used third-party mining library, we isolate the web socket communication between its in-browser mining module and the remote MSP server. In order to compare this PoW-related communication of miners with the corresponding ad-related traffic of ad-supported websites, we utilize the open-source blacklist of the Disconnect browser extension to isolate all advertising related content. In Fig. 7, we plot the distribution of the total transmitted volume of bytes per website for the visit duration of 3 minutes. Although the web socket communication of miners consists of small packets of 186 Bytes on average, we see that in total the median PoW-related communication of miner-supported websites transmitted 22.8 KBytes, when the median ad-traffic volume[10] of ad-supported websites was 6.7 KBytes. This means that the median miner-generated traffic volume is 3.4× larger than the median ad-generated. In this experiment, we see that the network utilization patterns depend not only on the throttling of the miner but also on the different implementations. For example, while using the same portion of CPU, the miner of coinhive.com transmits on average 0.6

[10] Traffic that ad-related domains transmit to the user's browser (includes: impressions, ad/tracking related scripts, etc.).

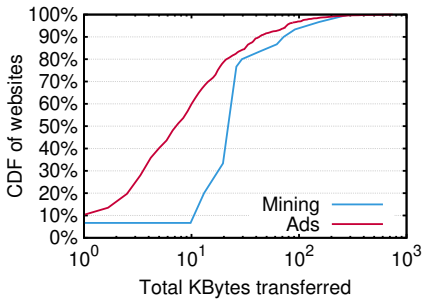

Fig. 7: Distribution of the transmitted volume of bytes per website for a visit duration of 3 minutes. The median miner-generated traffic volume is 3.4× larger than the median ad-generated.

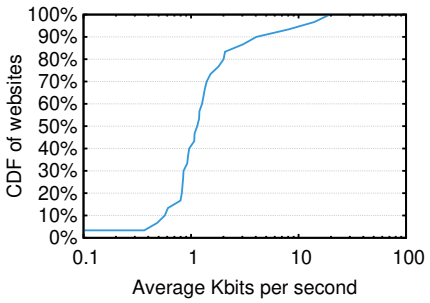

Fig. 8: Distribution of transmitted bit rate per miner-supported website in our dataset. The median in-browser miner communicates with its MSP by transmitting 1,168 bits per second.

packets/sec, webmine.cz: 2.2 packets/sec, cryptoloot.com: 4.7 packets/sec and jsecoin.com: 1.3 packets/sec.

In Fig. 8, we plot the distribution of the average data transfer rate per miner-supported website in our dataset. As shown, for the median case, the communication between the miner and the MSP has a transfer rate of 1 Kbit per second (or 142 Bytes/sec). As in the previous experiment, the rate highly depends on the mining library, with some of them reaching up to 14 Kbit per second. At this point, recall that the PoW-related communication between the miner and the MSP holds for as long as the miner is running, and as shown in Fig. 4, a miner must run for longer than 5.53 minutes to produce revenues higher than ads. This means that for the median case, the total volume of transferred bytes will exceed 46 KBytes.

In the case of a user who use consumption-based Internet service pricing[11] the monetary cost imposed is 0.000219$ per minute on average when browsing a miner-supported website. On the other hand, a publisher who includes a coinhive miner in its website earns 0.000409$ per minute from that user (considering that the user provides a hash rate of 227 Hash/sec as in [8]). Hence, we see that users with metered connections, when they visit miner-related websites, pay a monetization cost that is 53% less than the revenue of the publisher.

Power Efficiency: Of course the intensive resource utilization of cryptominers affects also the power consumption of the visitor's device, which has a direct impact on its battery autonomy. In the next experiment, we measure the power consumed by (i) main memory and (ii) CPU and network adapter components of the user's device while visiting miner- and ad- supported websites for a 3 minute duration. In Table 3, we report the average median, 10th and 90th

[11] Considering the average prices per byte in USA and Europe [16,58,2]

Table 3. Distribution of the average consumption of power for the different monetization methods. The median miner-supported website forces the user's device to consume more power than the median ad-supported website: 2.08× and 1.14× more power for the CPU and the memory component, respectively.

Component	Type	10^{th} Percentile	Median	90^{th} Percentile
CPU & Network adapter	Advertising	31.88 Watt	32.39 Watt	34.17 Watt
	Mining	63.35 Watt	67.60 Watt	71.22 Watt
Main Memory	Advertising	4.37 Watt	4.46 Watt	5.35 Watt
	Mining	4.76 Watt	4.99 Watt	5.67 Watt

percentile values for all websites in our dataset. As shown, there is a slightly increased (1.14× more than ad-supported websites) consumption of the memory component in miner-supported websites. However, we see that the heavy computation load of cryptominers significantly increases the CPUs and network adapters consumption, making miner-supported websites consume 2.08× more energy than ad-supported websites! This means that a laptop able to support 7 hours of consecutive traditional ad supported browsing, would support 3.36 hours of mining-supported browsing.

System Temperature: The increased electricity powering of the visitor's system results to an increased thermal radiation. During the above experiment, we measure the distribution of the per-core temperatures while visiting each website in our dataset for 3 minutes. In Fig. 9 we present the average results for the percentiles: 10th, 25th, 50th, 75th, 90th. As we can observe, the core temperatures for miner-supported websites are constantly above the optimal range of $45-60°$ Celsius [33]. In particular, the visitor's system operates for most of the time in the range of $43-50°$ Celsius while visiting ad-supported websites. When the visited website includes a miner, the average temperature of the cores reaches up to 52.8% higher, in the range of $73-77°$ Celsius, when in 10% of the cases it may reach higher than 84° Celsius.

To that end, with regards to the costs imposed to the user, high temperatures may lead to degraded system performance and poor user experience. Apart from that, constantly running a commodity device (e.g., mobile phone, laptop or desktop PC) at high temperatures, without a proper cooling mechanism, may significantly decrease the hardware's lifespan in the long term or even cause physical damage by thermal expansion.

Impact on Multi-tab Browsing and Parallel Processes: The heavy utilization of the visitor's CPU can affect the overall user's experience not only in the visited website, but in parallel processes and browser tabs, too. Indeed, for as long as the browser tab of a mining-supported website is open, the multi-threaded computations of the miner leaves limited processing power for the rest of the running applications. To make matters worse, as part of a PC's own cool-

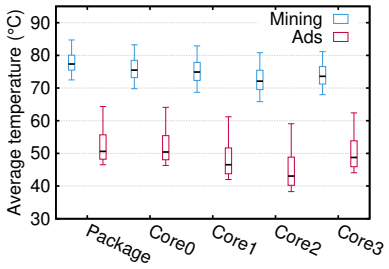

Fig. 9: Distribution of average temperatures per system's core. When the visited site has miner, the average temperature of the cores may reach up to 52.8% higher (73-77° Celsius) than when with ads.

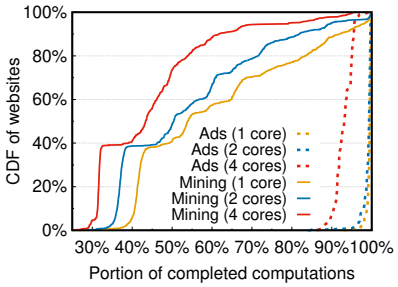

Fig. 10: Impact of background running miner- and ad-supported websites to a user's process. When the majority of ad-supported websites have negligible effect in other processes, the median embedded miner causes a performance degradation of up to 57% to a parallel running process.

ing system, the motherboard, in case of increased temperatures, may instruct the CPU component to slow down, or even force the system to shut down [4].

To assess how these factors may affect parallel running processes in the visitor's device, we use the interference measuring module of WebTestbench and we measure the performance overhead caused by background running miners. This module introduces computation workloads to the system to emulate a parallel running process of the user. Specifically, WebTestbench fetches each website in our dataset for the average visit duration (i.e., 1 minute), in parallel conducts multi-threaded MD5 hash calculations, and in the end reports the number of calculated hashes. To test the performance of parallel processing under different computation workloads, we visit each website using 3 setups for the MD5 process, utilizing in parallel 1, 2, and 4 cores of the CPU. In addition, we run the MD5 process alone for 1 minute to measure the maximum completed operations.

In Fig. 10, we plot the distribution of completed operations per website. As expected, when there is a miner-supported website running in the user's browser, the performance of the user's processes that run in parallel is severely affected. In particular, we see that the median miner-supported website forces the parallel process (depending on its computation intensity) to run in 54%, 50% or even 43% of its optimal performance, thus causing a performance degradation of 46% to 57%! Additionally, we see a 39% of miners greedily utilizing all the system's CPU resources causing a performance reduction of 67% to the parallel process.

Additionally, we measure the interference that ad-supported websites introduce to the parallel processes. As expected, the impact is minimal and practically only processes with full CPU utilization are affected, facing performance degradation of less than 10% for the majority of websites. This is mainly the result of JavaScript code responsible for ad serving, user tracking, analytics, etc.

Such severe performance degradation when the user is visiting a mining-supported website can cause glitches, or even crushes to parallel applications (like movie playback, video games), thus ravaging the user's experience. Of course, this performance degradation does not only affect parallel running applications but also mining operations from other open browser tabs. Indeed, a miner can achieve full utilization when the user has visited the miner-supported *website*1. However, when the user opens a second miner-supported *website*2, the utilization for both, as well as the revenues for *publisher*1 and *publisher*2, drop to a half. It is easy thus to anticipate, that **the scalability of cryptomining is limited since the more websites rely on web-cryptomining for funding, the less revenues will be generated for their publishers.** While this monetization model has that apparent drawback, in digital advertising each ad-supported website is totally independent from any parallel open browser tabs.

5 Discussion

User awareness: The recent years, directives like GDPR aim to bring transparency on the web. Website owners need to request user's consent before placing any cookies on the user side but also clearly present in a privacy policy statement what they do and with whom they share the user's data.

The lack of similar policies and directives regarding the proper use of cryptomining has raised a big controversy regarding the lack of transparency in miner-supported websites [31,9]. Many miner-supported websites do not inform the user about the existence of a miner, neither ask for the visitor's consent to utilize their system's resources for cryptocurrency mining.

In one of the first law cases about web-based mining, the Attorney General John J. Hoffman stated that "*no website should tap into a person's computer processing power without clearly notifying the person and giving them the chance to opt out*" [26]. As a result, whenever a user visits a website and she is not aware about the background web-cryptomining, irrespectively whether the mining code has been legitimately deployed by the publisher or a malicious actor that hijacked the website, this is considered as a *cryptojacking* attempt.

Letting the users choose: Since both digital advertising [23,46] and web-cryptomining impose costs on the user, a new paradigm could be to inform the user about these costs and give them the option to choose which of the two monetization schemes is more suitable for them (as already happens with paywalls [48,17]). In the case of advertising, the costs are associated with the network bandwidth and the privacy implications of targeted advertising [41,45,3], while the cost of web-cryptomining is associated with higher energy consumption (and battery drainage, overheating, etc.). A viable option for publishers would be to inform the users about these costs, and provide two different versions of their website (i.e., one that serve ads and one that uses cryptoming), thus allowing the user to choose between the two schemes. Indeed, such examples have already tested by various publishers [52].

Of course there are users using ad-blockers who may also deploy mine-blockers (e.g., Coin-Blocker, No Mining, CoinBlock, etc.) to avoid both approaches. However, publishers will eventually deny access to such users (similar to what already happened with ad-blocking users [39]) in an attempt to mitigate declining revenues.

6 Related Work

Eskandari et al. [15], in one of the first web mining related studies, analyzes the existing in-browser mining approaches. In particular, they measured the growth of cryptomining by looking for mining libraries in Internet archive services. In addition, they collected a set of 33K websites by querying for popular mining projects the Censys.io BigQuery dataset, and they studied the CPU utilization of the included miners. In [1] authors analyze the top 100,000 websites for cryptocurrency mining scripts in an attempt to measure the adoption of cryptominers in the web. The analysis revealed 220 of these websites using cryptomining scripts with their aggregated audience being around 500 million people. The content of these hosting websites were usually content that could keep the user on the website for long, and specifically movie/video/tv streaming (22.27%), file sharing (17.73%), Adult (10%) and News & Media (7.73%).

In [27], authors propose OUTGUARD: a system for automated cryptojacking detection. This system uses an SVN classification model which uses as input features like CPU usage, page execution time, iframe source loads, etc., and the accuracy it achieves reaches 97.9% TPR and 1.1% FPR. Authors run a prototype of their system across Alexa Top 1M sites and they found 3,600 new cases of cryptojacking. Konoth et al. in [28] analyze Alexa Top 1 Million websites studying the prevalence and profitability of cryptojacking (or drive-by mining). Authors identified 20 active cryptomining campaigns and evaluate current blacklist-based countermeasures. The ineffectiveness of these approaches motivated authors to propose *MineSweeper*: a detection approach that leverages cryptographic functions identified through static analysis and monitoring of cache events during run time. Similarly, in [55], authors study the prevalence (measured as low as 0.08%) of cryptomining as a monetization model for publishers. They identify and classify mining websites in the three largest TLDs and the Alexa Top 1M (over 138M domains). They also discuss the inadequacy of block lists and present a WebAssembly-based fingerprinting method to identify miners. Their approach was able to identify 82% more mining websites than NoCoin filter list.

In [24], authors leverage inherent characteristics of cryptojacking scripts to build *CMTracker*: a behavior based detector to automatically track cryptocurrency web mining scripts and their related domains. They discovered 2770 unique cryptojacking samples in a dataset with Alexa top 100K sites. They estimate the cost of these miners to be more than 278K kWh extra power per day, and earnings of at least 59K USD/day for the attackers. Finally, they study the evasiveness of cryptojacking scripts with 26% of them use code obfuscation. They tested the

detectability of these obfuscation-using scripts using VirusTotal and they saw that only 28% of them could be detected by at least one anti-virus engine.

7 Summary And Conclusion

In this paper, we measure the costs cryptominers impose on the user side by analyzing the utilization patterns of miner-deploying websites on the visitor's system resources (e.g., CPU, memory, network). We study the impact of these utilization patterns (i) on the visitor's device by measuring the system's power consumption and temperature, and also (ii) on the visitor's experience while running other applications in parallel. As a next step, we investigate the ability of web-cryptomining to become a reliable alternative monetization model for benign web publishers. The findings of our analysis can be summarized as follows:

- for the average duration of a web visit, website generates more than 5.5× higher revenues by including 3 ad impressions than by including a cryptominer.
- to produce higher revenues with a miner than with ads, the user's browser tab must remain open for a duration longer than 5.53 minutes or use a hybrid approach.
- the median miner-deploying website utilizes up to 59× more of the visitor's CPU and require 1.7× more space in real memory than ad-supported websites.
- the transfer rate of the median miner-MSP communication is 1 Kbit/sec. For a user with a metered Internet connection, the monetary cost imposed is on average 0.000219$ per minute, when the publisher from the same user earns 0.000409$ per minute.
- the median miner-generated traffic volume is 3.4× larger than the corresponding ad-generated.
- a visit to an average miner-deploying website consumes on average 2.08× more energy than to an average ad-supported website.
- a visitor's system operates in up to 52.8% higher temperatures when visiting a website with miner than when with ads.
- In-browser miners severely affect parallel running processes. The median miner-deploying website when running in the background may degrade even 57% of the performance of parallel running applications.

After completing our analysis, we see that web-cryptomining can indeed constitute an alternative monetization mechanism for specific categories of benign publishers after obtaining consent from the visitors. In these days, where EU regulators [40] aim to reform the way user data are being collected and processed for targeted advertising, cryptomining provides a privacy-preserving monetization model that requires zero data from the users. However, this study shows that the intensive resource utilization of web-cryptomining libraries imposes a significant cost on the user's device, thus accelerating the deterioration of its

hardware. To make matters worse, this *heavy utilization also limits the scalability and profitability of web-cryptomining, since the more websites adopt miners the less portion of resources each of them can acquire from a user that keeps multiple tabs open*. By using a hybrid approach, publishers could increase their profits by monetizing with ads at the beginning and then switch to mining to leverage the time when their websites reside in idle tabs.

Acknowledgements: The research leading to these results has received funding from European Union's Marie Sklodowska-Curie grant agreement 690972 (PROTASIS) and Horizon 2020 Research & Innovation Programme under grant agreement 786669 (REACT). This work has been also supported by the EU H2020-SU-ICT-03-2018 Project No. 830929 CyberSec4Europe (cybersec4europe.eu). The paper reflects only the authors' view and the Agency and the Commission are not responsible for any use that may be made of the information it contains.

References

1. AdGuard Research. Cryptocurrency mining affects over 500 million people. and they have no idea it is happening. https://adguard.com/en/blog/crypto-mining-fever/, 2017.
2. AT&T. Create your mobile share advantage plan. https://www.att.com/shop/wireless/data-plans.html, 2018.
3. Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. Tracing information flows between ad exchanges using retargeted ads. In *Proceedings of 2016 USENIX Security Conference*, SEC'16, 2016.
4. Philip Bates. How heat affects your computer, and should you be worried? https://www.makeuseof.com/tag/how-heat-affects-your-computer-and-should-you-be-worried/.
5. Catalin Cimpanu. Coinhive cryptojacking service to shut down in march 2019. https://www.zdnet.com/article/coinhive-cryptojacking-service-to-shut-down-in-march-2019/, 2019.
6. Coinhive. Monetize your business with your users' cpu power. https://coinhive.com/#javascript-api.
7. CoinWarz. Monero network hashrate chart and graph. https://www.coinwarz.com/network-hashrate-charts/monero-network-hashrate-chart.
8. Maxence Cornet. Coinhive review: Embeddable javascript crypto miner - 3 days in. https://medium.com/@MaxenceCornet/coinhive-review-embeddable-javascript-crypto-miner-806f7024cde8, 2017.
9. Cryptocurrency Posters. The ethics of javascript mining. https://www.cryptocurrencyposters.com/the-ethics-of-javascript-mining/, 2019.
10. CryptoMining24.net. Cpu for monero. https://cryptomining24.net/cpu-for-monero/, 2017.
11. Ernesto Van der Sar. The pirate bay website runs a cryptocurrency miner (updated). https://torrentfreak.com/the-pirate-bay-website-runs-a-cryptocurrency-miner-170916/.
12. Deepen Desai, Dhruval Gandhi, Mohd Sadique, and Manohar Ghule. Cryptomining is here to stay in the enterprise. https://www.zscaler.com/blogs/research/cryptomining-here-stay-enterprise.

13. Brannon Dorsey. Browser as botnet, or the coming war on your web browser. Radical Networks., 2018.
14. eMarketer Podcast. emarketer releases new us programmatic ad spending figures. https://www.emarketer.com/Article/eMarketer-Releases-New-US-Programmatic-Ad-Spending-Figures/1016698, 2017.
15. Shayan Eskandari, Andreas Leoutsarakos, Troy Mursch, and Jeremy Clark. A first look at browser-based cryptojacking. In *Proceedings of IEEE S&B'18*, S&B'18, 2018.
16. FANDOM Lifestyle Community. Prepaid data sim card wiki - spain. http://prepaid-data-sim-card.wikia.com/wiki/Spain, 2017.
17. Frederic Filloux. Paid or ad-supported: Pick one model, and stick with it. https://mondaynote.com/paid-or-ad-supported-pick-one-model-and-stick-with-it-a0c7f8db8db8, 2016.
18. Brian Fung. Hackers have turned politifact's website into a trap for your pc. https://www.washingtonpost.com/news/the-switch/wp/2017/10/13/hackers-have-turned-politifacts-website-into-a-trap-for-your-pc/, 2017.
19. Edward Gately. Webroot: Beware the worst malware of 2018. https://www.channelpartnersonline.com/2018/10/30/webroot-beware-of-the-worst-malware-of-2018/.
20. Global Coin Report. Here's how monero (xmr) gets to $1,000. https://globalcoinreport.com/heres-monero-xmr-gets-1000/, 2018.
21. Dan Goodin. Ad network uses advanced malware technique to conceal cpu-draining mining ads. https://arstechnica.com/information-technology/2018/02/ad-network-uses-advanced-malware-technique-to-conceal-cpu-draining-mining-ads/.
22. Patrick Greenfield. Government websites hit by cryptocurrency mining malware. https://www.theguardian.com/technology/2018/ feb/11/government-websites-hit-by-cryptocurrency-mining-malware, 2018.
23. Jiaping Gui, Stuart Mcilroy, Meiyappan Nagappan, and William G. J. Halfond. Truth in advertising: The hidden cost of mobile ads for software developers. In *Proceedings of the 37th International Conference on Software Engineering*, ICSE '15, 2015.
24. Geng Hong, Zhemin Yang, Sen Yang, Lei Zhang, Yuhong Nan, Zhibo Zhang, Min Yang, Yuan Zhang, Zhiyun Qian, and Haixin Duan. How you get shot in the back: A systematical study about cryptojacking in the real world. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1701–1713. ACM, 2018.
25. IBM Security. Ibm x-force report: Ransomware doesn't pay in 2018 as cybercriminals turn to cryptojacking for profit. https://newsroom.ibm.com/2019-02-26-IBM-X-Force-Report-Ransomware-Doesnt-Pay-in-2018-as-Cybercriminals-Turn-to-Cryptojacking-for-Profit.
26. Steve C. Lee John Hoffman, Jeffrey S. Jacobson. New jersey division of consumer affairs obtains settlement with developer of bitcoin-mining software found to have accessed new jersey computers without users' knowledge or consent. http://nj.gov/oag/newsreleases15/ pr20150526b.html, 2015.
27. Amin Kharraz, Zane Ma, Paul Murley, Charles Lever, Joshua Mason, Andrew Miller, Nikita Borisov, Manos Antonakakis, and Michael Bailey. Outguard: Detecting in-browser covert cryptocurrency mining in the wild. In *Proceedings of 2019 World Wide Web Conference*, WWW'19, 2019.
28. Radhesh Krishnan Konoth, Emanuele Vineti, Veelasha Moonsamy, Martina Lindorfer, Christopher Kruegel, Herbert Bos, and Giovanni Vigna. Minesweeper: An

in-depth look into drive-by cryptocurrency mining and its defense. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS'18, 2018.
29. John Leyden. More and more websites are mining crypto-coins in your browser to pay their bills, line pockets. https://www.theregister.co.uk/2017/10/13/crypto_mining/.
30. John Leyden. Real mad-quid: Murky cryptojacking menace that smacked ronaldo site grows. http://www.theregister.co.uk/2017/10/10/ cryptojacking/.
31. John Leyden. Security opt-in cryptomining script coinhive 'barely used' say researchers. https://www.theregister.co.uk/2018/02/27/ethical_coinhive/, 2018.
32. Natasha Lomas. Cryptojacking attack hits 4,000 websites, including uk's data watchdog. https://techcrunch.com/2018/02/12/ico-snafu/, 2018.
33. Jim Martin. What's the best cpu temperature? https://www.techadvisor.co.uk/how-to/desktop-pc/cpu-temp-3498564/, 2018.
34. Kieren McCarthy. Cbs's showtime caught mining crypto-coins in viewers' web browsers. http://www.theregister.co.uk/2017/09/25/ showtime_hit_with_coinmining_script/.
35. Donny Nadolny. Bitcoin plus miner. https://wordpress.org/plugins/bitcoin-plus-miner/.
36. Lily Hay Newman. The year cryptojacking ate the web. https://www.wired.com/story/cryptojacking-took-over-internet/, 2018.
37. Lily Hay Newman. Now cryptojacking threatens critical infrastructure, too. https://www.wired.com/story/cryptojacking-critical-infrastructure/, 2019.
38. JAKOB NIELSEN. How long do users stay on web pages? https://www.nngroup.com/articles/how-long-do-users-stay-on-web-pages/.
39. Rishab Nithyanand, Sheharbano Khattak, Mobin Javed, Narseo Vallina-Rodriguez, Marjan Falahrastegar, Julia E. Powles, Emiliano De Cristofaro, Hamed Haddadi, and Steven J. Murdoch. Adblocking and counter blocking: A slice of the arms race. In *6th USENIX Workshop on Free and Open Communications on the Internet (FOCI 16)*, Austin, TX, 2016. USENIX Association.
40. Official Journal of the European Union. Directive 95/46/ec (general data protection regulation). http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679.
41. Lukasz Olejnik, Minh-Dung Tran, and Claude Castelluccia. Selling off user privacy at auction. In *21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014*, 2014.
42. Charlie Osborne. Ransomware has been abandoned in favor of cryptojacking attacks against the enterprise. https://www.zdnet.com/article/ransomware-has-been-abandoned-in-favor-of-cryptojacking-attacks-against-the-enterprise/, 2019.
43. Pierluigi Paganini. Thousands of websites worldwide hijacked by cryptocurrency mining code due browsealoud plugin hack. https://securityaffairs.co/wordpress/68966/ hacking/browsealoud-plugin-hack.html, 2018.
44. Panagiotis Papadopoulos, Panagiotis Ilia, Michalis Polychronakis, Evangelos P Markatos, Sotiris Ioannidis, and Giorgos Vasiliadis. Master of web puppets: Abusing web browsers for persistent and stealthy computation. In *Proceedings of the Network and Distributed System Security Symposium*, NDSS'19, 2019.
45. Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos Markatos. Cookie synchronization: Everything you always wanted to know but were afraid to ask. In *The World Wide Web Conference*, WWW '19, pages 1432–1442, New York, NY, USA, 2019. ACM.

46. Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P. Markatos. The cost of digital advertisement: Comparing user and advertiser views. In *Proceedings of the 2018 World Wide Web Conference*, WWW '18, pages 1479–1489, Republic and Canton of Geneva, Switzerland, 2018. International World Wide Web Conferences Steering Committee.
47. Panagiotis Papadopoulos, Nicolas Kourtellis, Pablo Rodriguez Rodriguez, and Nikolaos Laoutaris. If you are not paying for it, you are the product: How much do advertisers pay to reach you? In *Proceedings of Internet Measurement Conference*, IMC '17, 2017.
48. Panagiotis Papadopoulos, Peter Snyder, and Benjamin Livshits. Another brick in the paywall: The popularity and privacy implications of paywalls. *CoRR*, abs/1903.01406, 2019.
49. Panagiotis Papadopoulos, Giorgos Vasiliadis, Giorgos Christou, Evangelos Markatos, and Sotiris Ioannidis. No sugar but all the taste! memory encryption without architectural support. In Simon N. Foley, Dieter Gollmann, and Einar Snekkenes, editors, *Computer Security – ESORICS 2017*, pages 362–380, Cham, 2017. Springer International Publishing.
50. Phidgets Inc. What is a phidget? https://www.phidgets.com/docs21/What_is_a_Phidget.
51. PublicWWW. Source code search engine. https://publicwww.com/, 2019.
52. Adi Robertson. Salon asks ad-blocking users to opt into cryptocurrency mining instead. https://www.theverge.com/2018/2/13/17008158/salon-suppress-ads-cryptocurrency-mining-coinhive-monero-beta-testing, 2018.
53. Thomas Robitaille. psrecord: Record the cpu and memory activity of a process. https://github.com/astrofrog/psrecord.
54. Guenter Roeck. Overview of the lm-sensors package. https://github.com/groeck/lm-sensors.
55. Jan Rüth, Torsten Zimmermann, Konrad Wolsing, and Oliver Hohlfeld. Digging into browser-based crypto mining. In *Proceedings of the Internet Measurement Conference 2018*, IMC '18, pages 70–76, New York, NY, USA, 2018. ACM.
56. ShafayI. Javascript mining — best coinhive alternative for 2019. https://coinogle.com/javascript-mining/, 2019.
57. Liam Tung. Windows: This sneaky cryptominer hides behind taskbar even after you exit browser. https://www.zdnet.com/article/windows-this-sneaky-cryptominer-hides-behind-taskbar-even-after-you-exit-browser/, 2017.
58. WhistleOut Inc. Compare the best cell phone plans. https://www.whistleout.com/CellPhones, 2018.
59. zerodot1. Coinblockerlists - simple lists that can help to prevent illegal mining in the browser or other applications. https://zerodot1.gitlab.io/CoinBlockerListsWeb/index.html, 2019.