

A First Look at the Crypto-Mining Malware Ecosystem: A Decade of Unrestricted Wealth

Sergio Pastrana
Universidad Carlos III de Madrid
Leganes, Spain
spastran@inf.uc3m.es

Guillermo Suarez-Tangil
King's College London
London, UK
guillermo.suarez-tangil@kcl.ac.uk

ABSTRACT

Illicit crypto-mining leverages resources stolen from victims to mine cryptocurrencies on behalf of criminals. While recent works have analyzed one side of this threat, i.e.: web-browser cryptojacking, only commercial reports have partially covered binary-based crypto-mining malware.

In this paper, we conduct the largest measurement of crypto-mining malware to date, analyzing approximately 4.5 million malware samples (1.2 million malicious miners), over a period of twelve years from 2007 to 2019. Our analysis pipeline applies both static and dynamic analysis to extract information from the samples, such as wallet identifiers and mining pools. Together with OSINT data, this information is used to group samples into campaigns. We then analyze publicly-available payments sent to the wallets from mining-pools as a reward for mining, and estimate profits for the different campaigns. All this together is done in a fully automated fashion, which enables us to leverage measurement-based findings of illicit crypto-mining at scale.

Our profit analysis reveals campaigns with multi-million earnings, associating over 4.4% of Monero with illicit mining. We analyze the infrastructure related with the different campaigns, showing that a high proportion of this ecosystem is supported by underground economies such as Pay-Per-Install services. We also uncover novel techniques that allow criminals to run successful campaigns.

CCS CONCEPTS

• **Security and privacy** → **Malware and its mitigation**; • **Social and professional topics** → *Malware / spyware crime*; • **General and reference** → *Measurement*.

KEYWORDS

Malware, Mining, Monero, Cryptocurrency, Cybercrime

ACM Reference Format:

Sergio Pastrana and Guillermo Suarez-Tangil. 2019. A First Look at the Crypto-Mining Malware Ecosystem: A Decade of Unrestricted Wealth. In *Internet Measurement Conference (IMC '19)*, October 21–23, 2019, Amsterdam, Netherlands. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3355369.3355576>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IMC '19, October 21–23, 2019, Amsterdam, Netherlands

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6948-0/19/10...\$15.00
<https://doi.org/10.1145/3355369.3355576>

1 INTRODUCTION

Mining is a key component responsible for the wealth of Blockchain-based cryptocurrencies. This process requires a network of interconnected miners to solve a complex mathematical problem in order to link blocks and maintain the integrity of the transactions. In exchange, miners receive an amount of the mined cryptocurrency as a reward.

The high value of cryptocurrencies has attracted a large number of malicious actors that use hijacked resources to mine these currencies. The illicit crypto-mining threat has grown considerably over the recent years [1], and it is considered one of the top-most cybersecurity threats, even surpassing ransomware according to recent reports [46].

Illicit crypto-mining is typically conducted using either one of these two modes: (i) by using browser-based crypto-mining programs (dubbed cryptojacking [30]), where the mining process is run in scripts (typically JavaScript) embedded in web content; or (ii) by using binary-based crypto-mining malware, where the mining process is embedded in the payload of a malware running in infected machines that are connected to the Internet. In both cases, by using hundreds of hijacked machines, perpetrators can obtain a hash-rate similar to medium-sized mining farms. Each mode has different characterizing features and unique challenges, specially when it comes to devising effective countermeasures. For example, in browser-based cryptojacking the damage ceases when the victim stops browsing the site. Also, users can reduce the threat by restricting the use of JavaScript. Meanwhile, crypto-mining malware entails classical malware-related challenges, such as persistence and obfuscation. Also, since mining increases the CPU load, thus reducing the computer's performance, it might be noticed by end-users. Thus, we observe a new paradigm aimed at evading user- rather than AntiVirus-detection using techniques such as *idle mining* (mining only when the CPU is idle) or reducing CPU consumption when monitoring tools (e.g., Task Manager) are running.

Motivation. While illicit crypto-mining has been less notorious than other threats such as ransomware, it poses nonetheless an important threat to users and organizations; and its presence is an indicator of weaknesses in security practices that must be addressed [24]. First, the profits generated by their miners introduce massive incomes to cyber-criminals. These incomes fuel the underground economy and gear other cyber-criminal activities [41]. Second, this threat causes important economical loses to their victims. By draining the CPU-usage, corporations see how their electricity bills increase and how their hardware rapidly deteriorates [39, 46]. Finally, this indirectly causes a non-negligible environmental footprint [15]. Due to these concerns, browser-based crypto-mining has been widely studied recently, both analyzing it

as a crime [8, 13, 17, 36] and as an alternative business model to monetize web content [29, 31, 35]. However, the literature lacks of a systematic approach to measure the binary-based mining threat at scale.

The first and only seminal work putting this threat in perspective is from 2014 [9]. Authors analyzed 2K malware samples mining Bitcoin and their methodology relied on the analysis of public transactions. However, there has been a significant increase in the number of malware samples monetizing this threat since 2014 [1, 7]. Also, criminals' attention has shifted to other cryptocurrencies, mainly motivated by: i) the proliferation of ASIC mining, which uses dedicated hardware and renders the use of desktop computers no longer profitable for mining bitcoins, and ii) the development of protocols that provide transaction anonymity (such as those used in Zcash or Monero). Anonymous currencies are used by criminals to thwart traceability and they are on demand in underground markets. Commercial reports, in the form of blog posts [7, 30] or white papers [24] provide a further, but limited, view of the magnitude of the problem and the landscape. Security firms have analyzed isolated cases of decontextualized mining operations [10, 11]. However, these studies are limited by the simplicity of the analysis.

In this paper, we aim to bridge these gaps by addressing the following research questions: (1) What are the preferred cryptocurrencies mined by criminals? (2) How many actors are involved in this ecosystem and what are their profits? (3) What is the level of sophistication used in different campaigns and how does this affect the earnings? (4) What is the role of underground markets and what are the tools and techniques adopted from them? (5) How can we improve current countermeasures and intervention approaches? Due to potential *ethical concerns* arisen from this work (see Appendix), we obtained approval from our REB office.

Contributions. Our work focuses on crypto-mining malware as key novelty. By looking at a wide-range of underground communities, where knowledge and tools are shared, we have observed increased interest in this malware. This suggests that cybercrime commoditization plays a key role in the wealth of illicit crypto-mining. We design a measurement pipeline to automatically analyze malware samples observed in the wild and to extract information required to identify the miners and pools, using both dynamic and static analysis. Then, we build a graph-based system that aggregates related samples into campaigns based on a series of heuristics. The system is designed to distinguish campaigns using third-party infrastructure such as Pay-Per-Install (PPI) services or binary obfuscators. This allows to analyze to what extent this threat is sustained by different underground markets [41]. Our analysis system enables the research community to leverage crypto-mining measurements at scale. To the best of our knowledge, this paper presents the largest systematic study of malicious binary-based crypto-mining, providing a reliable lower bound of the earnings made by this criminal industry.

In summary, the main contributions of this paper are:

- (1) We analyze and describe the role of underground communities for the proliferation of the illicit crypto-mining business (§2).

- (2) We present a system that uses both static and dynamic analysis to extract relevant mining-related information from crypto-mining malware, such as wallet addresses and pool domains (§3). Our system uses different techniques to aggregate related samples into larger campaigns represented as a graph that is then mined for further analysis. Additionally, we feed the system with information gathered from various Open-Source Intelligence (OSINT) repositories to further classify and analyze the campaigns.
- (3) We present a longitudinal study of the crypto-mining malware threat using data spanning over more than a decade (§4 and §5). Then, by focusing on Monero, we rely on information gathered from mining pools to measure the earnings gained by each campaign. We also analyze the infrastructure used by criminals and extract the attribution to stock mining software.
- (4) We propose a number of countermeasures, and discuss the efficacy of existing ones together with the open challenges (§6). Then, we contextualize the most important findings of our study with respect to relevant works in the area (§7).

Finally, to foster research in the area, we release our dataset in our online repository.¹ We encourage readers to visit this repository and the extended version of this paper [32] as it provides a wider presentation of the measurements left out of this paper due to space constraints.

2 BACKGROUND

In this section, we first provide an overview of the cryptocurrency mining process. Then we describe the underground economy supporting the illicit crypto-mining threat.

Cryptocurrency Mining. Cryptocurrencies are digital assets that can be exchanged in online transactions. These transactions are grouped into blocks and added to a distributed database known as the blockchain. Each block is linked to its previous block, and the addition of new blocks is done by voluntary miners. In order to provide integrity, miners compute a cryptographic hash of the block together with the solution to a complex mathematical puzzle known as 'Proof-of-Work' (PoW). As a reward, miners receive a certain amount of the cryptocurrency.

The increased value of cryptocurrencies such as Bitcoin or Ethereum has led to the growth of mining farms using specialized hardware known as ASICS, which makes end-user machines useless for mining. However, in 2014 a new PoW known as Cryptonote required not only CPU power but also memory, turning ASIC-based mining inefficient. Additionally, the PoW algorithm changes periodically, thus discouraging ASIC development (which is optimized for specific algorithms) [19]. This allows individuals to mine with their end-user machines. Examples of cryptocurrencies using Cryptonote as PoW are: Monero and Bytecoin.

When mining a block, only the first one solving all PoWs gets the reward. Thus, mining has become a race which highly depends on the hashrate (i.e., number of hashes computed per second) of a miner. A higher hashrate increases the probability of mining a block. A common approach is to mine through public mining pools, which

¹<https://github.com/gsuareztangil/cryptomining-malware>

can be viewed as partnership between workers where the complexity of the mining challenge is distributed among the partners. Each partner contributes with a given hashrate aimed at solving the puzzle, and when the pool successfully mines a block, the reward is divided proportionally to the workers' hashrate. To receive rewards, the workers are identified in the pool. Some pools use proprietary site-keys, like CoinHive, which was the major provider of browser-based mining services until March 2019. Other pools use emails or wallet addresses. The communication between each miner and the pool is done by using the Stratum protocol. This is a clear-text protocol that uses JSON-RPC format to transmit information, including authentication tokens, block puzzles and mined hashes. We refer to the work by Recabarren and Carbutar for a complete analysis of this protocol [34].

Illicit crypto-mining refers to mining carried out by criminals using resources stolen from their victims. This threat appeared with Bitcoin in 2009 [9], but it has increased notably since 2014 due to the inception of PoW algorithms resistant to ASIC-based mining. Illicit crypto-mining is performed in the device of a victim by either embedding a payload in web resources (scripts) executed by browsers [8, 17], or distributing the payload in the form of malware.

The Underground Economy. Underground markets play a key role in the business of malicious crypto-mining. Users with few technical skills can easily acquire services and tools to set up their own mining campaign. Forums are also used for sharing knowledge. To put our study in context, we have analyzed a dataset of posts collected from various underground forums [33], looking for conversations related to crypto-mining. We observe that crypto-mining malware can be easily purchased online, for a few dollars (e.g., the average cost for an encrypted miner for Monero is 35\$).

In particular, we have seen an online service which allow to create customized binaries (e.g., for a particular cryptocurrency and/or a given pool) to mine cryptonote based currencies, for \$13.² It provides several stealthy-related techniques such as idle mining or the use of execution-stalling code [16] targeted to certain conditions (e.g., when the Task Manager is running). Other providers opt to share their miners for free, in exchange for a donation: “*Miner is free, we charge a fee of 2% to cover the time coding.*” The extended version of this paper [32] shows a longitudinal analysis of posts related to crypto-mining in these forums. Here, we show that Monero is the most prevalent currency nowadays.

We also observed that a common topic of conversation concerns (i) “friendly” pools, i.e.: pools that do not generally ban users displaying botnet-like behaviors, or (ii) how to remain undetected otherwise. For instance, users claim that a good trade-off between profitable hash-rates and a long-lasting mining strategy is using botnets with less than 2K bots. For bigger botnets, many discussions and tutorials explain how to configure proxies and provide advice on how to reduce the risk of being exposed: “*The best option is to use a proxy and you can use any pool. Contact me for PM, I am willing to help.*” Also, we found various conversations with users looking for partners and offering custom (private) mining pools: “*In my pool there is no ban by multiple connections.*”

²For ethical reasons we do not disclose the URL of this service

Finally, we note that it is also possible to purchase *all-you-need* packages, including tools and services, with a guarantee period and maintenance (e.g., re-obfuscation when the miner is detected, or updates to new versions).

Take-Away: The support offered by underground communities to criminals explains the sharp growth on the amount of malware monetizing their victims. This motivates the need for a longitudinal measurement of this threat. We show that Monero is currently the most discussed crypto-mining coin by underground forum users.

3 MEASUREMENT METHODOLOGY

A general overview of the measurement methodology is presented in Figure 1. For the sample collection, we query both public and private repositories of malware and different intelligence feeds as described in §3.1. We make a number of sanity checks for each sample to ensure that we only feed crypto-mining malware to our pipeline (see §3.2). We also collect OSINT related to running botnets and relevant Indicators of Compromise (IoCs) observed in malware samples.

A key phase in our pipeline is to analyze relevant samples both statically and dynamically as described in §3.3 (*Binary, Sandbox, and Network Analysis* in Figure 1). The goal of this multi-step phase is to extract the following information from each miner: (i) the pool or address which the crypto-mining malware connects to for mining, and the identifier used to authenticate themselves into this pool, (ii) the addresses of the e-wallets where mined cryptocurrencies are paid to³, which in most cases coincide with the identifier, (iii) URLs where the malware connects to or is seen at, and (iv) other metadata obtained from intelligence feeds such as when the sample was first seen or related samples (e.g., dropped binaries).

The next step is to analyze the mining pools that the miners work with. We decouple connections made to proxies (that in turn connect to pools) from the connections to the actual pools. We then look at the profit reported by each of the wallets in a pool. These two steps are described in §3.4 and are referred to as *Profit Analysis* in Figure 1.

Finally, we aggregate related samples into campaigns and analyze them separately in the *Aggregation* step as described in §3.5. For this, we create a graph which interconnects crypto-mining malware that: (i) share a common execution ancestor (i.e., dropper) or are packed together, (ii) accumulate their earnings into the same wallet, (iii) share common infrastructure (e.g., proxies or hosting servers), or (iv) relate to the same IoC related with mining campaigns — gathered both from OSINT reports and our own investigation. We then enrich every interconnected sub-graph (campaign) to include details about related infrastructure used in each campaign (i.e., stock mining software or Pay-Per-Install services).

3.1 Data Gathering

We collect malware samples, metadata and OSINT information, and known mining tools from various sources.

Malware. We rely on public and private feeds from:

³The terms *address* and *wallet* are used interchangeably in the literature.

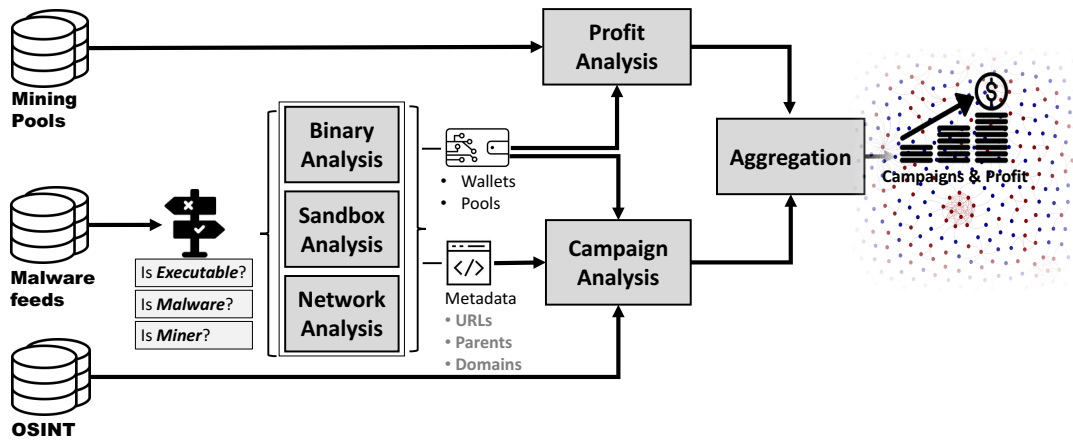


Figure 1: Overview of our processing pipeline and measurement methodology.

- (1) **Virus Share.** This is an online community that shares torrents to malicious binaries. We use it to gather our initial dataset of raw binaries.
- (2) **Virus Total.** This is an online service containing both publicly and privately available services to the security community through an API (Application Programming Interface). This service is a subsidiary of Google that runs multiple AntiVirus (AV) engines and offers an unbiased access to resulting reports. We use the private API to download malware binaries. We then query the public API to obtain threat intelligence (which we refer to as metadata in the paper). In particular, we collect metadata from all the samples obtained through all sources listed in this section.
- (3) **Hybrid Analysis.** This is an online community providing malicious binaries and threat intelligence obtained from static and dynamic analysis of the samples. Thus, we use this community to fetch readily available intelligence when possible.
- (4) **Other Sources.** We have developed a crawler to fetch samples from a variety of online communities such as *malc0de.com* or *vxvault.net*. Our pipeline also aggregates malware feeds from cybersecurity companies. For the purpose of this paper, we have received feeds from Palo Alto Networks.

Refer to [32] for details about the dataset overlaps.

Metadata. When available, we primarily rely on the following metadata to put our study in context: (i) the first time the sample was seen in the wild, (ii) the URLs where the sample was seen, (iii) the list of parents that are known to have dropped the binary under analysis, and (iv) the list of contacted domains.

Stock Mining Software. We also collect binaries from *known mining frameworks*, such as *xmrig*⁴ or *xmr-stak*⁵, that are hosted in various public repositories. While these binaries are not badware *per se*, their usage is deemed malicious when run by malware. Our assumption is that the usage of proprietary software to mine is not the norm. Anecdotal evidence observed during the course of a preliminary investigation has shown that miscreants rely on legitimate

— open-source — mining tools. The modus operandi of the malware is to fetch one of these tools (i.e., acting as a dropper) and run it in infected machines. Mining is configured with the wallet of the miscreant, where the rewards are paid by the network. One of our goals then is to understand if this assumption holds true and how many campaigns are using stock mining software illicitly.

Summary. Our data collection registers over 4.5 million samples (see §4.1 for a breakdown), which have been active between early 2007 and early 2019. This includes about 1K versions of known mining tools from 13 different frameworks. Our initial data contains a wide-range of samples, many of which are irrelevant to this study (e.g., web-based cryptojackers). Thus, we next describe the rules we use to consolidate the dataset where we report our findings with.

3.2 Sanity Checks

One important aspect when systematizing the analysis of malware is properly curating the dataset [20].

We perform the following sanity checks for each sample: (i) *is it malware?* (ii) *is it a miner?*, and (iii) *is it an executable sample?*

First, we rely on Virus Total reports to learn if a sample is malware. Virus Total have been shown to perform remarkably well when providing malware feeds according to a recent comparative analysis of Threat Intelligence [21]. In particular, Virus Total was able to detect 99.94% of the threats over one of the largest non-targeted⁶ malware aggregators. We assume that a sample is malware if at least 10 AV vendors flag the sample as malicious. While this is a common practice in other works in the area [12, 25, 26], we acknowledge that having a solid ground-truth is essential (see discussion in §6). Thus, we use a white-list with the hashes of known mining tools, to ensure that they are not considered as malware samples in our study. This white-list is compiled from binaries collected from various online open-source repositories. A wallet extracted from a malware sample is considered an ‘illicit’ wallet throughout our dataset. Therefore, we exceptionally keep samples with less than 10 AV positives when it contains an illicit wallet.

⁴<https://github.com/xmrig/xmrig>

⁵<https://github.com/fireice-uk/xmr-stak>

⁶Meaning that they target malware threats to generic platforms. Other targeted malware aggregators focus on threats that specifically target platforms like Facebook and “that are not as relevant to most Virus Total users” [21].

Second, we assume a malware is a crypto-mining tool when there are IoCs that reveal activity related to mining, such as connections made through the Stratum protocol or DNS resolutions for web mining pools. To this end, we apply publicly available YARA rules⁷ to our samples. Additionally, we compare OSINT information related to known mining campaigns with IoCs extracted from the samples (e.g., file hashes or network data). We also use advanced queries from Virus Total and Hybrid Analysis to look for malware that meet the following criteria: (i) samples that contact domains of known mining pools, (ii) communicate through the Stratum protocol, and (iii) are labeled as “Miner” (or related variants) by more than 10 AVs.

Finally, to understand whether a malware is executable, we rely the magic number from its header, and consider only those related to executables like PE, ELF or JARs. §6 provides discussion of the limitations behind these assumptions.

3.3 Extraction of Pools and Wallets

With our dataset of crypto-mining malware, we rely on:

- (1) *Static Analysis*: we perform binary inspection to extract evidences of mining activity embedded *into the binary*.
- (2) *Dynamic Analysis*: we then use environmental information obtained from the *execution of the binary* in a sandbox. Specifically, we obtain the network traffic, the dropped files, processes opened, and command line parameters passed to the binaries. When available, we rely on reports provided by *Virus Total* and *Hybrid Analysis* through their API service.

In some cases we are able to find identifiers (e.g., wallets or emails) and pool names using static analysis. In other cases, we rely on dynamic analysis to extract these identifiers from the network activity or the command line processes. In both cases, we process the output of these two analyses using heuristics and regular expressions to extract the following information:

Cryptocurrency wallets. Miners connect to the pools using the Stratum protocol [34]. Upon connection to the pool, they send a request-for-work packet with the identifier of the miner in a ‘login’ parameter. This identifier can be extracted from the command line options passed to the mining tool or directly from the network traffic. We also process the type of wallet to understand the cryptocurrency (e.g.: Monero, Bitcoin or Ethereum) the malware is intending to mine.

Mining pools. We collect additional information such as domains and IPs of mining pools and proxies. Similarly to wallet addresses, this information is typically extracted from either the command line of the process invoking the mining tool or from the network traffic. Typically, miners connect to a known pool.⁸ In some cases, the miner either uses a proxy or mines against a private/unknown pool.⁹ We consider that a miner is using a proxy if we record mining activity for the corresponding wallet in a known pool (see §3.4).

⁷<https://github.com/Yara-Rules>

⁸We consider known pools as those listed in public sources, e.g.: <http://moneropools.com/> or <http://www.blockchain.com/pools>.

⁹While the use of private pools is encouraged in certain underground communities, we have observed few samples using private pools.

3.4 Collecting Mining Activity

One of the main challenges when measuring the impact of the malicious crypto-mining campaigns is the difficulty to accurately estimate the profits. In browser-based cryptojacking, recent works use estimations of the number of visitors per hour for similar websites and the average hashrate of a single visitor (victim) [5, 8, 17]. This is highly inaccurate as evidenced by the variances reported by concurrent related works (see §7). In the case of crypto-mining malware, the actual wallet which the mining reward is paid to can be extracted. We leverage public information obtained from mining pools (which include total reward paid to wallets) to get a more approximate estimation of the profits.

For all the extracted wallets, we queried the most prevalent mining pools to collect activity associated with these wallets. While the amount of information offered by each pool varies, it always contains the timestamp of the last share, the current (last) hashrate and the total amount of currency paid to the wallet. Additionally, some pools also provide the historic hashrate of the wallet and the list of payments done to the wallet (including timestamp and amount). While the total paid is always available, some pools only provide payment data for the last period (e.g. a week or month). Since we are interested on studying how the payments evolve across time, we use public APIs to collect this information periodically for a period of 10 months (July’18-April’19). As a single wallet can use more than one mining pool, we queried all the wallets against all the pools. Then, to estimate profits, we aggregate all the payments sent by the pools to the wallets. In general, we report payments using XMR. To ease readability we also report the equivalent in US dollars (USD). However, we note that we do not have information about when the criminals have cashed-out their earnings (if ever). Thus, it is hard to extract an exact figure in USD (and other currencies) due to the fluctuations on the value of Monero. To approximate this value, we dynamically extract the exchange rate between XMR and USD of the date when the payments were made, if available. We use the average exchange rate of 54 USD/XMR in cases where historical payments are unavailable.

3.5 Campaign Analysis

Two major limitations in related works are: i) the simplicity in which they analyze related mining campaigns, and ii) the inability to study anonymous cryptocurrencies such as Monero (as discussed in §7). Thus, we aggregate samples into campaigns following a novel methodology that leverages various characterizing features observed in the wild.

Spreading Infrastructure. We distinguish two types of infrastructure used to spread the malware: one that can be owned and another one that belongs to a third-party and can be rented (e.g., botnets that are monetized as PPI services and that are used for mining). When available, we link samples to known botnets by querying OSINT information with IoCs extracted from the samples. We refrain from using these botnets to aggregate samples as we detail later. However, we use them to enrich the information of the campaigns in a post-aggregation phase. This way, we can draw conclusions about the number of campaigns using *known* third-party infrastructure. However, since we rely on public intelligence feeds,

a limitation of this approach is that samples using *unknown* third-party infrastructure (e.g., offered in underground markets) might be aggregated together in a single campaign. In these cases, we can guarantee that the campaign runs on top of the same infrastructure. This is relevant to law enforcement agencies when devising take-downs strategies. Thus, our analysis considers campaigns that are either from the same actor or a group of actors that use the same infrastructure, independently from the monetizing approach used by the operators of the infrastructure that spreads the samples. Analyzing whether profits from a campaign are given to a single actor or a group of actors is out of the scope of this paper.

Grouping Features. We rely on the following features to group samples into campaigns:

Same identifier: In order to get rewards from the mining pools, workers must mine using a unique identifier, which in most cases corresponds with the wallet address to which payments are made. In other cases, these are e-mails or other identifiers, like user-generated names. If two samples contain the same identifier, it means that they are accumulating earnings in the same wallet and thus they are grouped together. Some mining tools contain donation wallets to reward the developer, which is done by mining for a certain time (typically 2-5%) using the donation wallet. While this is configurable and can be turned off, we have observed a few samples doing donations. We note that the CPU cycles donated are also hijacked from the victim and therefore inflict harm to her. However we are primarily interested in measuring the earnings of the miscreants, and thus donation addresses are excluded from the aggregation. For this reason, we create a white-list by manually extracting donation wallets from known mining software repositories. We also enrich our white-list using Google searches (e.g., looking for “Monero” and “donation”) and manually analysing the results. We have white-listed 14 donation wallets directly obtained from the developers’ sites. Due to limitations in the manual extraction process, we could be missing donation wallets. This can result in the over-aggregation of two independent campaigns as discussed in §6. Non-white-listed donation wallets display a characterizing pattern: the same wallet (the donation wallet) appears together with different wallets (from the miscreants) in multiple samples across our dataset. However, we do not observe this pattern after white-listing all donation wallets we account for. This suggest that we have effectively white-listed all donation wallets.

Ancestors: In many cases, the same sample is used to download additional malware. This is the case of droppers, which adapt based on information gathered from the infected host, e.g.: operating system or processor capabilities. Accordingly, if a sample is parent of two samples with different wallets, these are grouped together. Ancestors and other dropped files that are not directly intended for mining are considered auxiliary binaries and we refer to them as ancillaries. This includes samples that do not have a wallet.

Hosting servers: We use metadata from the samples to extract the URL from where the malware was downloaded. A common approach is to host the malware (or even stock mining software) in public cloud storage sites such as Amazon Web Services (AWS), Dropbox or Google Drive (see §4.2). Thus, we aggregate two samples if either they are downloaded from the same IP address which does not resolve to a domain from a public repository, or if they are

downloaded from exactly the same URL, e.g: `hxxp://suicide.mouzze-had.su/gpu/amd1.exe`. We also include the parameters to avoid those cases where a parameter is used to uniquely identify the resource being hosted, e.g. `hxxp://file8desktop.com/download/get56?p=19363`. This approach has as limitation that we are not aggregating resources where a URL contains ephemeral information (e.g., timestamps or click-IDs), even when they point to the same resource in the server. However, this limitation is partially overcome due to other sources for aggregation.

Known mining campaigns: As mentioned, we collect IoCs (e.g., domains or wallets) from mining operations reported publicly. We look at IoCs that are known to belong to a given mining operation, and look for matches against samples in our dataset. We group two samples if they belong to the same operation. In our analysis, we have collected IoC for the following mining operations: Photominer [6], Adylkuzz [10], Smominru [11], Xbooster [43], Jenkins [3] and Rocke [23]. However, our methodology is designed to easily include data collected from new operations.

Domain aliases (CNAMEs): During our investigation, we observed many samples using domain aliases (i.e., CNAMEs) that resolve to known mining pools. In these cases, miscreants create one or various subdomains for a domain under their control, and set these subdomains to be aliases of known mining pools. Since the resolution is done for the CNAME rather than for the mining pool, they thwart defenses blacklisting mining pools (see §6 for a discussion on anti-analysis techniques). To address this evasion method, we perform DNS requests for all the domains extracted from our samples, and look for responses pointing to known mining pools from a CNAME. Since CNAMEs might have changed, we also query a DNS history-resolution service provided by AlienVault (<https://www.threatcrowd.org>) Accordingly, we aggregate samples using the same domain alias.

Mining proxies: Mining using a large number of machines (i.e., more than 100) with the same wallet raises suspicion of botnet usage, and mining pool operators might opt to ban the miner. To prevent this situation, offenders use mining proxies that gather all the shares from the different bots and forward the aggregated to the pool. Thus, pool operators only receive responses from a single machine, the proxy. As described in §3.3, we identify various samples using proxies. We aggregate together samples that use the same proxy.

Aggregation. To measure the number of related campaigns and how they are structured, we build a graph where nodes are elements of a given resource (e.g., malware samples, proxies, or wallets) and the edges are determined by the relationships mentioned above. We consider each connected component of the graph as a single campaign, where the internal nodes of the graph represent the crypto-mining malware together with the infrastructure used by the campaign.

Enrichment. After the aggregation, we enrich each campaign with samples related to known Pay Per Install (PPI) services, and mining tools. We emphasize that these features are only informative and they are not used to aggregate campaigns. We next explain the rationale behind this.

Botnets and PPI: A common approach to spread malware is through PPI services, where customers pay a fee to botnet operators in order to spread their malware [2]. Due to commodization of cybercrime

services, purchasing a botnet to spread malware is simple and open to anyone with few technical skills, e.g.: by leveraging underground markets [44]. During our analysis, we have observed samples belonging to various botnets that are commonly used as PPI services, such as Virut or Nitel. Since these are known third-party infrastructures, two samples using these services are not necessarily related to each other and thus are not aggregated together.

Stock Mining Software: During our exploratory analysis, we have observed that many campaigns use stock mining software. This is, the hash of a file dropped by the malware matches with one of the hashes in our collection of mining tools. Actually, we have observed that some crypto-mining malware fetch this executable directly from the official GitHub repository. However, we have also observed that some miscreants fork these projects and make minor modifications to the mining tool, e.g. to remove donation capabilities.

We use Fuzzy Hashing (FH) to pick up on the aforementioned modifications and to relate these samples with known mining tools. FH is a similarity preserving hash function that allows to compare binary files. Specifically, FH computes a fingerprint of each binary in such a way that any two binaries that are almost identical map to a “similar” hash value. Fuzzy hashing has been shown to be an effective way of comparing malware [22]. In our pipeline, we use context triggered piecewise hashing [18] and compute the distance between the FH of all samples in a campaign and the FH of all known mining tools. We choose a conservative distance threshold of 0.1 as it performs well when comparing malware [22]. Thus, samples with a distance lower than 0.1 are considered as stock mining tools.

4 THE BINARY-BASED MINING THREAT

In this section, we present the analysis of our measurement. We first present our dataset (§4.1), which contains malware seen for over a decade. Next, we perform a longitudinal analysis through the lens of our dataset (§4.2). Then, we characterize the type of mining pools and currencies we have seen (§4.3) and study the earnings of the campaigns in the most prevalent crypto-currency, i.e.: Monero (§4.4).

4.1 Dataset

Our study results from the collection of 4.5 million malware samples from the range of sources described in §3.1. We then apply sanity checks to tailor our analysis to crypto-mining malware only, resulting in a total of: (i) 1,017,110 miner binaries, and (ii) 212,923 ancillary binaries.

The samples in (i) are samples where we have observed mining capabilities together with an associated wallet and a pool address. The samples in (ii) are samples used by the miners to run the mining operation (e.g., bot clients or loaders). In total, our study leverages over 1.2 million crypto-mining malware samples. Table 1 shows a summary of our dataset, together with the breakdown of data sources and the type of resources we resort from.¹⁰ Our largest source of miners is Virus Total and the smallest is Virus Share. As for the resources, we collect the largest number of wallets and pools through dynamic analysis (sandbox and network analysis).

¹⁰See [32] for details about other sources crawled.

Category	Type	#Samples
Summary	ALL EXECUTABLES	1,230,033
	Miner Binaries	1,017,110
	Ancillary Binaries	212,923
Sources	Virus Total	956,252
	Palo Alto Networks	628,915
	Hybrid Analysis	857
	Virus Share	519
Resources	Sandbox Analysis	1,143,384
	Network Analysis	258,564
	Binary Analysis	10,204

Table 1: Our dataset of miners and ancillaries, with the collection sources and the number of resources.

The data collection dates to March 2007 to capture the structure of the third party infrastructure from their early stages. However, malicious mining activity starts getting traction in 2011.

4.2 Longitudinal Analysis

We extract 16,050 different crypto-mining identifiers from a total of 103,894 samples. As mentioned, these mostly include addresses of wallets from various cryptocurrencies, but we also find emails and other identifiers used to authenticate the miner in the pool to later pay them the corresponding reward. In the case of wallets, we use regular expressions to detect the associated currency. Regarding the emails, we observe that the majority (97%) are used as identifiers of one of the most popular mining pools, i.e.: *minergate*.

Overall, we aggregate samples into 11,387 different campaigns. Leftmost side of Table 2 shows the breakdown of the number of campaigns per type of identifier (i.e., wallets and other identifiers). Recall that wallet addresses represent the public key of an electronic wallet in a given cryptocurrency. Thus, we show the breakdown for the different cryptocurrencies for which we have wallets. Note that two or more identifiers can be used in the same campaign, for example due to a change of a previous wallet address after being banned [11]. Monero is the cryptocurrency most frequently used, followed by Bitcoin. There are at least 18 campaigns using two or more currencies. While most of the campaigns are composed by one or few wallets, we observe campaigns having up to 304 different identifiers.

Out of all the samples with wallets, we rely on the *first seen* value (obtained from the metadata of the samples provided by Virus Total) to analyse the evolution of number of samples mining Bitcoin and Monero. Rightmost side of Table 2 shows the number of samples with wallets by year for the most prevalent cryptocurrencies, i.e.: Monero (XMR) and Bitcoin (BTC). Overall, the dataset contains 7.6K BTC and 62K XMR malware samples.¹¹ Judging by the number of samples and the distribution across time, we can confirm the decreasing interest in Bitcoin in favour of Monero. Moreover, we have queried available Bitcoin pools with the BTC addresses, and observed negligible earnings (i.e., less than 5K USD). Note that the data collection ended in early 2019 and thus data from this year is partial. Also, due to constraints in the Virus Total rate limit we

¹¹These are samples with embedded wallets and does not include ancillaries.

# campaigns with wallet addresses for:			# samples seen in:			
Monero	2,449	Aeon	57	Year	BTC	XMR
Bitcoin	1,535	Sumokoin	18	2012	9	1
zCash	178	Intensecoin	8	2013	23	3
Electroneum	150	Turtlecoin	3	2014	223	281
Ethereum	132	Bytecoin	1	2015	115	1.6K
Mixed			17	2016	461	8.7K
<i>Sub-total</i>			4,548	2017	3.8K	31K
With other identifiers:				2018	1.3K	6.2K
Email			5,008	2019	1*	49*
Unknown			2,195	~19?	1.7K	14K
TOTAL			11,751	ALL	7.6K	62K

Table 2: Leftmost side of the table: Number of campaigns per currency, amount of e-mails and unknown identifiers (i.e., not associated with a known currency). Rightmost: Number of samples (*partial data) seen in a given year for Bitcoin (BTC) and Monero (XMR).

could not retrieve the *first seen* entry for some of the most recently discovered samples. However, we attribute this samples to 2019 (denoted as ~19? in Table 2). As it can be observed, there are 4 samples that were seen in 2012 and in 2013 and that have later been mining Monero, which was released in 2014. This is due to malware reuse, i.e., malware samples that dynamically update their code and execute components downloaded at a later stage (after installation time), namely *droppers*. See [32] for further details on this.

We provide more details about the number of samples, wallets, and earnings per campaign in the extended version [32].

We further look at the most popular URL domains hosting crypto-miners. We observe that GitHub is the chosen site used to host the crypto-mining malware. This is because GitHub hosts most of the mining tools, which are directly downloaded – for malicious purposes – by droppers as discussed before. Additionally, GitHub is also used to host modified versions of the miners (e.g., by removing the donation capabilities or adding further capabilities). It is also used to host ancillary malicious tools [40]. We also observe that there are other public repositories and file sharing sites such as Bitbucket or 4sync, and web hosting sites such as Amazon (AWS), Google, or Dropbox. One can also find mining malware hosted through torrent sites (*b-tor.ru*), entertainment sites (*telekomtv-internet.ro*), or hosted as attachments in the Discord app, a voice and text chat (*cdn.discordapp.com*). There are also URL-shortener sites (*goo.gl*). This altogether shows that crypto-miners largely rely on publicly available third-party servers. The use of these services provides an economical incentive when compared to other approaches that use dedicated infrastructure such as bullet-proof servers – that are more resilient against take-downs.

Our longitudinal analysis confirms previous reports positioning Monero as the preferred currency used by miscreants for crypto-mining malware [7]. Thus, in the rest of the paper we focus our attention on campaigns using Monero.

Pool	XMR Mined	#Wallets	USD
crypto-pool	429,393	487	47,261,821
dwarfpool	168,796	461	1,088,516
minexmr	74,396	608	5,320,397
poolto	29,044	38	35,815
prohash	12,833	54	275,471
nanopool	5,205	375	858,949
monerohash	4,046	217	477,557
ppxxmr	3,860	185	518,487
supportxmr	3,217	241	443,087
Others (8)	2,797	314	325,034

Table 3: Popularity of the different mining pools ranked by the amount of XMR mined by malware.

4.3 Mining Pools

There are two possible strategies for mining: joining a pool or mining alone (which we call *solo-mining*). Using mining pools instead of “*solo-mining* strategies” has several advantages: it increases the chances of receiving payments for mining and reduces the need for specialized mining equipment. Selecting a mining pool is not straightforward because it depends on many dynamic factors such as the current hashrate of the pool, or the complexity required for mining. Pools with a high number of workers are more likely to mine a block faster, but the reward received is lower. To understand the popularity of the different mining pools among criminals, we look at the number of wallets and the amount of XMR mined over the most consolidated pools (according to various benchmarks such as <http://moneropools.com>, or <https://minexmr.com/pools.html>) that provide public information about the wallets. Table 3 provides a list of these pools ranked by popularity among criminals (in terms of earnings). We show that the most popular pools are *crypto-pool* and *dwarfpool*, with more than 429K and 168K XMR mined respectively. When looking at the number of wallets observed, the most common pool used is *minexmr*, with (at least) 608 wallets. An interesting pool not included in our analysis is *minergate*. We have found 4,980 emails mining at this pool in our dataset. Since *minergate* does not provide public information about the rewards paid to the miners, we are unable to estimate profits from this pool.

Our analysis show that 49.3% of the campaigns use or have used more than one pool. Indeed, 97% of the campaigns with largest earnings (i.e., over 1K XMR) have used more than one pool. However, seven campaigns with earnings over 10K are using just one pool. Out of these, six use *dwarfpool* and one uses *crypto-pool*. This suggests that mining in different pools depends on different strategies, probably driven by the revenues from each pool and their banning policies.

4.4 Monero-based Campaigns

As shown in Table 2, we find 2,360 campaigns mining Monero. Out of those, we are able to get payments to 2,145 campaigns through querying the various mining pools.

We summarize the results of our aggregation in Table 4 and show some demographics for the top 10 campaigns. Note the difference between USD and XMR in some campaigns. As explained before,

Campaign	#S	#W	Period	XMR	\$
C#627	66	7	06/16 to active*	163,756	20 M
C#3027	20	2	10/16 to 04/18	59,620	8 M
C#268	134	4	01/15 to 02/19	42,069	323 K
C#102	59	1	09/14 to 04/18	32,886	53 K
C#693	106	2	08/14 to 02/19	27,985	95 K
C#1290	91	14	06/16 to active*	27,086	1 M
C#10465	6	1	09/16 to 04/18	23,300	2 M
C#3311	9	1	06/16 to 05/18	22,520	5 M
C#2642	46	1	09/14 to 04/18	21,389	42 K
C#2202	25	1	09/14 to 04/18	20,694	38 K
TOP-10	562	34	2014/08/30 - *	441,305	38 M
ALL-2235	64 K	2532	2014/07/18 - *	740,927	58 M

Table 4: Top 10 campaigns ranked by amount of XMR mined. C=Campaign #S=Num. of samples, #W=Num. of wallets, and active* on April 2019. Recall that the exchange rate to USD is computed dynamically based on when the payments were made.

this is due to fluctuations of the XMR value and depends on when payments were made. A note of precaution when looking at the USD figures as we are unaware of when criminals cash-out their moneros. Thus, we prefer to report our findings in primarily in XMR. Overall, we estimate that there are at least 2,235 campaigns that have accumulated about 741K XMR (58M USD). Some of them are still active. Interestingly, just a single campaign (C#627) has mined more than 163K XMR (20M USD), which accounts for about 22% of the total estimated. This campaign is still active at the time of writing and it is later studied in §5. We observe that only the top 10 campaigns mine more than the remaining 2,225 ones. Overall, we observe that 99% of the campaigns earn less than 100 XMR. We also observe that while majority of the campaigns earn very little, there are a few campaigns overly profitable. This indicates that the core of this illicit business is monopolized by a small number of wealthy actors.

There are campaigns with a large number of samples, with up to 12K in the case of C#4 (see further details in our repository¹²). However, some of the most profitable campaigns have few samples (e.g., C#10465 or C#3311). This means that either the samples in those campaigns have infected a large number of victims, or that other samples from the campaign are not detected by any AV. In either case, it suggests that there are some miscreants that are proficient in remaining undetected. In the next section we analyze the infrastructure and stealth techniques used by the different campaigns, and how this affects their efficiency.

While most of profitable campaigns started in 2016 or earlier, we observe recent campaigns with large earnings. In particular, 21 campaigns that started in 2018 have mined more than 100 XMR, 12 of which are active at the time of writing (April 2019).

4.5 Infrastructure

We next analyze the third-party infrastructure used in the different Monero campaigns.

¹²<https://github.com/gsuareztangil/cryptomining-malware>

Mining software. We show that `xmrig`, `claymore` and `niceHash` are the most popular tools we account for. With the current distance threshold in our Fuzzy Hashing algorithm, we found no evidence pointing to the use of other less popular tools such as: `cast-xmr`, `jceMiner`, `srbMiner`, or `yam`. When using a higher threshold, we found one campaign using `xmr-stak`. Overall, the top most popular frameworks account for approximately 18% of the Monero campaigns.

Domain aliases for mining pools. A common mitigation strategy often suggested in commercial reports [24] is to block known mining pools, using blacklists. Criminals create CNAME domain aliases to evade this mitigation. In our analysis, we observe 215 different CNAMEs. Most of these are aliases of *minexmr* (176), *crypto-pool* (21) and *dwarfpool* (14). Interestingly, there are two aliases (*x.alibuf.com* and *xmrf.fjhan.club*) which have been eventually used to hide two different pools each. This suggests again dynamic changes in the mining strategy used by criminals to maximize their revenue. We note that the former alias is actually part of the most profitable campaign (C#627), which is detailed in §5.

Pay-Per-Install services. In order to spread malware, criminals use commodity botnets offered as PPI services in underground markets [2, 44]. We find samples from 3 different botnets offering PPI services. In particular, we observe 511 samples associated with the Virut botnet (in 44 different campaigns), 46 from Ramnit (in 10 campaigns) and 27 from Nitel (in 3 campaigns). Also, in one of the biggest campaigns (C#8), known as Photominer [6], we find 346 samples (3.01% of the samples belonging to this campaign) using Virut to deploy the mining operation. Recall that campaigns are automatically extracted. Observing campaigns from botnets that are known to the community shows that our heuristics provide a reliable aggregation. Yet, our framework steps up finding novel campaigns as shown in §5.

Obfuscation. A common practice when spreading malware is to obfuscate the binary to avoid detection. Criminals typically use existing tools, such as well-known packers (e.g., UPX) or crypters. Packers can be fingerprinted more easily than crypters, but crypters – which are usually purchased in underground markets – increase the cost of the operation. By leveraging the F-Prot unpacker [4], we extract packer information associated with each sample (when applicable). This tool also identifies compression algorithms, which are not considered obfuscation. Then, we look at the entropy to detect whether some other unknown obfuscation is applied in samples where no packer or compression algorithm is detected. In our implementation, we choose a conservative threshold of 7.5 (where 8 means total randomness) to decide when a sample is obfuscated, which is more restrictive than values tested in previous works [27, 42]. We found that around 30% of the samples are obfuscated. We consider that a campaign uses obfuscation if a large proportion of their samples (i.e., 80%) are obfuscated. While this is the ratio in the overall dataset, we found that only 4.16% of Monero campaigns use obfuscation. Table 5 summarizes the number of samples using obfuscation together with the tool used to obfuscate it. UPX is by far the most common tool used. Interestingly, we have seen many binaries created using AutoIt (a Windows-based scripting language) which by default packs the script into a PE file

UPX	328,493	eval	2,032
NSIS	17,464	docwrite	1,490
maxorder	5,988	ARJ	858
SFX	3,928	CAB	721
INNO	2,423	Enigma	710
Others			4,019
Not packed			862,712

Table 5: Packers used for binary obfuscation.

using UPX. In §6 we discuss the limitations of analyzing obfuscated binaries.

Analysis. Table 6 shows the third-party infrastructure, stealth techniques and period of activities for the different Monero campaigns (both divided according to their profits, and overall). While only 1.1% of the campaigns use domain aliases, a higher proportion is found in most profitable campaigns (9.4% of those mining between 1K and 10K XMR, and 26.7% of those mining more than 10K XMR). A similar situation happens with proxies and PPI services, which are more common in successful campaigns (i.e., with larger earnings).

Most profitable campaigns have longer period of activity (46.7% have been active since 2014). However, we also observe a high portion of campaigns (26.7%) operating only for 1 or 2 years and still having large profits. We also note the percentage of campaigns active before and after changes in the PoW (Proofs-of-Work): on 06/04/2018, 18/10/2018 and 09/03/2019. These changes require the update of mining software. This means that either botnet operators have to update their bots, or customers of PPI services must buy new installs. We show that most of the campaigns stopped due to PoW updates: around 72.4% in April 2018, 89.3% in October 2018 and 96.5% in March 2019.¹³ This means that changes in the PoW algorithm might be an effective (though unwitting) countermeasure, as discussed in §6.

4.6 Take-Aways

In summary, the main take-aways of our analysis include:

(1) We observe that it is no longer profitable to mine Bitcoin, and current criminal efforts focus on mining ASIC-resistant currencies. We also show that there are a small number of actors that monopolize the crypto-mining malware ecosystem.

(2) We note that some successful mining campaigns are very complex in terms of the size and infrastructure supporting the campaign. Our data shows that about 11% of these are supported by other underground economies such as third-party Pay-Per-Install botnets. On the contrary, we also observe very profitable mining campaigns that do not appear to use a large supporting infrastructure. Instead, they are effective campaigns (due to their long lifetime) with obfuscation and novel evasion techniques, e.g.: using CNAMEs to bypass blacklist-based detection.

(3) We estimate that the malicious ecosystem has currently mined at least 4.37% of the total Monero in circulation (approximately 58M USD). These numbers must be added to estimations made using web-browser cryptojacking in parallel work.

¹³Given that our data is from April 2019, some of these campaigns might not be defunct, since it might take some time to update the miners.

XMR Mined	< 100	[100-1k)	[1k-10k)	>10k	ALL
#Campaigns	2,013	154	53	15	2,235
THIRD PARTY INFRASTRUCTURE					
PPI	1.3%	3.2%	9.4%	13.3%	1.7%
Mining SW	8.6%	14.9%	30.2%	13.3%	9.6%
Both	0.4%	1.3%	7.5%	0.0%	0.7%
STEALTH TECHNIQUES					
Obfuscation	4.0%	5.2%	3.8%	0.0%	4.0%
CNAMEs	0.3%	5.2%	9.4%	26.7%	1.1%
Proxies	2.6%	6.5%	3.8%	20.0%	3.0%
PERIOD OF ACTIVITY					
+ Apr-18	24.6%	57.8%	49.1%	33.3%	27.6%
+ Oct-18	9.1%	27.3%	18.9%	33.3%	10.7%
+ Mar-19	2.7%	13.0%	5.7%	13.3%	3.5%
Start: 2014	0.2%	4.5%	11.3%	46.7%	0.2%
Start: 2015	0.2%	1.9%	3.8%	13.3%	0.2%
Start: 2016	5.5%	26.0%	41.5%	40.0%	5.0%
Start: 2017	37.3%	51.3%	41.5%	0.0%	33.6%
Start: 2018	51.7%	13.0%	1.9%	0.0%	46.6%
Start: 2019	0.5%	2.6%	0.0%	0.0%	0.4%
Years: 0	69.6%	11.0%	1.9%	0.0%	62.7%
Years: 1	28.0%	57.8%	41.5%	6.7%	25.2%
Years: 2	2.2%	24.7%	39.6%	20.0%	2.0%
Years: 3	0.2%	3.2%	7.5%	20.0%	0.2%
Years: 4	0.0%	3.2%	9.4%	53.3%	0.0%

Table 6: Summary of infrastructure, techniques and period of activity for the different campaigns targeting Moneros grouped by profit.

(4) It is common to see campaigns mining in various pools. We observe that the most popular mining pools are *crypto-pool*, *dwarf-pool* and *minexmr*. We show that a large number of samples mine to *minergate*, an opaque mining pool for which there is no publicly-available information about the rewards received.

(5) When looking at the activity period, we find long-lasting campaigns – some of which are active at the time of writing. In particular, we can see multi-million campaigns operating for a continuous period of more than four years (see Top-10 in Table 4). This shows that AVs have not addressed this threat appropriately. We argue that crypto-mining malware has not been given enough attention by the industry and the research community and novel countermeasures are required as discussed in §6.

5 CASE STUDY

We next present a case study related to a high-profit campaign that has not been previously reported. We refer the readers to the extended version of the paper [32] for additional case studies. Fig. 2 presents an overview of how the campaign is structured. In the graph, nodes in blue represent wallets and nodes in light-green represent malware miners. Thus, various light-green nodes connected to a blue node represent a group of samples using the same wallet. Nodes in gray and pink represent the infrastructure of the campaign, with gray nodes portraying contacted domain servers and the pink ones the malware hosts. Finally, ancillary

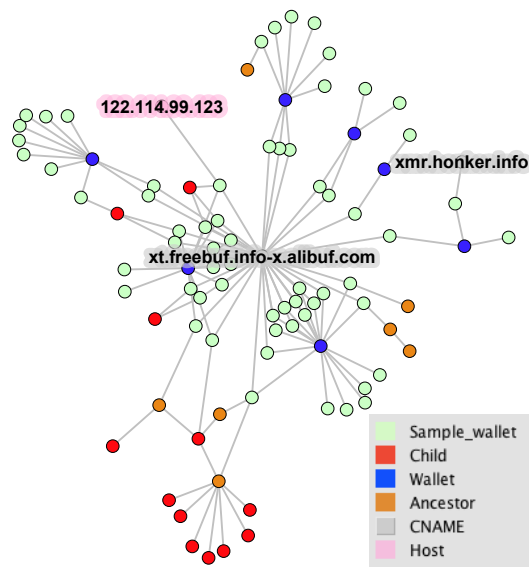


Figure 2: Overview of the analysis of our case study. The Freebuf campaign (C#627).

malware are depicted in red and orange. The edges represent the connections described in §3.5.

The most profitable Monero campaign (C#627 in Table 4) has mined more than 163K XMR in 3 years using 7 wallets. We have named it ‘Freebuf’ because the main grouping feature is the domain *xt.freebuf.info*, which is an alias (CNAME) of the *minexmr* pool.

Structure. Figure 2 shows how the campaign is structured. We observe that the aggregation is dominated by three grouping features: (i) same identifier, (ii) ancestors, and (iii) domain aliases (CNAMEs). Interestingly, the graph of this campaign reveals groups of samples with wallets that reach out to one another through paths traversing these three grouping features. In other words, the three grouping features are the key to map the structure of the campaign. In particular, there are two domains linked through a common wallet: *x.alibuf.com* and *xmr.honker.info*, which in turn are connected to *xt.freebuf.info*. Note that both have been aliases of *minexmr*, and *x.alibuf.com* has also been an alias of *crypto-pool*. We can observe that the backbone of the graph is established through connections of samples linked to CNAMEs. From there, there are other clusters that are linked via *same identifier*, and to those, there are some samples which are connected by common ancestors.

Payments. By analyzing the different payments received through time, we observe that before the update of the PoW in April 2018, this campaign was mining in various pools simultaneously. However, after the update all mining efforts were put into *minexmr*. In September 2018 we reported the wallets to the largest pools, resulting in two wallets being banned in October 2018. Upon request, one pool operator kindly provided us with statistics regarding the number of different IPs behind the wallets. The two banned wallets connected from 5,352 and 8,099 different IPs and had mined 362.6 and 1,283.7 XMR respectively. As a consequence of banning, we observe that the campaign operator decided to move their mining

efforts to another pool (*ppxmr*) – which indeed was used before the update on April 2018. We visualize these payments in [32]. We have seen that as a result of this intervention, together with the change in the PoW algorithm in October 2018, the payments received by the wallets associated to this campaign have been considerably reduced.

6 DISCUSSION

In this section, we analyze existing countermeasures, presenting their main challenges and weaknesses and looking at potential directions to address this challenge. We also discuss the limitations of our work.

Reporting illicit wallets. Reporting illicit wallets to the pools, while being a common practice [7], is not an effective countermeasure. First, it is costly and requires cooperation and coordination from all (or at least the main) pool operators. Second, criminals have developed mechanisms to bypass detection (e.g., using mining proxies). During our study, we have reported the illicit wallets we found to the largest pools, together with evidence of criminal behavior. There are non-cooperative pools that chose not to ban wallets; and those that are cooperative pools tend to err on the safe side. For example, the pool *minexmr* only blocks wallets with a large number of associated connections. Since criminals can use proxies to hide botnet-related mining activity, just relying on the number of connections proves ineffective.

Additionally, we found that successful campaigns mine in several pools. This makes them more resilient to take-down operations. Criminals respond to take-downs by changing the mining pool being used (as we have observed in our study) or by creating new wallets and setting proxies up [11].

Changes in the Proof-of-Work algorithm. ASIC-based mining uses customized hardware to compute faster the PoW algorithms [19]. Changes in these algorithms are intended to hinder ASIC-based mining, due to the cost of creating new hardware. This also requires updating the software, which is straightforward for benign miners. However, in the case of crypto-mining malware, each change in the algorithm requires botnet operators to update their bots. In turn PPI users have to purchase further installs to push the updated version of their miners.

We have monitored three changes in the PoW of Monero: April 6th, 2018, October 18th, 2018 and in March 9th, 2019. In each change, around 72%, 89% and 96% of the campaigns ceased their operations respectively. Although changes in the algorithms do not dismantle consolidated campaigns, they can dissuade new ones. Thus, a potential countermeasure against crypto-mining malware is to increment the frequency of such changes, and design these changes to not only be anti-ASIC, but also anti-botnet.

Third-party Infrastructure. We observe that most mining campaigns use third-party infrastructure such as PPI services or malware packers. To associate samples to such infrastructure we rely on IoCs gathered from public OSINT repositories, i.e. which have been reported previously. Thus, a limitation of our approach is that campaigns that use such *unknown* third-party infrastructure (e.g. custom packers being offered in underground markets) might be grouped together. Detecting such new third-party infrastructure would require applying other type of intelligence, which is out of

the scope of this paper. While this hinders our ability to account for the number of individual actors, our methodology still allows to understand the magnitude of the problem and devising novel mitigation strategies.

Quality of the ground-truth. As with other malware types, a countermeasure to this threat is to use updated AVs. Given the magnitude of the threat, AVs are required to have a large dataset of signatures. However, judging by the activity period of some of the campaigns we have seen that AV vendors have not been up to this task. This is a known limitation of the Threat Intelligence industry [21].

Likewise, our work relies on AV vendors to set our ground-truth. This might introduce both false positives (when legitimate samples are considered malware) and false negatives (when missing actual malware samples). The boundaries between ‘malicious’ and ‘legitimate’ samples are unclear, since a legitimate mining software turns malicious when used in infected computers (i.e., without user consent). Indeed, most AV classify mining tools as malware. We note that establishing an optimal trade-off between benign and malicious mining is not straightforward [8, 17]. We deal with FP by setting up a relatively high number of positive detections (i.e., 10 AVs). By doing this we are also introducing FN. In this paper, we err on minimizing the number of FP knowing that our findings have to be seen as an under-approximation of the current threat.

Finally, it is worth noting that we assume that a wallet is *illicit* when it seen together with a malware. However, not all the mining might be illicit. For example, the criminal could start mining from his own PC before/while she is mining illicitly. This is unlikely, but we would not have technical means to make this distinction should it were to happen. While some of the mining could be done licitly, we however argue that the revenues obtained from this activity revert on a criminal anyway. Thus, it can be used to fuel other illicit activities and the figures associated to this wallets are therefore relevant to our study.

Quality of the aggregation. Our methodology is used to group samples and wallets into campaigns. We leverage the data collected mostly for this purpose, but we also collect additional data to measure the magnitude of the problem. We use manual verification to evaluate that our heuristics are coherent. When campaigns relate to know mining botnets (e.g., Photominer, or Adylkuzz), we verify that the structure of the campaign matches with what has been reported about that botnet. When looking at novel campaigns, we verify that the resulting aggregation matches with the one provided by our heuristics. This can be seen from the case studies shown in §5. However, we acknowledge that our heuristics are not complete and can be subject to errors in the aggregation. This can lead to both campaigns that are under- or over-aggregated. In general, we have design our heuristics to be conservative to avoid over-aggregation. For example, we aggregate based on the full in-the-wild URL (rather than by domain name) to avoid aggregating campaigns that use common third-party infrastructure like Amazon WS or GitHub.

Anti-analysis techniques. Our study inherits the limitations of malware analysis. On the one hand, malware uses obfuscation to thwart static analysis, e.g. by using packers and crypters. We have partially addressed this by looking both at the usage of known packers and also at the entropy of the binaries. On the other hand,

malware uses evasion techniques to thwart dynamic analysis, (e.g., sandbox detection [28]). To partially prevent this, we leverage the use of various sanboxes (e.g. Virus Total and Hybrid Analysis). Also we have attempted to de-anonymize domain aliases that masked connections to mining pools (c.f., §3.5).

Mining-tailored solutions. Miners have a distinctive CPU usage, which allow for anomaly detection based on either modeling the CPU usage [17] or on instrumenting web browsers to detect suspicious activity [13, 45]. However, since the malware controls the infected computer, it might evade local defenses (e.g., by tampering with the CPU monitoring module). Other works propose to monitor the CPU usage of a computer from a hypervisor [39]. However, this approach focuses on protecting cloud providers, and it is not applicable to end-users. An alternative is to look at the energy consumption fingerprint. While power-aware anomaly detection systems have been proposed in other domains [14, 38], we are not aware of a solution tailored to crypto-mining malware for general-purpose computers. We position that these solutions could be deployed by electric-companies to end-users with smart-meters.

7 RELATED WORK

Illicit mining has been a threat since the emergence of Bitcoin in 2009. However, it has not been properly addressed in academia until recently. The first analysis of crypto-mining malware was published in 2014 by Huang et al. [9]. Authors analyzed botnets and campaigns mining bitcoins. However, mining non-ASIC resistant currencies (such as Bitcoin) is no longer profitable without dedicated hardware. Thus, most of the illicit mining focuses on Monero nowadays [7, 8, 17]. Recent works analyzed web-based mining, both as an alternative to advertisements to monetize web content [31, 35] and as cryptojacking [8, 17, 36]. To distinguish cryptojacking from benign mining, it is important to properly identify user consent. One approach is to search for keywords indicating mining activity [8]. This approach misses informed consent acquired by other means, such as images or additional documents. Thus, some works also look for AuthedMine scripts, which require explicit action from users to start mining [13, 17]. In our work, we rely on AV reports and other heuristics to classify binaries into malware or goodware.

Previous works are characterized by the simplicity in which they aggregate campaigns. In particular, related works mostly look at mining pool identifiers (e.g., wallets) alone [8, 13]. However, criminals use concurrent miners with different identifiers to retake operations when wallets are banned [11]. Konoth et al. includes information about the servers when performing the aggregation [17]. Unfortunately, this does not scale as it requires manual efforts vetting the code of the scripts (i.e., to get the verification code). In their analysis of Bitcoin, Huang et al. use information gathered from the Blockchain to aggregate campaigns [9]. However, this approach is not valid for cryptocurrencies that obfuscate transactions (e.g., Monero or Zcash).

The overall earnings obtained from malicious mining have increased in the last years. For example, Konoth et al. discovered 1,735 domains, estimating overall revenues of \$188,878 per month. In parallel work, Hong et al. detected 2,770 domains, estimating overall revenues of \$1.7M. However, estimations obtained from
















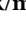
	Currency	Size		Profits
		Analyzed	Detected	
[9](2014)	BTC	Unknown	2,000 	14,979 BTC
[35](2018)	XMR	10M 	2,287 	1,271 XMR/m
[8](2018)	XMR	548,624 	2.2k 	7,692.30 XMR
[17](2018)	XMR	991,513 	1.7k 	746.55 XMR/m
[31](2018)	XMR	3M 	108K 	N/A
[29](2018)	XMR	1M 	2,500 	N/A
Our work	Various	4.4M 	1M 	741K XMR 13.7K XMR/m

Table 7: Summary of related work, where the last row shows our measurement. Legend:  stands for Web-based mining,  for Web-based cryptojacking,  for binary-based crypto-mining malware. XMR/m means monthly moneros

web-browser cryptojacking are not reliable. This is because analyzing profits from web activity relies on estimates of i) the number of monthly visitors, ii) the time spent by each visitor on average, and iii) the type of device they use. Instead, we are able to get wallets used by the malware and the payments given by the pools as a reward. This allows us to analyze not only the earnings of each wallet, but also the pools used for mining and the exact dates of the payments. Our findings increase the understanding of this threat. In particular, we estimate that earnings are — at least — 58 million USD obtained in 4.5 years of operation (more than 1M/month). Table 7 summarizes the related works and compares each of the measurements.

Unlike other type of malware threats, crypto-mining malware is well bounded by a unique characterizing behavior, i.e.: the use of mining activity. Thus, simple pattern matching (like the one done by YARA) has proven to be effective. As malware increases in sophistication (see discussion in §6), related works leverage machine learning approaches to systematically identify specific threats [37]. However, this is out of the scope of our work.

8 CONCLUSION

In this paper, we have presented a longitudinal large-scale measurement study of crypto-mining malware — analyzing the underlying infrastructure relied upon for over a decade. We show that Monero is currently the preferred currency used by criminals, who have obtained massive earnings. Our profit analysis on Monero reveals that a small number of actors hold sway the mining illicit business. We also show that campaigns that use third-party infrastructure (typically rented in underground marketplaces) are more successful. However, this is not always the case. Some of the most profitable campaigns rely on complex infrastructure that also uses general-purpose botnets to run mining operations without using third-party infrastructure. Here, we have discovered novel malware campaigns that were previously unknown to the community and we have present technical details about the way this ecosystem operates. We have leveraged the insights obtained to discuss countermeasures. Due to the need of updating mining software, our findings suggest that regular changes in the PoW algorithm discourage criminals. As future work, we plan to extend our study to

emerging campaigns as oppose to looking at the most consolidated ones, and with an extend array of heuristics.

ACKNOWLEDGMENTS

The authors would like to thank our shepherd, professor Ben Y. Zhao, for the assistance in the process of improving the quality of this paper, and the anonymous reviewers for their positive and helpful comments. We thank PaloAlto for sharing data feeds related to cryptomining malware. We also thank VirusTotal for granting access to its private API and the Cambridge Cybercrime Centre for sharing the CrimeBB dataset. Special thanks to Alexander Vetterl and Daniel Thomas for their assistance and valuable feedback.

This work is partially supported by MINECO (grant TIN2016-79095-C2-2-R), and by the Comunidad de Madrid (P2018/TCS-4566, co-financed by European Structural Funds ESF and FEDER). The opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect those of any of the funders.

REFERENCES

- [1] Christiaan Beek, Taylor Dunton, Steve Grobman, Mary Karlton, Niamh Minihaane, Chris Palm, Eric Peterson, Raj Samani, Craig Schumgar, ReseAnne Sims, Dan Sommer, and Bing Sun. 2018. McAfee Threat Report.
- [2] Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. 2011. Measuring pay-per-install: the commoditization of malware distribution. In *20th USENIX Security Symposium*. USENIX Association, San Francisco, CA, USA, 13–13.
- [3] Checkpoint. 2018. Jenkins miner: one of the biggest mining operations ever discovered. <https://perma.cc/SVN4-C5B4>
- [4] Cyren. 2014. F-prot Antivirus. <http://www.f-prot.com>
- [5] Shayan Eskandari, Andreas Leoutsarakos, Troy Mursch, and Jeremy Clark. 2018. A first look at browser-based Cryptojacking. In *IEEE Security & Privacy on the Blockchain*. IEEE, London, UK, 58–66.
- [6] Daniel Goldberg. 2016. Guardicore: The Photominer campaign. <https://perma.cc/JE3Y-F42L>
- [7] Josh Grunzweig. 2018. Palo Alto Networks: The Rise of the Cryptocurrency Miners. <https://perma.cc/4VZL-J45Q>
- [8] Geng Hong, Zheming Yang, Sen Yang, Lei Zhang, Yuhong Nan, Zhibo Zhang, Min Yang, Yuan Zhang, Zhiyun Qian, and Haixin Duan. 2018. How You Get Shot in the Back: A Systematical Study about Cryptojacking in the Real World. In *Proceedings of the 25th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, Toronto, Canada, 1701–1713.
- [9] Danny Yuxing Huang, Hitesh Dharmdasani, Sarah Meiklejohn, Vacha Dave, Chris Grier, Damon McCoy, Stefan Savage, Nicholas Weaver, Alex C Snoeren, and Kirill Levchenko. 2014. Bitcoin: Monetizing Stolen Cycles. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. Citeseer, Internet Society, San Diego, CA, USA, 1–16.
- [10] Kafeine. 2017. Proofpoint: Adylkuzz cryptocurrency mining malware Spreading for weeks via EternalBlue/DoublePulsar. <https://perma.cc/3V7G-CDEN>
- [11] Kafeine. 2018. Proofpoint: Smominru Monero mining botnet making millions for operators. <https://perma.cc/V5UR-TDLU>
- [12] Alex Kantchelian, Michael Carl Tschantz, Sadia Afroz, Brad Miller, Vaishaal Shankar, Rekha Bachwani, Anthony D Joseph, and J Doug Tygar. 2015. Better malware ground truth: Techniques for weighting anti-virus vendor labels. In *Proceedings of the 8th ACM Workshop on Artificial Intelligence and Security*. ACM, Denver, Colorado, USA, 45–56.
- [13] Amin Kharraz, Zane Ma, Paul Murley, Charles Lever, Joshua Mason, Andrew Miller, Nikita Borisov, Manos Antonakakis, and Michael Bailey. 2019. Outguard: Detecting In-Browser Covert Cryptocurrency Mining in the Wild. In *Proceedings of The Web Conference (WWW)*. ACM, San Francisco, USA, 840–852.
- [14] Hahnsang Kim, Joshua Smith, and Kang G Shin. 2008. Detecting energy-greedy anomalies and mobile malware variants. In *Proceedings of the 6th International Conference on Mobile Systems, Applications, and Services*. ACM, Breckenridge, CO, USA, 239–252.
- [15] Dhilung Kirat, Giovanni Vigna, and Christopher Kruegel. 2014. BareCloud: Bare-metal Analysis-based Evasive Malware Detection. In *23th USENIX Security Symposium*. USENIX Association, San Diego, CA, USA, 287–301.
- [16] Clemens Kolbitsch, Engin Kirda, and Christopher Kruegel. 2011. The power of procrastination: detection and mitigation of execution-stalling malicious code. In *Proceedings of the 18th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, Chicago, Illinois, USA, 285–296.

- [17] Radhesh Krishnan Konoth, Emanuele Vineti, Veelasha Moonsamy, Martina Lindorfer, Christopher Kruegel, Herbert Bos, and Giovanni Vigna. 2018. MineSweeper: An In-depth Look into Drive-by Cryptocurrency Mining and Its Defense. In *Proceedings of the 25th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, Toronto, Canada, 1714–1730.
- [18] Jesse Kornblum. 2006. Identifying almost identical files using context triggered piecewise hashing. *Digital investigation* 3 (2006), 91–97.
- [19] Kay Kurokawa. 2018. Forging for ASIC resistance: A Monero case study. <https://perma.cc/5JL6-RPPS>
- [20] Chaz Lever, Platon Kotzias, Davide Balzarotti, Juan Caballero, and Manos Antonakakis. 2017. A Lustrum of malware network communication: Evolution and insights. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, San Jose, CA, USA, 788–804.
- [21] Vector Guo Li, Matthew Dunn, Paul Pearce, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2019. Reading the Tea leaves: A Comparative Analysis of Threat Intelligence. In *28th USENIX Security Symposium*. USENIX Association, Santa Clara, CA, 851–867. <https://www.usenix.org/conference/usenixsecurity19/presentation/li>
- [22] Yuping Li, Sathya Chandran Sundaramurthy, Alexandru G Bardas, Xinming Ou, Doina Caragea, Xin Hu, and Jiyong Jang. 2015. Experimental study of fuzzy hashing in malware clustering analysis. In *8th Workshop on Cyber Security Experimentation and Test (CSET 15)*, Vol. 5. USENIX Association, Washington, D.C, USA, 52.
- [23] David Liebenberg. 2018. Cisco Talos: Rocke, the champion of Monero Miners. <https://perma.cc/ZH4B-DBG3>
- [24] David Liebenberg, Charles McFarland, Michelle Martinez, Cruz Jerome, Fred Gutierrez, Anthony Giandomenico, Terrance DeJesus, Josh Grunzweig, Michelle Martinez, Andrew Brandt, Neil Jenkins, and Scott Scher. 2018. The Illicit Cryptocurrency Mining Threat. <https://perma.cc/Z5ZH-H5TG>
- [25] Martina Lindorfer, Matthias Neugschwandtner, and Christian Platzer. 2015. Marvin: Efficient and comprehensive mobile app classification through static and dynamic analysis. In *Proceedings of the IEEE 39th Annual Computer Software and Applications Conference (COMPSAC)*, Vol. 2. IEEE, Taichung, Taiwan, 422–433.
- [26] Martina Lindorfer, Matthias Neugschwandtner, Lukas Weichselbaum, Yanick Fratantonio, Victor Van Der Veen, and Christian Platzer. 2014. Andrubiis–1,000,000 apps later: A view on current Android malware behaviors. In *Proceedings of the 3th International Workshop on Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*. IEEE, Wroclaw, Poland, 3–17.
- [27] Chad McMillan and Jason Garman. 2011. System and method for determining data entropy to identify malware. US Patent 8,069,484.
- [28] Najmeh Miramirkhani, Mahathi Priya Appini, Nick Nikiforakis, and Michalis Polychronakis. 2017. Spotless sandboxes: Adapting malware analysis systems using wear-and-tear artifacts. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, San Jose, CA, USA, 1009–1024.
- [29] Marius Musch, Christian Wressnegger, Martin Johns, and Konrad Rieck. 2018. Web-based Cryptojacking in the Wild. *arXiv preprint arXiv:1808.09474* (2018).
- [30] Andy Norton. 2018. Crypto-malware - A Look At The Latest Malware Threat. <https://perma.cc/6K3K-E632>
- [31] Panagiotis Papadopoulos, Panagiotis Ilia, and Evangelos P Markatos. 2018. Truth in Web Mining: Measuring the Profitability and Cost of Cryptominers as a Web Monetization Model. *arXiv preprint arXiv:1806.01994* (2018).
- [32] Sergio Pastrana and Guillermo Suarez-Tangil. 2019. A First Look at the Crypto-Mining Malware Ecosystem: A Decade of Unrestricted Wealth (Extended Version). *arXiv preprint arXiv:1901.00846* (2019).
- [33] Sergio Pastrana, Daniel R. Thomas, Alice Hutchings, and Richard Clayton. 2018. CrimeBB: Enabling Cybercrime Research on Underground Forums at Scale. In *Proceedings of The Web Conference (WWW)*. ACM, Lyon, France, 1845–1854.
- [34] Ruben Recabarren and Bogdan Carbunar. 2017. Hardening Stratum, the Bitcoin pool mining protocol. *Proceedings on Privacy Enhancing Technologies* 2017, 3 (2017), 57–74.
- [35] Jan R uth, Torsten Zimmermann, Konrad Wolsing, and Oliver Hohlfeld. 2018. Digging into Browser-based Crypto Mining. In *Proceedings of the 2018 Internet Measurement Conference (IMC'18)*. ACM, Boston, MA, USA, 70–76.
- [36] Muhammad Saad, Aminollah Khormali, and Aziz Mohaisen. 2018. End-to-End Analysis of In-Browser Cryptojacking. *arXiv preprint arXiv:1809.02152* (2018).
- [37] Guillermo Suarez-Tangil, Santanu Kumar Dash, Mansour Ahmadi, Johannes Kinder, Giorgio Giacinto, and Lorenzo Cavallaro. 2017. Droidsieve: Fast and accurate classification of obfuscated android malware. In *Proceedings of the 7th ACM Conference on Data and Application Security and Privacy (CODASPY'17)*. ACM, Scottsdale, Arizona, USA, 309–320.
- [38] Guillermo Suarez-Tangil, Juan E Tapiador, Pedro Peris-Lopez, and Sergio Pastrana. 2015. Power-aware anomaly detection in smartphones: An analysis of on-platform versus externalized operation. *Pervasive and Mobile Computing* 18 (2015), 137–151.
- [39] Rashid Tahir, Muhammad Huzaifa, Anupam Das, Mohammad Ahmad, Carl Gunter, Fareed Zaffar, Matthew Caesar, and Nikita Borisov. 2017. Mining on someone else's dime: Mitigating covert mining operations in clouds and enterprises. In *Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID)*. Springer International Publishing, Cham, 287–310.
- [40] Threat Intelligence Team. 2018. Avast: Greedy cybercriminals host malware on GitHub. <https://perma.cc/DV7R-Q7CQ>
- [41] Kurt Thomas, Danny Yuxing, Huang David, Wang Elie, Bursztein Chris Grier, Thomas J. Holt, Christopher Kruegel, Damon Mccoy, Stefan Savage, and Giovanni Vigna. 2015. Framing dependencies introduced by underground commoditization. In *In Proceedings of the Workshop on Economics of Information Security (WEIS)*. Delft, Netherlands.
- [42] Xabier Ugarte-Pedrero, Igor Santos, Borja Sanz, Carlos Laorden, and Pablo Garcia Bringas. 2012. Countering entropy measure attacks on packed software detection. In *Proceedings of the IEEE Consumer Communications and Networking Conference (CCNC)*. IEEE, Las Vegas, NV, USA, 164–168.
- [43] Ashwin Vamshi. 2018. Netskope: Technical analysis of Xbooster parasitic Monero miner. <https://perma.cc/8RZG-5QBS>
- [44] Rolf Van Wegberg, Samaneh Tajalizadehkhoob, Kyle Soska, Ugur Akyazi, Carlos Hernandez Ganan, Bram Klievink, Nicolas Christin, and Michel Van Eeten. 2018. Plug and prey? Measuring the commoditization of cybercrime via online anonymous markets. In *27th USENIX Security Symposium*. USENIX Association, Baltimore, MD, USA, 1009–1026.
- [45] Wenhao Wang, Benjamin Ferrell, Xiaoyang Xu, Kevin W Hamlen, and Shuang Hao. 2018. SEISMIC:Secure In-lined Script Monitors for Interrupting Cryptojacks. In *European Symposium on Research in Computer Security (ESORICS)*. Springer International Publishing, Barcelona, Spain, 122–142.
- [46] Webroot. 2018. Webroot Threat Report: Mid-Year Update. <https://perma.cc/3M2Z-Q76Y>

APPENDIX

Ethical issues

Most of the data collected is publicly available. However, both Palo Alto and Virus Total shared some non-public information with us and we requested their permission to use it for this paper. Another ethical concern relates to the implications of reporting any misuse activity to the pools. In particular, we have reported evidence of wallets seen in crypto-mining malware. Our actions might produce an intervention over the reported users due to criminal activity. This entails potential ethical implications when non-criminal wallets are mistakenly banned. We have taken due precautions to guarantee that we only report wallets of samples used by malware as discussed in the paper. Thus, our study has been approved by the designated ethics officer at the Research Ethics Board (REB) of our institution.

In addition to our precautions, we have provided the pools with accompanying evidences proving illicit activity, including a pointer to the Virus Total report. However, despite our involvement, the final decision of banning the wallets relies on the pool operators. These operators have additional insights about the *modus operandi* of their users (e.g., the number of IP addresses that are currently mining with a wallet) that can be used to further corroborate any type of misuse. In fact, it is our understanding that the pools that took actions against some of the reported users based their decision solely on the number of connections per wallet. Multiple connections from the same wallet evidences the use of a botnet, which it is against the terms and conditions of some of the pools. However, we note that a botnet can always hide behind a single IP addressed using proxies. Also, we have observed that the pools do not proactively ban wallets that display botnet-like activity. We discuss the ethical implications of the banning policies of the pools in §6.