Botcoin: Monetizing Stolen Cycles

Danny Yuxing Huang, Hitesh Dharmdasani†, Sarah Meiklejohn
Vacha Dave, Chris Grier*, Damon McCoy†, Stefan Savage, Nicholas Weaver*
Alex C. Snoeren and Kirill Levchenko

UC San Diego †George Mason University *International Computer Science Institute

***Abstract*—At the current stratospheric value of Bitcoin, miners with access to significant computational horsepower are literally printing money. For example, the first operator of a USD $1,500 custom ASIC mining platform claims to have recouped his investment in less than three weeks in early February 2013, and the value of a bitcoin has more than tripled since then. Not surprisingly, cybercriminals have also been drawn to this potentially lucrative endeavor, but instead are leveraging the resources available to them: stolen CPU hours in the form of botnets. We conduct the first comprehensive study of Bitcoin mining malware, and describe the infrastructure and mechanism deployed by several major players. By carefully reconstructing the Bitcoin transaction records, we are able to deduce the amount of money a number of mining botnets have made.**

I. Introduction

Compromised end-user PCs—bots—find many uses in the modern cyber criminal ecosystem; e.g., they are used to send spam, commit click fraud, carry out denial-of-service attacks, and steal personal user data. To a botmaster, bots are a resource, and each of these activities are a means of extracting value from this resource. How a botmaster chooses to monetize these resources is of considerable interest to the security community, as the profitability of these schemes that drives demand for compromised PCs. The more money a botmaster can extract from a botnet, the greater the demand for bots. This demand plays a powerful role in shaping the computer security landscape, driving the evolution of malware and new attacks.

This paper examines the practice of using compromised PCs to mine Bitcoin, a monetization scheme that has recently gained popularity. Bitcoin is a decentralized virtual currency that can be generated computationally through a process called *mining*, which requires repeatedly computing the SHA-256 cryptographic hash function over a large range of values. Although initially regarded as a novelty, interest in Bitcoin has exploded. The value of a bitcoin, fueled largely by speculation, has climbed from USD $5 in May 2012 to just over $1,100 at the end of November 2013, catching the eye of regulators [7], the courts [28], the popular media [19], [30], [32], and botmasters alike. The mining process is essentially a state-space search that can be conducted in parallel, making it an excellent candidate for distribution across a set of compromised PCs, allowing an enterprising botmaster to generate bitcoins at scale.

A particularly appealing feature of Bitcoin mining is that it requires little additional investment on the part of the botmaster. Each monetization scheme has resource and infrastructure requirements that must be balanced against the expected revenue from the activity. To monetize spam, for example, a spammer needs an audience and a product to sell. Developing a viable advertising channel requires additional investment, while operating an online store and accepting customer payment requires significant infrastructure, today handled by affiliate programs paying the spammer a commission. Even extracting money from bank accounts (using stolen user credentials) requires a network of money mules and carries with it substantial criminal liability.

In contrast, Bitcoin mining can be carried out without any additional infrastructure. Compared to existing monetization schemes, the cost of Bitcoin mining is thus very low. These costs are not zero, and must be balanced by revenue for Bitcoin mining on a botnet to be profitable. Because Bitcoin mining is almost entirely computational, however, it is unlikely to interfere with other monetization activities, allowing a botmaster to add Bitcoin mining to the set of revenue-generating activities carried out by a botnet without adverse impact. On the other hand, the resulting high CPU utilization may very well increase the likelihood that a bot's true owner will detect the compromise. Understanding the balance of added cost and risk versus potential revenue from Bitcoin mining is the motivation for our work.

Toward this end, we collect Bitcoin mining malware from multiple sources, including security industry malware databases ThreatExpert and Emerging Threats. For each executable, we identify how it mines bitcoins—using both sandboxed execution and binary analysis—and extract the botmaster's mining credentials. Using other data sources (public data published via the Web, communication with mining pool operators, the public blockchain, leaked data, passive DNS, and proving the mining proxies) we identify, where possible, the infrastructure used by each operation, when each operation was active, and how much each earned, providing a comprehensive view of existing botnet Bitcoin mining activity.

In brief, the contributions of this paper are:

- ❖ We identify malware engaged in Bitcoin mining and report how it operates.

NDSS '14, 23-26 February 2014, San Diego, CA, USA
 ISBN 1-891562-35-5
http://dx.doi.org/doi-info-to-be-provided-later

- ❖ We present nine case studies of botnet Bitcoin mining operations ranging in sophistication. Where possible, we estimate the infected population and its geographic distribution.
- ❖ We estimate the total value extracted from compromised PCs through mining. In particular, we find that the compromised PCs used by the operations we identify have mined *at least* 4,500 bitcoins.
- ❖ We discuss the profitability of Bitcoin mining on botnets and conclude that the potential revenue from Bitcoin mining alone is unlikely to cover the costs of a botnet, but may be attractive as a secondary activity for large botnets with already established primary monetization schemes.

The rest of this paper is organized as follows. Section II surveys related work, while Section III provides the technical background on Bitcoin and Bitcoin mining necessary for the remainder of the paper. In Section IV we describe our measurement methodology before presenting our results in Section V. We discuss the economics of malware mining in Section VI before concluding in Section VII.

II. Related Work

There are three main areas of research related to our investigation of botnet mining malware. The first performs analysis of the transactions in the Bitcoin network [20], [25], [26], to measure activity and test the limits of the anonymity it provides. We apply some of the same methods employed by these studies to trace payouts to botmasters. Another recent study performs an analysis of Silk Road [3], an underground drug market that—until its shutdown in October 2013—accepted payments only in bitcoin.

Two recent studies focus specifically on Bitcoin mining [6], [18]. Both studies examine how Bitcoin mining can be "gamed" by an appropriately powerful adversary, and find that sufficiently motivated adversaries can disrupt the Bitcoin economy. Our study, while also focused on mining, is complementary to these, as it focuses exclusively on the phenomenon of using compromised hosts to mine bitcoins.

There are many studies that explore the dynamics of different methods of profiting from malware; recent examples focus on pay-per-install [2], fake anti-virus [29], pharmaceutical spam [16], and click fraud [21], among others. While these related efforts characterize the dynamics of malware and its profit motives, we believe ours is the first study to explore the dynamics of Bitcoin mining malware in particular.

III. Background

Bitcoin is a decentralized peer-to-peer virtual currency based on a proposal published in 2008 by an unknown author under the pseudonym Satoshi Nakamoto [22]. Bitcoin is not backed by any government or physical commodity: it is a purely virtual currency. Any user can generate bitcoins through a computational process called *mining*. Bitcoin itself is simply a global, public ledger of balances that associates a *wallet*, identified by a *wallet address* (or simply *address*), with each balance. A wallet address is the hash of a public key in an elliptic curve digital signature scheme; the owner of the wallet has the corresponding private key, allowing her and only her to sign *transactions* transferring money out of the associated wallet address and into another. All Bitcoin transactions are logged in a public ledger known as the *blockchain*. The blockchain is maintained by a peer-to-peer network. The peer-to-peer network appends only valid transactions to the blockchain, meaning the sending wallet addresses must contain sufficient funds and the transaction must be properly signed by the wallet owner(s).

A. Bitcoin Mining

The security of Bitcoin depends upon the integrity of the blockchain. While individual transactions can be validated simply by reading the chain, preventing double-spending and other misbehavior requires ensuring that there is only one append-only ledger. The integrity of the blockchain is ensured through the mining process, which serves a dual role: it maintains the blockchain and requires participants to execute a proof-of-work algorithm in order to generate new bitcoins.

A Bitcoin miner[1] groups new valid transactions received via the peer-to-peer network into blocks consisting of a set of transactions and a header containing a hash of the previous block and a nonce. The miner then computes a SHA-256 hash value of this block. If the binary representation of the hash value contains a sufficient number of leading zeroes, the miner disseminates the newly mined block to other users via the peer-to-peer network; each peer verifies the validity of the new block by computing the SHA-256 hash value of the block and checking that it contains at least the required number of leading zeroes. This acts to confirm the transactions in the block, protecting them from future tampering. The new block also contains a special transaction, the *coinbase*—which serves as an additional nonce—a comment field inserted by the miner, and a special transaction that pays all transaction fees and the block reward (initially 50 BTC, currently 25 BTC, and halving approximately every 4 years) to the miner's wallet.

The common case, however, is that the miner's choice of nonce leads to a SHA-256 hash value without a sufficient number of leading zeroes. The miner then repeats the process with a different nonce value until some miner finds a block with the proper hash value and publishes it via the peer-to-peer network. After a new block's discovery, all miners remove the newly confirmed transactions from their pool of work and continue the process with another group of transactions. The number of leading zeroes required to mine a block controls the difficulty of mining bitcoins, and is recomputed by global consensus every 1024th block to maintain an average mining rate of one block every ten minutes.

The random nature and fixed block-creation rate make mining competitive: every miner's chance of discovering a valid block is proportional to both the number of SHA-256 calculations (the *hash rate*) it can perform per second (usually measured in millions of hashes per second (MH/s), billions (GH/s), or trillions (TH/s)) and the hash rate of the Bitcoin network as a whole. An average desktop PC can perform anywhere from 2 to 10 MH/s, while a dedicated ASIC mining system can reach 500 GH/s or more. On November 30, 2013,

[1] In the interest of space, we have simplified the technical details of the block mining process that are not relevant to the technical content of this paper. The interested reader may consult Nakamoto's original paper [22] and other publicly available information on Bitcoin.

the Bitcoin network's hash rate was approximately 6,000 TH/s, which implies that a single 10-MH/s PC would have expected to receive less than 0.0000002% of all Bitcoins produced globally during the period it mined.

B. Pooled Mining

At today's difficulty level, a desktop PC mining at 10 MH/s can expect to mine 425 years before finding a winning block. Even with a top-of-the-line GPU capable of 500–1,000 MH/s, or one block every 4 years at the current difficulty level, mining becomes a lottery. To overcome this uncertainty, a miner can join a *mining pool*, which combines the mining power of a large number of individual miners and pays a small amount for each unit of work performed toward mining a block. Essentially, by parallelizing the search for a winning block, a pool can be thought of as buying multiple lottery tickets for any given drawing. With a typical mining pool, each miner is paid in proportion to his hashing power, but the income is significantly steadier due to the decrease in variance in the expected time required (for some member of the pool) to successfully mine a block.

In pooled mining, the pool server manages all pending transactions on the peer-to-peer network, and provides a starting point (the hash of all pending transactions) to each worker upon request. The worker then iterates the 32-bit nonce and, if the resulting hash has a sufficient number of leading zeroes, the worker reports the resulting block header back to the pool server. The pool server then checks if the resulting hash matches the difficulty and, if so, publishes the associated block. Otherwise, it simply uses the partial collision as proof that the worker is performing the necessary computation. The pool server then pays the workers based on their relative contribution to the pool server's computational effort.

This communication uses a simple a simple HTTP-based RPC protocol, called the *getwork* protocol after its main procedure call. An individual may mine using multiple machines; each host, called a *worker*, connects to the pool server, queries it for work, performs the necessary SHA-256 computations, returns the results to the server in case of a partial hash collision, and requests the next unit of work.

Most pools require miners to register a user name, password, and associate a payout address where the pool server sends the user's share, as the pool will periodically pay the miners based on the miner's contribution (often using a pool-specific payment formula). Some pools also support *pseudonymous mining*. In this case, the worker provides a wallet address rather than a user name to the pool server, and all earnings are sent to the specified address.

C. Botnet Mining

The ability to turn computation directly into money has given botmasters a new way to monetize the untapped computational capacity of their compromised hosts, while the rising value of Bitcoin has given them a strong incentive to do so. The first Bitcoin mining malware was observed in the wild in June 2011 [23]; since then, numerous families of malware have taken up Bitcoin mining.

The first family we identified with mining capability was NGRBot, a malware kit that has been available for several years. NGRBot is a generic malware platform with many

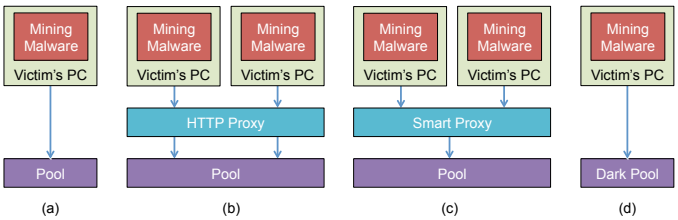

Fig. 1: Different ways in which mining malware connects to mining pools: (a) directly to a pool, (b) via an HTTP proxy, (c) via a smart proxy, and (d) directly to a dark pool.

different capabilities, such as stealing personal information, automatic spreading on USB and network disks, and DDoS. Instances of NGRBot have continued to mine, and recently an NGRBot variant spread through Skype messages was seen mining [1]. In mid-2012 several news stories documented ZeroAccess performing both bitcoin mining and click fraud at large scale [33]. Shortly after ZeroAccess and NGRBot, many other families of malware began to appear that installed (or dropped) bitcoin mining functionality.

We observe three distinct botnet mining pool structures in the wild, summarized in Figure 1.

Direct pool mining. At its simplest, mining with a botnet is no different from mining using one's own hardware. One of the more popular techniques for botnet-based Bitcoin mining is to simply distribute a mining executable (such as cgminer.exe or bfgminer.exe) inside a wrapper script that specifies all the parameters required to mine. This removes any cost associated with developing or modifying botnet software and is popular with Trojans distributed as pirated software. An example is the FeodalCash family of botnets (Section V-D) that mine directly at Eligius, a public pool.

A botmaster simply needs to specifying a mining pool and provide his own credentials. Each compromised PC will connect directly to the mining pool as a worker and start doing work on behalf of the pool; the pool will direct payments to the botmaster's account. We call this approach *direct pool mining* or simply *direct mining*.

Mining pool operators can easily detect direct mining, as it involves a large number of hosts, all using the same account, with each host providing very little CPU power for the mining task. Once detected, most pool operators will ban such users.[2] Once banned, the botnet becomes useless if there is no way to change the mining pool or credentials used by the bot.

Proxied pool mining. To overcome some of the drawbacks of mining directly, a botmaster can proxy connections to the pool through a server he controls, a mode we call *proxied pool mining* or simply *proxied mining*. Since the getwork protocol uses HTTP as a transport, a botmaster can simply employ an HTTP proxy (e.g., nginx). Using a proxy has two advantages. First, it hides the IP addresses of the bots: all connections appear to come from the proxy itself, making the botnet seem more like a single, powerful miner[3]. Using a proxy

[2]One pool operator reported having to relent after his pool servers came under DDoS attack from the botnet.

[3]It is still possibly detectable, as the time between sending a getwork request and providing a corresponding proof of work could be longer than then the gap produced by a single miner. The number of getwork calls is also unchanged.

also provides a level of indirection, allowing the botmaster to switch to new credentials or mining pools if banned.

Alternatively, the botmaster can design a more sophisticated *smart proxy* that does more than blindly pass through getwork requests. In this configuration, a proxy operated by the botmaster requests work from a pool as a normal worker, but then splits the work into smaller units provided to the bots. This architecture requires modifying existing mining pool software to support such operation, an additional investment. The Fareit botnet, however, uses a form of smart proxying. The server to which the bots connect operates as a mining pool server for the bots, but appears as a single worker to a P2Pool mining pool; more detail is provided in Section V-E.

The downside of either form of proxy mining is that it requires additional infrastructure—the proxy. Several mining operations we observe use this mechanism, such as DLoad.asia and ZeroAccess (Sections V-A and V-B).

Dark pool mining. The final option is for the botmaster to maintain his own pool server. In this mode, which we call *dark pool mining* or simply *dark mining*, bots connect to a mining pool controlled by the botmaster. In this configuration, the botmaster's dark pool must participate in the Bitcoin peer-to-peer network. In addition to the infrastructure investment in the pool server, the botmaster loses the consistency of payouts provided by a (larger) pool. The botnet now only generates revenue if it finds a block itself. A botnet of 10,000 compromised desktop PCs each capable of 10 MH/s running continuously mines one block every 16 days on average, as of August 2013.

IV. Methodology

Our goal is to identify Bitcoin mining malware, the size of the infected population, and how much value has been extracted through mining. The latter frequently requires us to understand a good bit about the botnet's mode of operation, including its mining pool credentials and payout wallet. Here, we describe the methodology we use to collect the data that underlies the analysis presented in Section V.

A. Identifying Mining Malware

To our knowledge, all malware currently engaged in Bitcoin mining uses the HTTP-based getwork protocol supported by existing mining pools. We therefore rely on this signal as our primary means of identifying mining malware: we use mining protocol traffic in network traces of a malware binary's execution as evidence that it is engaged in Bitcoin mining.

To obtain network traffic of various malware, we execute the binaries in our own malware execution environment or rely upon data from public and private sandboxes, including ThreatExpert[4] and Emerging Threats[5]. Some environments also provide OS-level monitoring such as logs of registry keys changed and files modified. We manually assess if a sample is performing Bitcoin mining by inspecting the traffic and looking for evidence that a particular sample is requesting work from a Bitcoin pool server. Then, using traffic and OS-level logs we construct queries to identify additional samples with similar characteristics. In total we identify over 2,000 executables that connect to pools and mine bitcoins.

B. Extracting Mining Credentials

Most mining malware relies on generic, off-the-shelf mining clients to do the actual mining. The malware executes the client and provides the pool name and worker user name—mining credentials passed as parameters to the miner—on the command line.

Command-line arguments. In many cases, we can extract these command-line arguments directly from the packaged binary statically. In other cases, we extract the mining credentials from the process execution environment; an example is the BMControl malware (Section V-C), from which we extract the usernames from the memory dump.

HTTP basic authentication. We can also extract the pool name and miner user name from the network trace of the malware. The getwork protocol relies on HTTP basic access authentication to provide the miner user name to the pool.[6] With HTTP basic authentication, a Base64-encoded user name and password are submitted in an HTTP header, making them easy to extract from a network trace. We use this method of extracting miner identifiers for binaries executed in third-party sandboxes.

Command-and-control channel. Some malware does not embed the pool or worker name into the binary. Instead, the mining credentials are obtained through a custom command-and-control channel. The Fareit botnet (Section V-E) uses the Dropbox and Pastebin Web services to disseminate mining credentials to bots. The contents of the Dropbox or Pastebin document are usually obfuscated using algorithms ranging from simple Base64 encoding to custom encoding schemes.

We manually reverse-engineer the malware to determine the technique used to obfuscate the data received through the command-and-control channel. For simple obfuscations, we can recreate the de-obfuscation algorithm and use it to continually retrieve the pool information and worker credentials. One example of this is the first version of the BMControl botnet that uses Pastebin to host Base64-encoded configuration information. The configuration includes the command-line parameters for the mining executable (in this case bfgminer[7]) as well as a list of the pools and worker credentials to use.

More complex obfuscation can be difficult to reverse-engineer; in this case we run the malware and take a memory snapshot after the malware has de-obfuscated the payload. An update to the BMControl botnet included a change in the obfuscation technique, so we use memory snapshots to capture the decoded payloads. The Fareit malware family also uses more substantial obfuscation, making memory snapshots a prudent technique for automatically decoding the configuration.

These techniques allow us to identify the mining credentials for all the samples of malware we find mining bitcoins. Based on the pools we observe the malware accessing, we find that 74% of the samples connect to well-known public pools

[4] http://www.threatexpert.com
[5] http://www.emergingthreats.net

[6] The password is ignored by all pools of which we are aware, since there is no benefit to doing work in another miner's name, nor is there any obvious harm to the miner in whose name the work is submitted.
[7] http://bfgminer.org/

(light pools), while the remaining 26% connect to unknown private pools (dark pools).

Pool operators. Finally, some user names and wallet addresses were provided to us by public pool operators as miners they believed to be using a botnet. We confirm their claims by locating the corresponding malware MD5 hashes (e.g. the gamer-targeting botnets in Section V-J).

C. Earnings

One of our main goals is to estimate the value extracted from compromised PCs through Bitcoin mining. To do so, we use a variety of techniques, described below.

Mapping miners to wallet addresses. Mining pools generally require registration to mine with the pool. When mining, the worker supplies the user name created at registration; all earnings are credited to that user and periodically transferred to a wallet address specified at registration. (The exception are pools that support so-called pseudonymous mining, in which the worker specifies the payout wallet address—rather than a user name—when connecting; in this case, no mapping is necessary.)

Pools do not normally list miner wallet addresses publicly, making it difficult to connect mining activity to payouts. To obtain this information, we resort to non-technical means, contacting the pool operators directly to ask for information about specific accounts. Some operators kindly provided us with this information, either sharing with us the payout address or the total amount paid out to it. Operators are sensitive about privacy and only provided information about users they themselves had identified as botnet miners.

Publicly-visible pool statistics. One pool, Bitclockers, provides a leader board, showing total user earnings, and work contribution for each solved block. We use this information to determine the earnings of 38 users (Section V-J).

The Eligius and 50 BTC pools provide public statistics about users mining pseudonymously. For malware mining operations using these pools, we obtain earnings and other information directly from these public statistics.

Our source of information about the Fareit botnet is the botnet itself. This botnet operates its own mining pool servers, operating as a dark pool (Section III-C). The mining server software is a fork of the P2Pool mining server code base.[8] This particular mining server provides miner statistics, which we are able to obtain directly from the dark pool servers.

Blockchain analysis. Because all transactions are visible, knowing the addresses to which mining payments are sent allows us to estimate the earnings of a specific miner via examination of the blockchain. We use the payout addresses provided to us by various pool operators. Given these addresses, we first need to isolate mining payouts from other types of transactions. To identify mining pool payouts, we use the technique of Meiklejohn *et al.* [20] to identify the payout transactions of five major mining pools: 50 BTC, BTC Guild, Deepbit, Eligius, and P2Pool. Briefly, this technique relies upon knowledge of patterns or addresses specific to each pool. For instance, a Deepbit payout transaction always uses the same address as the sender, and BTC Guild always sends its initial mining reward to the same address (at which point it pays each miner in an identifiable chain of transactions).

Once we have a collection of transactions representing the mining payouts to the address, we then consider all mining revenue to be derived from botnet mining. This number forms a lower bound on the actual mining revenue, as the techniques of Meiklejohn *et al.* [20] may fail to identify certain payout transactions in order to avoid false positives. While one might argue that it is possible for a botmaster to re-use this same address for legitimate mining operations, we view this possibility as unlikely. First, re-using the same Bitcoin address for multiple purposes has the potentially negative effects that it confuses bookkeeping for the owner and serves to de-anonymize her (as two users now know her by the same pseudonym). Second, re-using the same address has essentially no positive effect, as generating a new Bitcoin address requires generating only a signing keypair, and thus has virtually no computational cost. Finally, and perhaps most importantly, re-using the same payout address jeopardizes legitimate mining revenue, as botnet miners are routinely banned by pool operators.

Clustering wallet addresses. While re-using a single Bitcoin address might be unattractive to a botmaster, there are a number of reasons why one might use multiple addresses. For example, using different pool credentials for each malware distribution campaign would allow her to track earnings for each campaign separately. Using separate addresses also offers some protection against detection by a pool operator, as it spreads the activity across several accounts; even if one address were blocked by a pool operator, only those bots mining to that banned address would be affected.

To identify addresses belonging to the same botmaster, we rely on the observation—due to Satoshi Nakamoto himself [22]—that addresses used as inputs to the same transaction are controlled by the same user. This technique is employed frequently in studies of anonymity within the Bitcoin network [20], [25], [26], and we use it to cluster otherwise distinct malware. This clustering is especially useful for smaller mining operations: e.g., in the case of the BMControl malware, we first identified the family using this technique, and later confirmed the clustering by identifying and decoding its Pastebin-based command-and-control channel.

Clustering also allows us to identify other wallet addresses used by the botmaster. We refer to wallet addresses directly associated with malware mining as *primary* wallet addresses. We refer to wallet addresses in the same cluster as a primary wallet address, but which are not themselves primary addresses, as *secondary* wallet addresses. The income received by secondary wallet addresses may include mining income from other malware mining operations of the same botmaster that are unknown to us. It may also include, however, other sources of income, including some that may be legitimate. For this reason, we report the included income of secondary wallet addresses separately.

D. Estimating Infected Population

We contacted a top anti-virus software vendor (with an install base of millions across the world) with the MD5 hashes of the mining malware, and obtained from them, for each of our 976 samples, an aggregate list of countries from which the

[8] http://github.com/forrestv/p2pool

mining malware was seen to be operating along with the count of unique machine infections detected, over a period of about one and a half months around July 2013. The vendor also provided us with the count of their install base per country. Based on this information and the distribution of computers across the world, we can extrapolate the total population of malware infections per malware family as follows.

Let I_i be the number of infections observed by the vendor in country i, and M_i be the number of machines in that country that subscribe to the vendor's monitoring service. Using an approximate number of computers T_i for country i, the estimated bot population E_i can be computed as

$$E_i = (I_i/M_i) \times T_i.$$

The CIA factbook [4] reports the number of Internet users for 2009; we assume one computer per Internet user for T_i. In practice, this assumption holds for known data points (e.g. Worldwide PCs deployment is projected[9] to be 1.9 billion in 2013, while the CIA factbook estimates the total Internet population as 1.8 billion).

We expect the estimates here to be lower bounds for the following reasons: first, computers that do not have anti-virus protection from the vendor are not counted, including computers with no anti-virus protection at all—such computers are likely to be infected with malware over the long term, contributing to more mining. Second, the estimates are only for the specific binaries we collect. Many of the malware families involved in mining are polymorphic, so we expect many more samples that are not considered here.

E. Identifying Pool Proxies

The techniques described above work for malware engaged in direct public pool mining; that is, where the malware connects directly to a public pool, or for malware where the dark pool provides information, as in the case of the Fareit botnet. In some cases, however, the pool server to which the malware connects is not a known public pool nor does it report any useful information via a statistics Web page. These types of malware mining operations are the hardest to measure. If the server is no longer in operation, our options are limited still further. Here we describe the techniques we use to glean what information we could about such mining operations.

Cross-login test. Since the getwork protocol uses HTTP as a transport, it can be proxied by an HTTP proxy such as nginx without modification. In the simplest case, incoming connections are transparently proxied to a public mining pool. Such a proxy passes through all HTTP headers unchanged, including the Authorization header used by HTTP basic access authentication. In this configuration, bots must use credentials that are valid for the destination pool. To detect this form of proxying, we create miner accounts at several major mining pools and attempt to connect via the suspected proxy using the registered user names, as well as one randomly-generated name we confirmed did not correspond to an exist user name at any of the major pools. If the suspected proxy proxies to one of the major pools, then exactly one user name should succeed in authenticating—the user registered with the public pool to which to proxy is pointed. We identify one transparent proxy: `domain-crawlers.com` transparently proxies all connections to the 50 BTC public pool.

We also test whether pool credentials found in malware could successfully authenticate to a public mining pool. We find this to be the case for a number of worker user names; however this test is not conclusive so we draw no conclusions from this test alone. Rather, we use this information to engage with pool operators; in cases where the pool operator independently confirms that the miner was suspected of mining using a botnet, we include the miner in the analysis.

Passive DNS. The lifetime of a dark mining pool is usually indicative of the lifetime of the corresponding botnet. To determine when such pools were first and last seen, we use the passive DNS data from the ISC Security Information Exchange.[10] The passive DNS dataset contains DNS lookups issued by recursive resolvers at several vantage points, including a major US consumer ISP. We use this dataset for the purpose of discovering the DNS A-records historically returned for domain names of interest between October 2011 and April 2013. In addition to showing the first- and last-seen dates of dark pools, this data set also illustrates the overlap of A-records across different domains. For instance, two dark pools, `dload.asia` and `aquarium-stanakny.org`, pointed to the same IP addresses in the past. This coincidence suggests that the same botnet operation may be behind both domains.

Block reversal. In some cases, we can attribute a successfully mined block to a particular mining pool. The getwork call needs to provide different work for each worker, but if there are no new transactions added between getwork requests, subsequent calls would produce the same value. Pool servers counter this problem by changing the coinbase value for each call to getwork, using it as an additional nonce. Here, we use it as a signature of a particular pool: while the pool clearly will not provide the same coinbase to two different workers, many pools provide similar coinbases across workers.

We repeatedly poll all known pool servers several times a second for a period of three weeks. Then, for every block published during our monitoring, we perform a brute-force search modifying the coinbase of the published block (based on changing only the bits which change when examining the most-similar coinbase in the blockchain at the time) and checking whether the modified coinbase corresponds to one of our recorded getwork requests. A match indicates that the monitored pool is likely to have mined that particular block.

Although this approach is only effective against pools with low coinbase entropy, we are able to attribute blocks to both the Deepbit and 50 BTC pools. It also confirms that `domain-crawlers.com` was proxying to 50 BTC, as we discover multiple blocks where getwork calls to both a 50 BTC pool server and the `domain-crawlers.com` server correspond to the blocks published in the blockchain, suggesting that `domain-crawlers.com` simply forwards the getwork request on to 50 BTC.

Leaked data. In one case, we could glean information about a Bitcoin botnet mining operation from leaked data. Specifically, information about FeodalCash, an affiliate-based program that

[9] According to the Computer Industry Almanac, Worldwide PC use executive summary

[10] https://sie.isc.org/

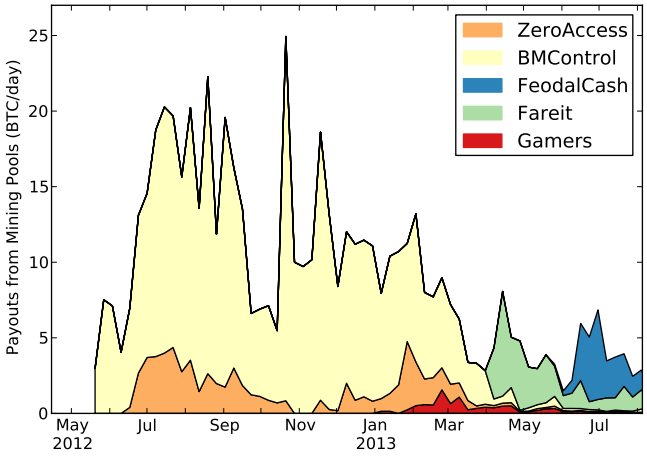

(a) Total daily bitcoin payouts for each botnet.

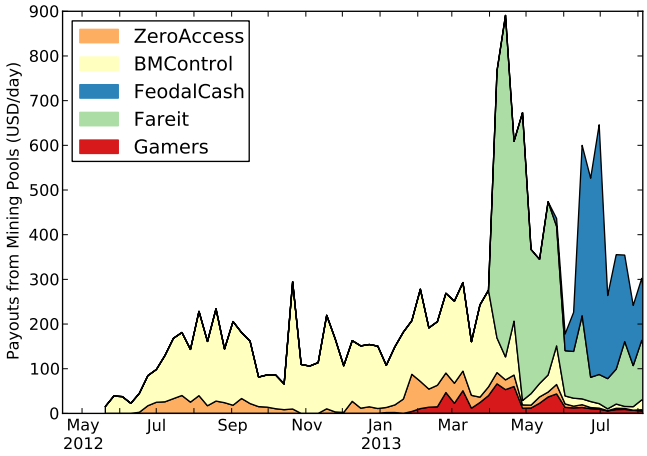

(b) The same payouts expressed in USD at the day's exchange rate.

Fig. 2: Two stacked line graphs showing the amount of mining payouts that botnets received over time, in BTC and in USD. The aggregate mining of all the operations never exceeds 0.4% of the bitcoins generated each day.

TABLE I: Bitcoin mining operations covered by our study.

Family	*Sec.*	*EXEs*	*Wlts*	*Active period*	*BTC*	*USD*
DLoad.asia	V-A	322	—	Dec '11 – Jun '13	10,000?	10,000?
ZeroAccess	V-B	976	3	Dec '11 – Today	486	8,291
BMControl	V-C	54	47	May '12 – May '13	3,097	46,301
FeodalCash	V-D	—	238	May '13 – Today	168	15,941
Fareit	V-E	5	1	Apr '13 – Today	265	30,448
Zenica	V-F	67	—	—	170?	—
HitmanUK	V-G	5	1	Mar '13 – Today	4	362
Xfhp.ru	V-H	42	—	—	—	—
Skype Miner	V-I	17	—	—	250?	—
Misc.	V-J	—	—	Dec '11 – Today	539	17,166

pays botnet operators to install their Bitcoin mining malware, was publicly posted on the Internet. This data enables us to identify earnings from the entire operation as well as earnings from individual affiliates of the program.

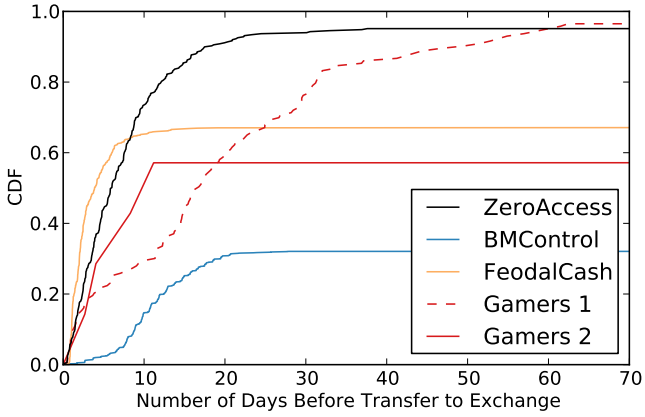

Fig. 3: Delay in transfer to exchanges for different botnets.

V. Analysis

Recall that our goal is to identify major Bitcoin mining operations, their scope, and revenue. In this section we describe nine major mining operations, including a Bitcoin mining affiliate program (Section V-D), as well as 80 smaller mining operations, most represented by a single executable found in the wild.

Table I summarizes our findings. The *EXEs* column shows the number of executables we observe engaged in mining. Several families of malware—ZeroAccess especially—are very aggressive about repacking binaries; it is likely that our sample does not represent the entire set of binaries in the wild. (Recall that our main means of identifying mining malware are reports by ThreatExpert, VirusTotal, and Emerging Threats.)

The *Wlts* column gives the number of wallet addresses known to us that receive payouts. FeodalCash, an affiliate program, has the largest number of wallet addresses (238) because each affiliate mines to a unique wallet, allowing earnings to be credited properly. BMControl also has a large number of wallets (47); we suspect it is also an affiliate program.

The *Active period* column shows the period when the mining operation was active, based on mining data and malware distribution activity.

The *BTC* column shows our estimate of each operation's total earnings. In most cases, earnings are measured directly, either from mining pool statistics (ZeroAccess, Fareit, BMControl, and FeodalCash), or from the blockchain based on payout address (DarkSons and HitmanUK), based on earnings reported by the pool (Skype Miner and Zenica), and finally based on order-of-magnitude estimates by the pool operator (Redem). Small mining operations covered in the miscellaneous section (Section V-J) use all four of the above types of estimates.

The *USD* column provides an estimate of the earnings in US dollars, using the exchange rate at the time of payout. Thus, although earlier mining operations (e.g., DLoad.asia) earned over 10 million dollars' worth of bitcoins at the exchange rates in effect on November 30, 2013, at the time of mining a bitcoin was worth considerably less. In two cases—Skype Miner and Zenica—we do not have accurate information about when the bitcoins were earned, so cannot estimate the

equivalent US dollar value accurately. We plot our estimate of the daily earnings of the five largest operations (in terms of revenue) in Figure 2, and their cumulative earnings in Figure 4. The latter further breaks down the earnings into just those transferred from the primary wallet addresses (i.e., the address(es) associated with the mining credentials used by the malware) as well as transfers from associated wallet addresses (see Section IV-C).

Our estimates notwithstanding, the true "takehome" earnings in terms of USD (or any other fiat currency) depend entirely on how—and when—the bitcoins are "cashed out", typically by transferring them to an exchange. Hence, transfers to exchanges are of particular interest, as they serve—with very few exceptions—as a necessary precursor to cashing out of the Bitcoin economy. Unfortunately, because most exchanges double as online banks we cannot claim definitively when—or if—all these earnings were converted to fiat currency.

Moreover, in some cases, the mining profits might travel through several intermediate addresses before arriving at an exchange. For simplicity, we consider only cases with no intermediate addresses; i.e., cases where the bitcoins earned from mining are spent immediately at an exchange. We define the transfer time as the interval between the mining payout and the actual transfer. We use the techniques of Meiklejohn *et al.* [20] to identify wallet addresses associated with exchanges. Figure 3 shows the delay between when a botnet receives payment for mining and when it transfers its earnings to an exchange. In most cases, botmasters liquidated their bitcoins shortly after mining.

In the remainder of this section, we describe each of the mining operations listed in Table I in greater detail.

A. DLoad.asia (Redem and DarkSons)

The DLoad.asia operation is one of the earliest major mining operations we encountered. More properly, the DLoad.asia operation consists of several mining operations using shared infrastructure. At least two individuals are behind the operation, known by the handles Redem (a.k.a. Mpower) and DarkSons (a.k.a. MrDD).

Operation. Based on information provided by one public mining pool operator, these individuals began mining in 2011, initially connecting to the pool directly and later via a proxy. These botnets continued mining using the same pool user names (variations of "Redem" and "DarkSons") even when connecting through a proxy. Later generations of malware used different miner user names and proxy domain names. Despite this, the server IP addresses and domain names were not changed in unison, making it possible to track the infrastructure as it evolved. Most recently, the DLoad.asia infrastructure was used as a mining proxy and NGRBot command-and-control channel. As documented by the "Inside Your Botnet" blog, the Redem and DarkSons names continued to appear in IRC channel and user names [9], [10], [11], [12], [13], [14], as well as in domain registration records.

Earnings. The pool operator shared with us the wallet address that DarkSons used. The wallet address was last active in November 2012, at which point it had amassed 2,403 BTC. The techniques of Meiklejohn *et al.* [20] are unable however, to identify any direct payments from mining pools. The

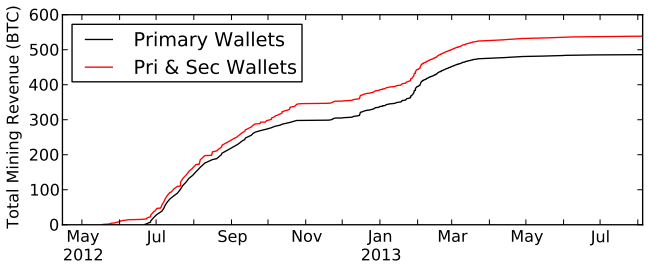

(a) ZeroAccess

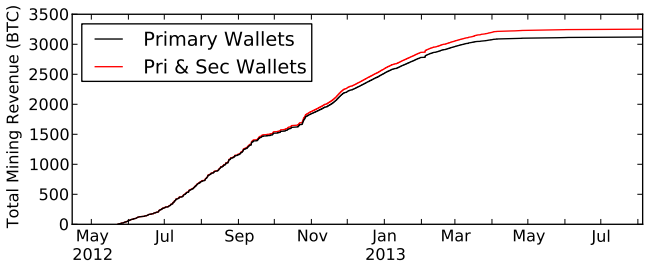

(b) BMControl

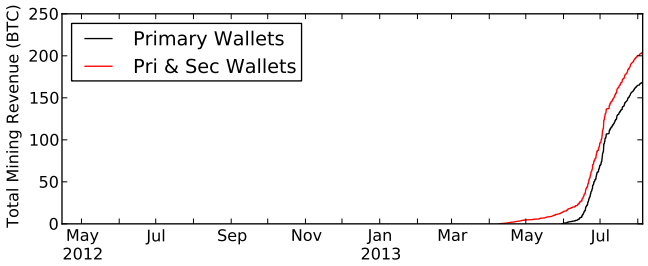

(c) FeodalCash

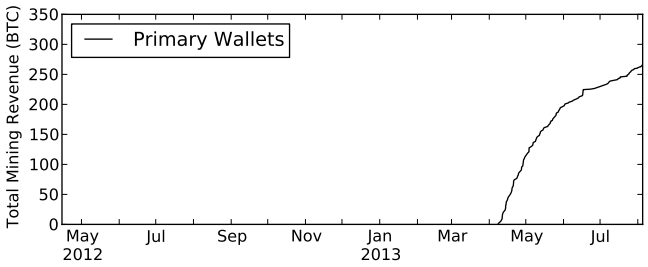

(d) Fareit

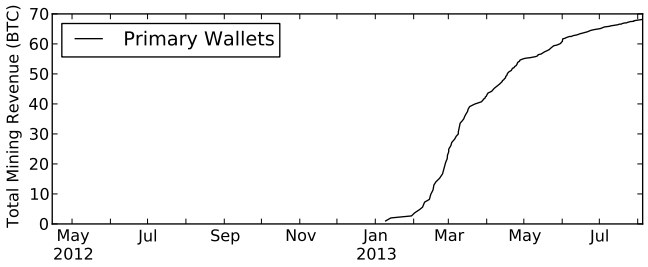

(e) Gamers

Fig. 4: Mining revenue by different botnet operations.

TABLE II: Distribution of DLoad.asia infections by country.

Country	*Share*	*Est.*
Brazil	16.0%	10,600
Malaysia	9.5%	19,300
Indonesia	8.7%	11,700
Russia	5.9%	4,200
South Korea	5.8%	7,100
Others	54.1%	71,800

blockchain does reveal a number of transactions in which the DarkSons wallet received block rewards through an intermediate wallet. The botnet received a total of 1,681 BTC through these transactions.

We are unable to locate a payout address for Redem. However, the pool operator recalls that the botnet connected to the pool using over 100,000 unique IP addresses and had a peak mining rate of over 100 GH/s. The operator estimates that the bot earned at least 10,000 BTC. During that time period, however, a bitcoin was worth only about $1. Our estimate of Redem's portion of DLoad.asia earnings in US dollars shown in Table I is therefore based on a 1:1 exchange rate.

Population. Although the DLoad.asia infrastructure is no longer active, infected hosts can still be found in the wild. Table II presents our estimation of the geographic distribution of infections based upon the data provided to us by a major anti-virus software vendor. The *Share* column shows the ratio of the number of infected hosts to the total number of hosts with the vendor's product in a given country. Based on this percentage and the number of computers in the country, we estimate the infection population. Brazil accounts for the largest share of infections; the vendor's coverage is smaller in Brazil than in Malaysia, however, so the estimated number of infected hosts is larger in Malaysia.

B. ZeroAccess

The ZeroAccess botnet is currently one of the largest botnets, with estimated 9 million infected PCs of which one million are online at a given time [33]. ZeroAccess uses drive-by downloads and other methods to infect victims [8]. The core of the botnet is a rootkit and peer-to-peer command-and-control (C&C) protocol. Using the C&C protocol, bots can fetch modules that enable the bot to carry out tasks such as mining bitcoins or committing click fraud. Bots can be configured to perform only one task. It is possible to update the modules as necessary through the same C&C protocol.

Operation. ZeroAccess began mining through a proxy server, `tang0-hote1.com`, and changed proxy servers several times to domains such as `google-updaete.com` and `great-0portunity.com`. According to our passive DNS data, these domains were first active in December 2011. They are currently not active; as of mid-June 2013 ZeroAccess is mining directly through Eligius, a public mining pool that offers detailed hash-rate graphs for every user. Using credentials embedded in the ZeroAccess malware, we find one wallet address, `1ASNjJ`, that it uses to mine at Eligius. Figure 5 presents the daily mining rates for ZeroAccess and two other operations that also use Eligius (discussed in subsequent sections).

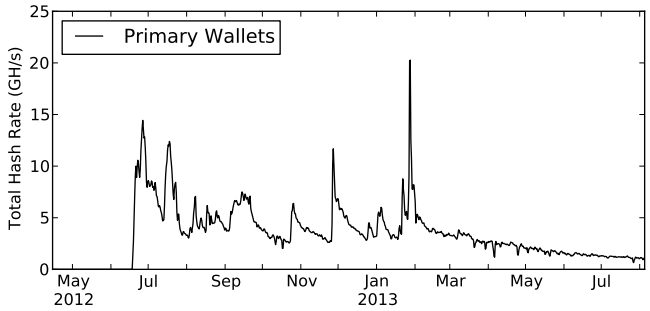

(a) ZeroAccess

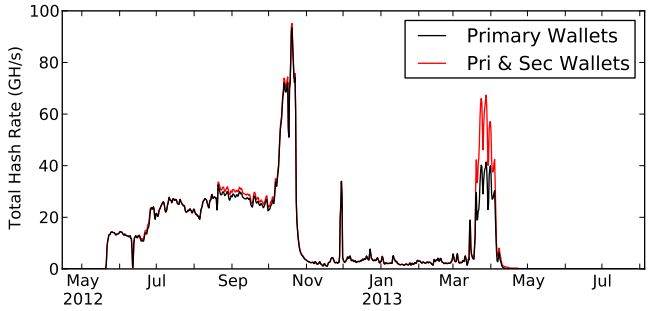

(b) BMControl

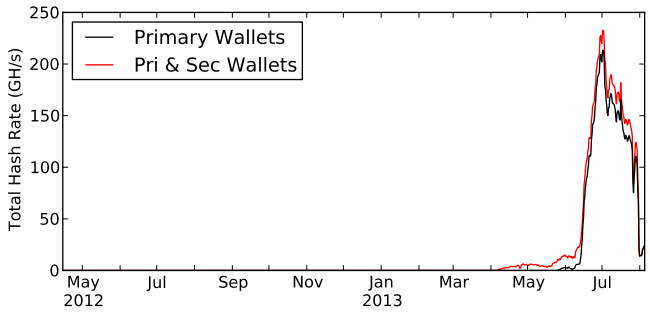

(c) FeodalCash

Fig. 5: Mining rates of three botnet operations at Eligius, a mining pool that publishes each user's hash rates over time.

Earnings. Our analysis of the earnings of ZeroAccess is limited to the most recent version that mines directly through Eligius. So far, the botnet has received more than 400 BTC from mining payouts (Figure 4a). The botnet is currently mining at less than 1 GH/s, although the peak in February 2013 was close to 20 GH/s (Figure 5a).

Population. Based on information provided by the security vendor, most of the Bitcoin-mining bots were located in Europe, with over 25 countries in Europe accounting for about 50% of all observed infections. On the other hand, the malware itself is widespread, with infections detected in more than 60 countries. Table III shows distribution of the observed bot population for 976 binaries for the top five infected countries, as a percentage of total infections observed in the *Share* column, while the *Est.* column gives the number of infections extrapolated as described in Section IV-D.

TABLE III: Distribution of ZeroAccess infections by country.

Country	*Share*	*Est.*
United States	14.9 %	2,600
France	12.2 %	1,800
Russia	8.2 %	800
Czech Republic	5.1 %	900
Canada	4.8 %	817
Others	54.8 %	10,600

The population estimate is much lower than previously published estimates [27], [33], which suggest the ZeroAccess bot population is between 1.2 to 9 million. However, our estimates are only for the 976 binaries we obtained and know to engage in Bitcoin mining for specific wallets. ZeroAccess is known to be polymorphic, so a large number of binaries are expected. Anti-virus vendors we checked with had well over a hundred thousand binaries labeled as ZeroAccess. Hence, we do not claim that our estimate represents the overall ZeroAccess bot population or its potential mining profits, only the subset we observe in action.

Another way to localize the botnet is to use the diurnal pattern of its operation [5]. To analyze the periodicity, we find the hours of each day (in UTC) at which the botnet's hashing rate reaches a local minimum. Then we compute the probability distribution of these relatively dormant hours, a histogram of which is shown in the leftmost portion of Figure 6. As shown in the graph for ZeroAccess, the botnet is the slowest around midnight UTC, suggesting that the majority of infected hosts that mine with the `1ASNjJ` wallet address are located in Asia [5]. However, Table III suggests that the US has the largest bot population. It is likely that the botnet uses multiple wallet addresses in its binaries, of which we are able to find only one.

Transfers to exchanges. The ZeroAccess line in Figure 3 shows the distribution of mining revenue that was transferred (within a single hop) to an exchange. The botnet transferred to an exchange more than 90% of the mining revenue that its primary wallet addresses received, using BTC-e as the primary exchange. The median time to do so is about a week. The botnet moved the remainder of its revenue to wallet addresses that we cannot identify as exchanges. These earnings might have been reinvested within the Bitcoin economy, or they might have been transferred to an exchange through intermediate wallet addresses.

C. BMControl

Another botnet that mines at Eligius is one we call BMControl, which can be identified by its command-and-control channel that uses specific PasteBin URLs to distribute configuration data to bots. We name this family of malware based upon the PasteBin user that uploaded the configuration data, BMControl. Upon startup, the malware retrieves and decodes the data contained in the PasteBin URL, and executes the mining binary. The configuration is a Base64-encoded string that includes the parameters to run the mining executable and credentials for logging into the pool servers. The BMControl botnet was documented online in September 2012 [34].

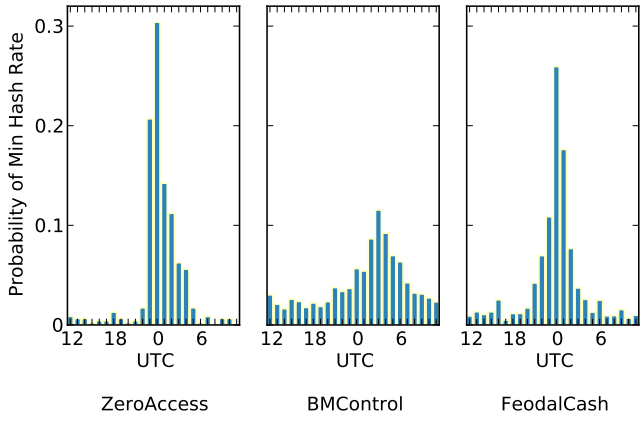

Fig. 6: The distribution of times (in UTC) at which Eligius-based botnets achieve minimal mining rate.

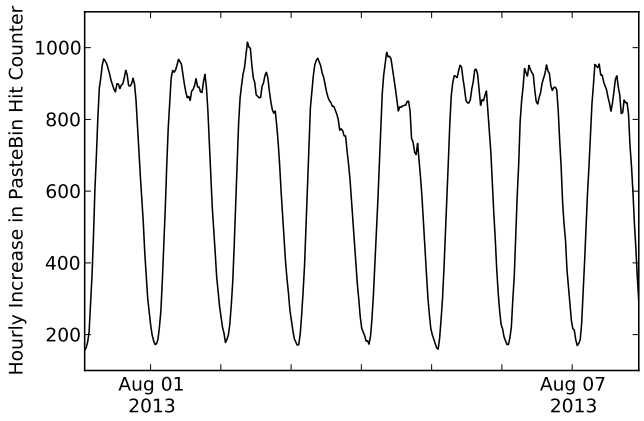

Fig. 7: PasteBin counters for BMControl configuration URLs.

Operation. BMControl has mined through proxies as well as directly through several pools. The configuration file at PasteBin contains a list of worker credentials for pools. When we first began monitoring this botnet, the configuration included only Bitcoin mining pools and used wallet addresses as worker names (a common feature of several Bitcoin pools). The primary pools that were used for Bitcoin mining were Eligius, 50 BTC and EclipseMC. Subsequent versions of the configuration file do not use Bitcoin wallet addresses as worker names, instead preferring to list usernames and passwords for the pools. The most recent configuration files for BMControl have included Litecoin pools and new worker credentials.

Earnings. For each distinct PasteBin URL, there is a counter of unique visitors that attempts to identify new visits based on cookies and IP addresses. For the two primary BMControl PasteBin URLs, there are over 8 million unique visits. This number increases between 200 and 1,000 every hour. The rate that the counter increases for one week at the beginning of August, 2013 is shown in Figure 7. Since this PasteBin post is only useful if one can decode the contents and it requires knowledge of the URL to find it, we can reasonably estimate that increases in the counter are due to new infections and daily check-ins by the malware. Using this, we estimate that there are around 16,000 bots online each day.

TABLE IV: Distribution of BMControl infections by country.

Country	*Share*	*Est.*
Bulgaria	50.6%	99,900
Turkey	28.8%	40,000
Macedonia	2.7%	7,000
Brazil	1.4%	2,200
Slovenia	1.4%	3,300
Others	15.0%	52,000

As seen in Figure 4b and Figure 5b we no longer see the BMControl botnet mining bitcoins. Instead, the botnet has changed to mining litecoins through `litecoinpool.org`. We have confirmed this by decoding the configuration file as well as through contacting the `litecoinpool.org` administrators who have acknowledged that the workers used by BMControl are earning litecoins. We discuss Litecoin mining further in the Epilogue.

Using the last known mining rates (Figure 5b) and the estimate of 16,000 bots active per day from the PasteBin counter increases, we estimate that the average mining rate per bot is 3.75MH/sec.

Population. Eastern European countries account for more than 80% of the BMControl infections, with Bulgaria dominating the list shown in Table IV. This matches well with the diurnal cycle of the mining rate shown in Figure 6. The minimum mining rate happens around 3:00 UTC and Bulgaria is on Eastern European Time (UTC +2 or +3).

Transfers to exchanges. Using the same methodology as for ZeroAccess, we examine how BMControl transferred its mining revenue—which it received in its primary wallet addresses—to exchanges. According to Figure 3, the botnet transferred around 30% of the revenue, with a median transfer time of around two weeks. Most of the transfers took place at the Bitcoin-24 exchange.

D. FeodalCash

Details about FeodalCash, the last of the major botnet operations we see mining at Eligius, were leaked and publicly posted onto the Internet [17]. From this leaked data, we can see that FeodalCash is an affiliate program that provides (GPU-capable) Bitcoin mining malware that affiliates install on their bots. In turn, FeodalCash then pays affiliates a fraction of the revenue earned by their bots. This type of labor division enables the affiliates to focus on gaining more bots while the affiliate program can focus on maintaining the malware and infrastructure.

Operation. According to the leaked data, the botnet started operating in May 2013 and there were 238 active affiliates at the time of the data leak. The Bitcoin mining malware was configured to directly mine with the Eligius Bitcoin mining pool and each affiliate was assigned an individual wallet.

Earnings. Since this botnet used Eligius, we can gather a complete profile of their earnings and hash rate over time, as shown in Figures 4c and 5c. The botnet did not start earning much until more affiliates joined the program around the end of June 2013. At this point the botnet reached a peak of almost 250 GH/s and has since experienced a steady decline in earnings as the difficulty level has increased while its hashing rate has fallen. At the time of writing the botnet has currently earned 168 BTC, which translates to approximately 15,941 USD.

Population. If an average PC can mine at about 4 MH/s, we estimate that the bot consisted of 62,500 hosts at its peak hashing rate. In addition, we analyze the diurnal patterns in the hash rate graph. We focus on the hours at which the hashing rate is the lowest every day. As shown in the rightmost portion of Figure 6, the botnet reaches minimal activity around midnight UTC. This suggests that the majority of the infected hosts are in Asia.

Transfers to exchanges. As shown in Figure 3, the botnet transferred more than 60% of the mining revenue to exchanges. The botnet almost exclusively used WebMoney as the exchange service. The median transfer time is less than five days.

E. Fareit Bots

The Fareit botnet originally focused on stealing passwords and DDoS attacks. However, on April 9th, 2013 it began distributing Bitcoin mining malware [31].

Distribution. This botnet uses the popular Black Hole exploit kit to install a small executable that contacts `kgtxdu.info` to download an open source Bitcoin mining client called CGMiner[11] onto the victim's system. CGMiner is disguised as a Flash.exe and once downloaded, a Visual Basic script is used to invoke the miner program with a predetermined command line string. The Visual Basic script is then copied onto the Startup directory of a windows system so that the miner will be persistent even when the victim reboots their computer.

Operation. The Bitcoin mining malware contacts a proxy server, `coonefix.ru`, which proxies connections to the public pool `p2pool.org`.[12] It is an example of a smart proxy, as shown in Figure 1. The proxy server reports fine-grained data, such as mean payout values, current hashing and stale share rates, which we plot in Figure 8. All of this information provides us with deeper insights into the inner workings of their botnet mining operation.

Earnings. To identify the botmaster's wallet address, we look for a wallet that receives payouts from P2Pool, such that the payout rate is consistent with the pool's hash rate, and that the first payout occurred on the same day (April 9th 2013) when the malware started mining. As of November 7th 2013, the Fareit botnet's wallet has received at least 265 BTC of mining revenue. As the global Bitcoin difficulty increases, Fareit has been receiving mining payouts at a slower rate (Figure 4d).

Population. We leverage the *stale share* rate to estimate the botnet's population. A share is the proof-of-work that miners submit to the mining pool. The share becomes stale when the another mining pool has mined the block. Whatever work the mining pool, along with its miners, has put in so far is essentially wasted. If a total hash rate of a pool is high, it is

[11] https://github.com/ckolivas/cgminer

[12] The pool server code on `coonefix.ru` is a fork of the original P2Pool open source software available at https://github.com/forrestv/p2pool

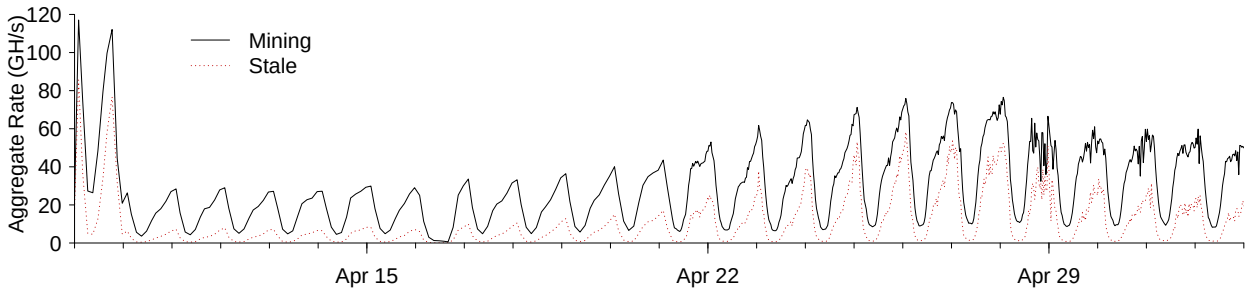

Fig. 8: Fareit hash rate and stale share rate as reported by the proxy pool server `coonefix.ru`.

TABLE V: Distribution of Zenica infections by country.

Country	*Share*	*Est.*
Vietnam	61.0%	119,800
Thailand	9.2%	16,000
Romania	4.3%	5,200
Taiwan	3.2%	5,100
United states	2.3%	3,300
Others	20.0%	32,200

TABLE VI: Distribution of infections by country for Xfhp.ru.

Country	*Share*	*Est.*
Indonesia	10.9%	3,200
Mexico	7.3%	1,200
Peru	6.1%	1,900
Thailand	5.5%	1,800
Brazil	4.8%	700
Others	65.5%	28,000

less likely that another pool will have mined the block first. The stale share rate will thus be lower.

To estimate the number of infected hosts based on the stale share rate, we perform the following experiment. We mine for P2Pool with a standard desktop computer, which was capable of mining at 4.6 MH/s. We observe that our stale share rate is 24%. Meanwhile, the Fareit proxy reports a stale share rate of 34%. Since the stale share rate goes up as hashing rate goes down[13] , at this point in time the average hashing rate of a bot is less than 4.6 MH/s. If we assume a compromised host mines at 4 MH/s, a low standard deviation in the hashing rate of bot and a total hashing rate of 50 GH/s (a long-term average of the mining rate shown in Figure 8), we can estimate there are about 12,500 bots mining in this botnet.

F. Zenica

Zenica is a botnet that mines at a major public pool. It appears to be operated by one person. Unlike the other major botnets, there are few activity reports of this botnet on anti-virus websites, security blogs or online forums. We are not sure how the malware is distributed or how the botnet operates. However, its sheer size and large earnings merit close scrutiny.

Earnings. We find 67 malware binaries that connected to the mining pool via the username `zenica@gmail.com`. We contacted the pool operator about this user. The operator claimed that the account "had 312,000+ active IPs" and was "paid out about 170 BTC in 3 months."

Population. Zenica bots are most prevalent in Southeast Asia (Table V), with Vietnam and Thailand accounting for over 70% of the sampled infections.

[13]We confirm that a higher hashing rate results in lower stale-share rates by mining with a CPU capable of 18 MH/s and observing a stale-share rate of 14%.

G. HitmanUK

HitmanUK is a botnet that mines at a major public pool. It has a relatively small mining income: 4 BTC to date. Even so, it makes an interesting case study, because the botmaster launched a DDoS attack on the pool when the pool first blacklisted the botnet.

Operation. We find five malware binaries with the username "hitmanuk." According to the pool operator, the account is associated with the wallet address `1ARHrS`. The binaries and the wallet address were first seen in February 2013. It appears that the botnet has remained active since; the wallet is still receiving mining payouts.

At some point, the pool operator blacklisted the botnet's account, possibly due to reports of malware. The botnet immediately retaliated by launching a DDoS attack on the pool's mining server, paralyzing the entire pool and preventing other users from mining for a few hours. In the end, the pool operator gave in and unfroze HitmanUK's account. This incident suggests that the botnet was—at least at the time—of considerable size.

Earnings. HitmanUK's wallet is active to this day. At the time of this writing, it has received 4 BTC, worth $362 at the time of payout.

H. Xfhp.ru Miner

This botnet uses ZBot, also known as Zeus, which connects to `xfhp.ru`. At the time of writing the domain is still active and runs a stratum proxy pool server. ZBot then downloads a plugin that does Bitcoin mining.

Population. Most of the infections for this malware come from Southeast Asia and South American countries, perhaps indicating that the botmaster chose to buy cheaper hosts. Table VI shows the distribution by country and extrapolated population.

Although the estimated infected population of the instance is rather modest, this is another example of a major malware family incorporating Bitcoin mining in addition to other activities.

I. Skype Miner

We name this botnet "Skype Miner" because at one point it used a combination of Skype and social engineering to distribute the malware. To carry out the attack, the bot sent a Skype Instant message from a compromised Skype account by the name of "Carolina Chapparo" [1]. If the victim clicked on the link in the message, she would be taken to a webpage that contained a drive-by-download exploit pack. The executable would attempt to install the Bitcoin mining malware.

Operation. The initial samples of this malware that was distributed beginning in July 2012 used the same credentials as the version that was distributed via Skype during April 2013. The original malware sample uses `keep.husting4life.biz` as its pool domain and the newer version uses `suppp.cantvenlinea.biz`. Information included in the Stratum headers indicates that both of these domains are proxying connections to the same public pool. In private conversations the pool operators confirmed that this botnet was proxying to their pool.

Earnings. According to the mining pool operators, the user received about 250 BTC. However, they did not provide a wallet for us to confirm these earnings.

J. Miscellaneous

In addition to the mining operations above, we also find numerous smaller mining operations, many of which mine directly using a fixed set of credentials embedded into the malware binary.

Mining at registration-based public pools. Bitclockers is the only registration-based mining pool that publishes each user's earnings. From malware reports, we extract all usernames that were associated with Bitclockers. We look up all 38 of them in Bitclocker's public records and examine their earnings. After summing up the individual payouts, we find that they have earned close to 30 BTC in total. The biggest earner accumulated 9.6 BTC between November 2012 to January 2013.

In contrast, most major registration-based mining pools do not publish user statistics. We have to manually contact the pool operators, via email or IRC, for user information. One pool operator reports to us a botnet that specifically targeted gamers (we therefore refer to it as Gamers). He provided us with four usernames and their wallet addresses. According to a forum post—purportedly written by an infected user—the malware disguised itself as a game executable, which connected to the mining pool via one of the four wallet addresses.[14]

We analyze the mining payouts for the primary and secondary wallets for the Gamers botnet and present their earnings in Figure 4e. It shows that the botnet first became active in January 2013. Mining activities have waned since mid-June, possibly after a crackdown by the pool's operator or anti-virus companies.

TABLE VII: Miscellanous mining operations.

Worker	*BTC*	*USD*
`ophelion` (Gamers 1)	67.45	4,552.64
`1HUVG8`	65.03	532.35
`1ES11K`	45.59	600.93
`13CnZa`	37.99	494.35
`19zKyp`	37.80	629.74
`18G7T7`	35.23	4,016.90
`1H1xa5`	29.14	357.54
`1PbPiV`	24.74	208.31
`1AfBS5`	24.37	323.08
`1FiPR4`	23.96	163.29
`17F8N9`	19.92	468.02
`1ByFLx`	17.70	208.87
`1AFVcM`	14.54	135.53
`12W29H`	11.51	839.97
`sarajevo`	9.56	119.02
`15p86j`	7.80	923.78
`boywonder`	7.67	103.71
`1a3dpd`	7.03	79.85
`1PwfoA`	6.82	828.91
`process1`	5.39	72.86
`17pdMw`	5.37	326.07
`15LuUP`	4.85	58.09
`archy10`	4.10	48.57
`1PyoNm`	2.08	250.61
`1Kjvxd`	2.03	25.17
`ridetohell` (Gamers 2)	0.55	50.57
Others	21.58	798.14
Total	539.24	17,166.30

We are able to trace how two of the botnet's wallets transferred the mining revenue to exchanges, as shown in Figure 3. The first wallet (Gamers 1) took a median of three weeks before transferring more than 90% of the mining revenue. The second wallet (Gamers 2), by contrast, transferred a little more than 55%. Both of the transfers happened at the Bitstamp exchange.

The first wallet was also associated with Eligius, another public mining pool. Its hash rate graph displays a typical diurnal pattern that is strongly suggestive of botnet activity. Moreover, the average hash rate is around 4 GH/s in the last two months, with a total of 70 BTC paid out by Eligius. Assuming that an infected host can range from an average CPU-only computer (4 MH/s) to a typical gamer's PC (50 MH/s), we can estimate the size of the botnet as somewhere between 80 to 1,000 infected computers.

In addition to the Gamers botnet, the pool operator also gave us four more malware wallet addresses. We do not know their mode of operation. The four wallet addresses alone have only earned 7.7 BTC from mining since December 2011. However, they are associated with more than 40,000 secondary wallet addresses. We believe that not all of them are involved in receiving mining payouts. One common practice is to have a small number of wallets for mining, while the rest are used for "mixers"—services that attempt to obfuscate the trail of transactions before cashing out, making it difficult, but not impossible, to trace the transactions. Using the techniques

[14] https://bitcointalk.org/index.php?topic=159307.0

in [20], we identify the transactions for mining payouts. We find that both the primary and secondary wallets have received 886 BTC of mining revenue.

Bots for no-registration public pools. We find four additional wallet addresses that mining malware uses to connect to 50 BTC, a public pool that supports both conventional registration-based and no-registration mining. Since receiving their first mining payouts in December 2012, these four addresses have only received 2.6 BTC from mining. If we are to examine all the secondary wallets—all 24 of them—the total revenue from mining amounts to 242 BTC.

At Eligius, we find 29 wallets that do not belong to any of the major botnet operations we study. These wallets alone have yielded an income of 332 BTC from mining since their initial mining payouts in March 2012. They are associated with more than 600,000 secondary wallet addresses. Again, we believe only a small fraction is directly involved in mining. Even so, the total mining revenue for these secondary addresses amounts to more than 30,000 BTC. Some major botnet operations may be behind this, but we leave it to other researchers to analyze.

Bots for proxies to light pools. Recall that in Section IV we identify `domain-crawlers.com`, a dark pool, as a proxy to 50 BTC. We find a total of three usernames associated with mining malware at `domain-crawlers.com`. The operators of 50 BTC confirm them as pool users, but tell us only that the accounts have a total balance of 0.1 BTC. This small amount suggests that the botnet may have already cashed out their mining earnings, but the exact revenue remains a mystery.

VI. Discussion

Bitcoin mining, as evidenced by the operations we examine, can generate non-trivial revenue for a botnet operator (see Tables I and VII). Still, these numbers are nothing like the spectacular earnings—millions of US dollars—estimated for spamming and click fraud [15]. Bitcoin mining as a botnet monetization activity is ultimately judged by its profitability, that is, the expected revenue from Bitcoin mining minus costs.

Mining revenue. Mining revenue—whether from a botnet or a legitimate mining operation—depends on two factors: hashing power and network difficulty. For revenue measured in US dollars, the BTC-to-USD exchange rate is a factor as well. Daily revenue is thus given by:

$$\frac{\text{USD}}{\text{day}} = \frac{\text{sec}}{\text{day}} \cdot \frac{\text{MH}}{\text{sec}} \cdot \frac{\text{BTC}}{\text{MH}} \cdot \frac{\text{USD}}{\text{BTC}}.$$

Here BTC/MH is the expected revenue, in bitcoins, per million SHA-256 computations. At the current difficulty level (November 30, 2013), this is $8.22 \times 10^{-12}\text{MH}/\text{sec}$.

Denote $D = \text{BTC}/\text{MH}$ for short. Denote the exchange rate $U = \text{USD}/\text{BTC}$, which was slightly over \$1,100 per bitcoin on November 30, 2013. Let R be aggregate hash rate in million hashes per second; $R = \text{MH}/\text{s}$. A low-end PC without a GPU is capable of about 4 MH/s, a newer PC without a GPU of about 20 MH/s, and a top of the line AMD Radeon 7970 GPU is capable of about 500 MH/s.

Fig. 9: Daily miner revenue per MH/s of mining capability. Daily revenue per MH/sec of hashing power is given by $86400 \cdot D \cdot U$, where D is the expected revenue in BTC per million hashes computed, and U is the USD:BTC exchange rate.

A botmaster's revenue per bot per day is thus given by

$$\text{USD Daily Revenue} = 86400 \cdot R \cdot D \cdot U. \tag{1}$$

At today's exchange rate and difficulty, this comes to $\$0.00078 \cdot R$. With $D = 8.22 \times 10^{-12}$ and $U = \$1,100$ as above, a low-end PC generates about 0.3¢ per day; a PC with a discrete GPU capable of 100 MH/s can generate about 7.8¢ per day. A network of 10,000 low-end PCs would generate about \$31 per day; if at least one in ten have a discrete GPU capable of 100 MH/s, it would generate another \$75 per day.

Figure 9 shows the daily revenue in USD per MH/s of hashing power as a function of time. That is, the graph plots Eq. 1 with parameter $R = 1$ and parameters D and U varying with time. Revenue per unit of hashing power is at an all-time low—nearly an order of magnitude lower than the previous lows in October 2011 and December 2012 (when the block mining reward halved to 25 BTC).

Botnet costs. We can divide the costs into the cost of acquiring the bots and the cost associated with the monetization scheme itself. Compromised PCs in Asia cost \$5 to \$10 per thousand, as reported by Caballero *et al.* [2], which agrees with our own informal survey of such services. (We note that the wholesale price may be much lower.) The cost of a bot is amortized over its lifetime. Unfortunately, we are not aware of reliable estimates of how long a bot remains infected. (The spikes of activity in Figure 2a suggest the median lifetime is on the order of a week.)

Much less still is known about the non-acquisition costs. These include the cost of infrastructure, development, and day-to-day operations. Neglecting non-acquisition costs, if we estimate that a low-end bot costs 0.2–1¢ to acquire and and can generate 0.2–1¢ per day from mining, then the time to break even ranges anywhere between 1 day and 25 days (operating continuously).

Profitability. The volatility of the BTC-to-USD exchange rate, increasing Bitcoin network difficulty, variance in PC hashing power, and unknown botnet acquisition and operating costs make it difficult to accurately estimate the profitability of

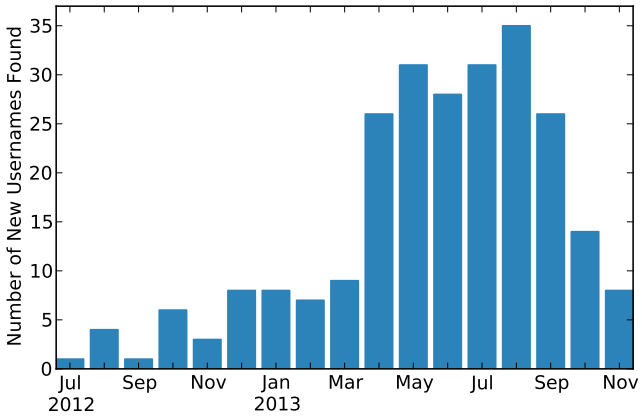

Fig. 10: Monthly new usernames for Litecoin-mining malware.

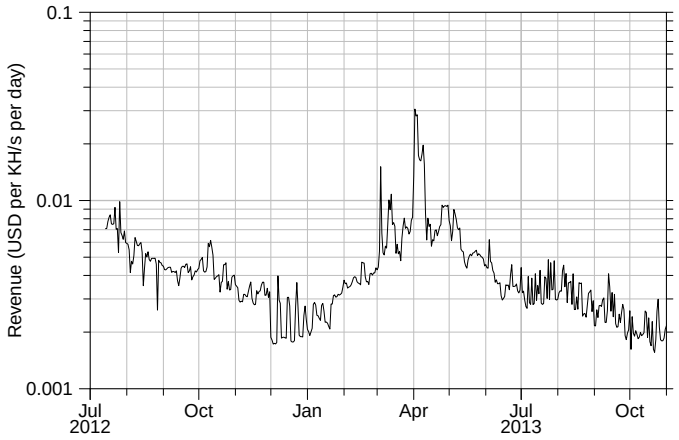

Fig. 11: Daily revenue per KH/s of Litecoin mining capability.

Bitcoin mining. Setting the question of determining profitability aside, let us examine the possible outcomes qualitatively. At any given time, the profitability Bitcoin mining, or more generally, any botnet monetization activity, falls into one of three classes:

Absolutely profitable. The revenue from Bitcoin mining exceeds the costs of operating a botnet solely for mining.

Marginally profitable. For an existing botnet, the additional mining revenue exceeds the additional costs associated with Bitcoin mining.

Unprofitable. Mining revenue does not cover the additional costs of Bitcoin mining.

Throughout most of 2012 and the first quarter of 2013, Bitcoin mining could generate over 1¢ per day from a low-end PC, so that even at retail pricing for bots, Bitcoin mining was an absolutely profitable botnet monetization activity. At least some of the operations dating back to that period, such as the DLoad.asia family, appear to be known only for their mining (of course, we cannot exclude other activities with absolute certainty).

Where we are today is less clear. We observe, however, that the marginal cost of Bitcoin mining, that is, the cost of Bitcoin mining on a botnet already engaged in another activity, is very low. Bitcoin mining does not interfere with other activities such as spamming or click fraud, since it exploits a previously untapped resource: computation. With acquisition and infrastructure costs paid for by another activity, the remaining costs are software development and additional management overhead associated with illicit Bitcoin mining. Because these costs can benefit from economies of scale, we expect Bitcoin mining to remain at least marginally profitable for large botnets.

VII. Conclusion

This paper provides an in-depth analysis of Bitcoin mining malware, which is one of the first methods to directly monetize the computational ability of a compromised computer. Among the results of this work, we show that it is often possible to track the earnings of these botnets due to the fact that all Bitcoin transactions are public. We also show that some of the larger botnets in our analysis have earned sizable amounts of bitcoins and have been in operation for years. Most of the infections of Bitcoin mining malware are traced to geographic regions of lower cost bots, thus increasing the profitability of bots that might not otherwise be valuable. Finally, we have developed a number of methods to trace the mining pool the malware is using even when proxy servers are used to hide the actual mining pool. Our analysis reveals that even larger botnets, such as Dload.asia, ultimately use public mining pools either directly or via proxies to coordinate their mining operations. However, it is challenging for public pool operators to disable their accounts due to the risk of retaliation. While Bitcoin mining might become unprofitable even for lower cost bots due to specialized hardware, other cryptographic currencies such as Litecoin might continue to be profitable.

VIII. Epilogue

We would be remiss if we concluded our inquiry without mentioning Litecoin. Litecoin is another decentralized virtual currency, based on Bitcoin code, that has garnered some interest in the Bitcoin community. Its slogan, "Litecoin is silver to Bitcoin's gold," suggests it is a lower-value complementary currency, and indeed it is currently valued at about 30 litecoins to 1 bitcoin[15]. The only significant changes for Litecoin are a difficulty parameter that produces blocks four times faster and replacing the SHA-256 proof of work with scrypt ($N = 1024$, $P = 1$, and $R = 1$) [24]. The scrypt hash function is not only slower than SHA-256 (a good rule of thumb is 1/1000 the speed for a CPU), but the selected parameters normally require random access to approximately 64 kB of memory.

Litecoin developers selected scrypt to lessen the advantage of specialized hardware (and therefore the ability of someone investing in specialized hardware to control a significant fraction of the mining network). Bitcoin mining ASICs on the market today are 3–4 orders of magnitude more efficient than CPU miners, while Litecoin's design should favor CPU and GPU miners. A typical CPU can mine litecoins at a rate between a few KH/s to tens of KH/s.

From a botmaster's point of view, Litecoin mining is the same as Bitcoin mining, differing only in the executable, mining pools, and profitability. The BMControl bot has already

[15]We obtain Litecoin:USD exchange rate data from the BTC-e exchange. The data set starts on July 13, 2012.

begun mining litecoins, and contacting the pool server operator indicates that the botnet has received 453 LTC (about $900) since April 2013. Another botnet reported online received 6700 LTC (present value about $13,000), mostly between December 2012 and March 2013.

The interest can be observed in other ways. Figure 10 shows the number of new usernames for Litecoin-mining malware discovered in the Emerging Threats database every month, where the first Litecoin-mining malware was found in July 2012. Beginning in April 2013, its clear that the significant uptick in Litecoin-mining malware suggests increased botnet interest in Litecoin. Figure 11 is the recent Litecoin analog to Figure 9, showing the revenue per day per KH/s for a Litecoin miner. In contrast to Bitcoin, which witnessed a collapse in revenue due to the influx of ASIC miners, the revenue per KH/s for Litecoin has experienced fewer drastic fluctuations.

Acknowledgements

This work was funded in part by the National Science Foundation through CNS-1237264. We are grateful to our system administrators, Cindy Moore and Brian Kantor, who have provided invaluable technical assistance.

References

[1] D. Bestuzhev, "Skypemageddon by Bitcoining," *Securelist - Information about Viruses, Hackers and Spam*, 2013. [Online]. Available: http://www.securelist.com/en/blog/208194210/Skypemageddon_by_bitcoining

[2] J. Caballero, C. Grier, C. Kreibich, and V. Paxson, "Measuring Pay-per-install: The Commoditization of Malware Distribution," in *Proceedings of the 20th USENIX Security Symposium*. Berkeley, CA, USA: USENIX Association, 2011.

[3] N. Christin, "Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace," *CoRR*, vol. abs/1207.7139, 2012.

[4] CIA, "CIA World Factbook, Country Comparison: Internet Users." [Online]. Available: https://www.cia.gov/library/publications/the-world-factbook/rankorder/2153rank.html

[5] D. Dagon, C. C. Zou, and W. Lee, "Modeling Botnet Propagation Using Time Zones," in *Proceedings of the 13th Annual Symposium on Network and Distributed System Security*, 2006.

[6] I. Eyal and E. G. Sirer, "Majority is not Enough: Bitcoin Mining is Vulnerable," http://arxiv.org/abs/1311.0243, 2013.

[7] Financial Crimes Enforcement Network, "Application of FinCEN's Regulations to Persons Administering, Exchanging, or Using Virtual Currencies," Mar. 2013, http://fincen.gov/statutes_regs/guidance/html/FIN-2013-G001.html.

[8] C. Grier, L. Ballard, J. Caballero, N. Chachra, C. J. Dietrich, K. Levchenko, P. Mavrommatis, D. McCoy, A. Nappa, A. Pitsillidis, N. Provos, M. Z. Rafique, M. A. Rajab, C. Rossow, K. Thomas, V. Paxson, S. Savage, and G. M. Voelker, "Manufacturing Compromise: The Emergence of Exploit-as-a-Service," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, Oct. 2012.

[9] Inside Your Botnet, "av.psybnc.cz (100k ngrBot hosted in France Paris Gandi)," http://www.exposedbotnets.com/2011/11/avpsybnccz100k-ngrbot-hosted-in-france.html, Nov. 2011.

[10] ——, "a.xludakx.com (ngrBot hosted in France Paris Gandi around 80k)," http://www.exposedbotnets.com/2011/10/axludakxcomngrbot-hosted-in-france.html, Oct. 2011.

[11] ——, "b.mobinil.biz (Silent BitCoin GPU Miner using Phoenix Miner)," http://www.exposedbotnets.com/2011/07/bmobinilbizsilent-bitcoin-gpu-miner.html, Jul. 2011.

[12] ——, "xD.a7aneek.net (80-100k ngrBotnet hosted in France Paris Gandi)," http://www.exposedbotnets.com/2011/11/xda7aneeknet80-100k-ngrbotnet-hosted-in.html, Nov. 2011.

[13] ——, "beast.darkogard.com (irc botnet hosted in Germany Frankfurt Am Main Sedo Domain Parking)," http://www.exposedbotnets.com/2012/07/beastdarkogardcomirc-botnet-hosted-in.html, Jul. 2012.

[14] ——, "d.xludakx.com (ngrBot hosted in Netherlands Amsterdam Leaseweb B.V.)," http://www.exposedbotnets.com/2012/01/9521116562ngrbot-hosted-in-netherlands.html, Jan. 2012.

[15] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, V. Paxson, G. M. Voelker, and S. Savage, "Spamalytics: an Empirical Analysis of Spam Marketing Conversion," in *Proceedings of the ACM Conference on Computer and Communications Security*, Alexandria, VA, Oct. 2008, pp. 3–14.

[16] C. Kanich, N. Weaver, D. McCoy, T. Halvorson, C. Kreibich, K. Levchenko, V. Paxson, G. M. Voelker, and S. Savage, "Show Me the Money: Characterizing Spam-advertised Revenue," in *Proceedings of the USENIX Security Symposium*, San Francisco, CA, Aug. 2011.

[17] B. Krebs, "Botcoin: Bitcoin Mining by Botnet," http://krebsonsecurity.com/2013/07/botcoin-bitcoin-mining-by-botnet/, 2013.

[18] J. Kroll, I. Davey, and E. Felten, "The Economics of Bitcoin Mining, or Bitcoin in the Presence of Adversaries," in *Proceedings of WEIS 2013*, 2013.

[19] P. Krugman, "The Antisocial Network," *The New York Times*, April 2013.

[20] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. Voelker, and S. Savage, "A Fistful of Bitcoins: Characterizing Payments Among Men with No Names," in *Proceedings of the ACM Internet Measurement Conference*, 2013.

[21] B. Miller, P. Pearce, C. Grier, C. Kreibich, and V. Paxson, "What's Clicking What? Techniques and Innovations of Today's Clickbots," in *Proceedings of the 8th International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 164–183. [Online]. Available: http://dl.acm.org/citation.cfm?id=2026647.2026661

[22] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008, http://www.bitcoin.org/bitcoin.pdf.

[23] R. Naraine, "Researchers Find Malware Rigged with Bitcoin Miner," 2011, http://www.zdnet.com/blog/security/researchers-find-malware-rigged-with-bitcoin-miner/8934.

[24] C. Percival and S. Josefsson, "The scrypt Password-Based Key Derivation Function," http://tools.ietf.org/html/draft-josefsson-scrypt-kdf-01.

[25] F. Reid and M. Harrigan, "An analysis of anonymity in the Bitcoin system," in *Security and Privacy in Social Networks*, Y. Altshuler, Y. Elovici, A. B. Cremers, N. Aharony, and A. Pentland, Eds. Springer New York, 2013, pp. 197–223. [Online]. Available: http://dx.doi.org/10.1007/978-1-4614-4139-7_10

[26] D. Ron and A. Shamir, "Quantitative Analysis of the Full Bitcoin Transaction Graph," in *Proceedings of Financial Cryptography 2013*, 2013.

[27] C. Rossow, D. Andriesse, T. Werner, B. Stone-Gross, D. Plohmann, C. J. Dietrich, and H. Bos, "SoK: P2PWNED: Modeling and Evaluating the Resilience of Peer-to-Peer Botnets," in *Proceedings of the IEEE Symposium on Security and Privacy*, 2013.

[28] Securities and Exchange Commission, "Memorandum Opinion Regarding the Court's Subject Matter Jurisdiction, Securities and Exchange Commission v Shavers et al, TXED 4:13-cv-416, Docket 23," Aug. 2013, http://www.archive.org/download/gov.uscourts.txed.146063/gov.uscourts.txed.146063.23.0.pdf.

[29] B. Stone-Gross, R. Abman, R. Kemmerer, C. Kruegel, D. Steigerwald, and G. Vigna, "The Underground Economy of Fake Antivirus Software," in *Proceedings of WEIS 2011*, 2011.

[30] C. Thompson, "BitCoin is Gold 2.0: Venture capitalist," *CNBC.com*, April 2013.

[31] J. Umawing, "Fareit Goes Bitcoin Mining (ThreatTrack Security Labs IT Blog)," 2013, http://www.threattracksecurity.com/it-blog/fareit-goes-bitcoin-mining/.

[32] G. Williams, "Should you invest in Bitcoin?" *U.S.News & World Report*, May 2013.

[33] J. Wyke, "The ZeroAccess Botnet - Mining and Fraud for Massive Financial Gain," SophosLabs, Tech. Rep., 2012. [Online]. Available: http://www.sophos.com/en-us/why-sophos/our-people/technical-papers/zeroaccess-botnet.aspx

[34] C. Xiao, "Bitcoin Miner Malware." [Online]. Available: http://www.antiy.net/p/bitcoin-miner-malware/