An Automatic Detection and Analysis of the Bitcoin Generator Scam

1st Emad Badawi
Faculty of Engineering
University of Ottawa
Ottawa, Canada
ebada090@uottawa.ca

2nd Guy-Vincent Jourdan
Faculty of Engineering
University of Ottawa
Ottawa, Canada
gjourdan@uottawa.ca

3ed Gregor Bochmann
Faculty of Engineering
University of Ottawa
Ottawa, Canada
Bochmann@uottawa.ca

4th Iosif-Viorel Onut
IBM Centre for Advanced Studies
Ottawa, Canada
vioonut@ca.ibm.com

***Abstract*—We investigate what we call the "Bitcoin Generator Scam" (BGS), a simple system in which the scammers promise to "generate" new bitcoins using the ones that were sent to them. A typical offer will suggest that, for a small fee, one could receive within minutes twice the amount of bitcoins submitted. BGS is clearly not a very sophisticated attack. The modus operandi is simply to put up some web page on which to find the address to send the money and wait for the payback. The pages are then indexed by search engines, and ready to find for victims looking for free bitcoins.**

We describe here a generic system to find and analyze scams such as BGS. We have trained a classifier to detect these pages, and we have a crawler searching for instances using a series of search engines. We then monitor the instances that we find to trace payments and bitcoin addresses that are being used over time. Unlike most bitcoin-based scam monitoring systems, we do not rely on analyzing transactions on the blockchain to find scam instances. Instead, we proactively find these instances through the web pages advertising the scam. Thus our system is able to find addresses with very few transactions, or even none at all. Indeed, over half of the addresses that have eventually received funds were detected before receiving any transactions.

The data for this paper was collected over four months, from November 2019 to February 2020. We have found more than 1,300 addresses directly associated with the scam, hosted on over 500 domains. Overall, these addresses have received (at least) over 5 million USD to the scam, with an average of 47.3 USD per transaction.

***Index Terms*—Cryptocurrency, scam analysis, cyberattack,fraud detection, bitcoin, blockchain analysis, data mining**

1. Introduction

In recent years, there has been a rise in the use of cryptocurrencies as an investment platform [1]. At the time of writing, there are 5,140 different cryptocurrencies, with a capitalization market of around 300 billion US dollars [2]. The most popular cryptocurrency is Bitcoin, which has a capitalization market of around 170 billion US dollars [2].

Bitcoin [3] is a decentralized cryptocurrency that has become popular in the last ten years. It is a peer-to-peer electronic currency that can be exchanged between users without the involvement of a trusted authority such as an administrator or a central bank [3]–[5]. Bitcoin has two key features: Transparency and Pseudo-anonymity [4]–[6]. It is transparent because the transactions are publicly announced in a decentralized ledger called a blockchain.

Bitcoin pseudo-anonymity comes from the fact that the users use pseudonyms (addresses). These pseudonyms are not related to individuals, and they are computed from the user's public key. Moreover, Bitcoin addresses can be generated at will. As a result, users can generate unique addresses for each transaction. This increases privacy by creating an additional firewall to prevent the addresses from being linked to a specific owner [3].

Cybercriminals have leveraged Bitcoin Pseudo-anonymity in their attacks. According to a 2019 report by CipherTrace[1], the value of thefts, hacks, and scams in 2019 was more than twice the value in 2018: more than $4.4 billion was siphoned away from users and exchanges through fraudulent activities in 2019 only.

Cryptocurrency-based attacks take many forms. "High Yield Investment Programs" (HYIP) is one of the popular examples of the scams that cybercriminals carry out [4], [5], [7], [8]. HYIP is a scam in which investors are promised a high-interest rate, e.g., more than 1-2% per day [5]. Perhaps the most famous HYIP scammer is Charles Ponzi, who claimed in the early 1920s to ran an arbitrage; the investors were promised a 50% profit within 45 days, or 100% profit within 90 days. Because of Charles Ponzi, HYIP is sometimes called *Ponzi scheme* [5].

Money laundering (ML) [9], [10] and ransomware [11]–[13] are other popular examples. ML describes the process of disguising the sources of illegal profits generated by criminal activity. Its aim is to hide the link between original criminal activities and the corresponding funds by passing the money through a complex sequence of commercial transactions or banking transfers [9].

Ransomware is a denial-of-access attack in which a malicious piece of software locks and encrypts a victim's device data until a sum of money is paid [13]. Cryptocurrencies, usually Bitcoin, are often used for these payments. Recently, Riviera Beach officials voted to pay 65 bitcoins, worth US$600,000 at the time, to a cybercriminal who seized and shut down the city's computer systems. The re-

1. CipherTrace cryptocurrency Intelligence Report, accessed through https://bit.ly/2ROPYCK.

DOI 10.1109/EuroS&PW51379.2020.00060

sulted outage forced the local fire and police departments to write down hundreds of 911 calls on paper [14].

The current state of the art for bitcoin scam detection usually relies on a classification model to detect scam addresses based on transactions history [4]–[6], [8]. These addresses are either collected manually, e.g., by searching on bitcoin discussion forums such as bitcointalk.org [8], or they come from semi-automated web crawls of the same forums, followed by manual addresses collection [4]–[6]. Once a set of addresses used in the scam has been collected, the transaction history of these addresses is used to train a classification model [5], [8]. The classifier is trained on features such as the frequency of transactions, the ratio of received/sent transactions to all transactions, the address lifetime, or the "payback" ratio, which is the ratio of addresses that appears in the input and output sides of address transactions.

However, the increasing number of transactions recorded on the blockchain[2] makes it difficult to extract meaningful patterns that can be used in fraud detection [8]. Additionally, these methods, based on transaction history, are by nature only able to detect a scam address after the fact, once some victims have been defrauded.

In this paper, we look at a scam that has emerged with the rise of cryptocurrencies. We call this attack the "Bitcoin Generator Scam" (BGS). In BGS, the attackers claim that they will provide free bitcoins in return for a small mining fee, using dubious claims such as their ability to "hack the blockchain ledger". BGS attacks start with an online website targeting their victims. We call these websites "generators". These generators are carefully designed web pages that attempt to convey to the victim the advanced technical abilities of the scammer and a large, satisfied user base for the BGS instance. Some BGS instances display a fake chat box, and a pop-up showing claimed current users and the number of mined bitcoins they supposedly gained.

BGS attacks can be directly advertised e.g., on social media. Still, victims can also be actively seeking easy profit by looking online for "Bitcoin hack services" using search engines, social media, streaming sites, blogs, etc. (Figure 1, image 1). The search results may link directly to BGS instances, or results may link to pages that have links to BGS instances. Once a BGS instance like the one shown in Figure 1 image 2 is accessed, the victim is asked to provide the number of coins they want to mine, and the bitcoin address in which the mined coins will be deposited. Once the victim provides the information, the BGS pretends to perform some "hacking" (Figure 1, image 3). Finally, some success message is displayed, and the victim is asked to pay a mining fee to collect the funds (Figure 1 image 4). In many cases, the fees are a fixed number of satoshis. In other cases, the attacker promise that the victim will receive some multiple of the amount they pay.

Some authors (e.g. [15]) characterize Ponzi scheme by their pyramidal structure and the payout to existing investors using funds from new investors. By this definition, BGS is not a Ponzi scheme, since most BGS do not require investors to enrol new investors, and as discussed in Section 4.2 we usually do not find any evidence of payout at all. However, some other authors characterize Ponzi schemes by their extremely high rates of return [7], [16], and BGS certainly fall under that category, with advertised return rate in the range of 100% in 24 hours.

We have created an automated system to find BGS and extract the bitcoin addresses they use. To do that, we have first collected 330 BGS instances from an initial manual search and by using blacklisted domains [17], cutestat.com, and the Internet archive [18].

We used this initial list to train a classifier to recognize BGS instances. The 10-fold cross-validation of the performance of our classifier on our initial dataset has shown that 98.7% of BGS instances are correctly classified while maintaining a false-positive rate of less than 1%. We have then generated hundreds of search queries related to BGS and used them daily on popular search engines for four months.

We use our classifier to detect BGS instances in the pages directly returned by the search engines, and in pages they link to. When we identify a new BGS instance, we interact with it to extract the bitcoin address(es) used to accept money from the victims. In the four-month period of our study, we have discovered more than 500 scam domains, and more than 1,300 bitcoin addresses associated with them. These addresses have received more than 5 million dollars, with an average of 47.3 dollars per transaction.

Since our approach is not based on existing transactions, we are able to detect scam addresses before they receive any money. Indeed, over half of the addresses that have eventually received funds in our study were detected before receiving any transactions.

The remainder of this paper is structured as follows. After this introduction, in Section 2, we detail our methodology. In section 3, we report some basic numbers obtained during our crawling period. In Section 4, we carry out various analyses and discuss the results. A literature review is provided in Section 5. In Section 6 we discuss some of the main limitations in our model, as well as future enhancement and analysis. Finally, we conclude in Section 7.

2. Methodology

We have developed a data-driven approach to collect, detect, and analyze BGS. We initially started our work with a manual search for BGS pages to obtain a representative dataset on which to train our model. This also helped us getting an initial broad understanding of the scam and provided the source we needed to automate effective BGS search queries. Figure 2 describes our complete system, which includes five modules: Search query generator, Web Crawler, Classifier, BGS instances triggering, and Analysis. In this Section, we provide a general overview of the general system before diving into more details in Sections 3 and 4.

2.1. Dataset Construction

The first step is to collect an initial set of BGS instances that can be used to train our classification model,

2. Over half a billion transactions at the time of writing:https://bit.ly/39VjOwT.

Figure 1. An Example of BGS Attack.

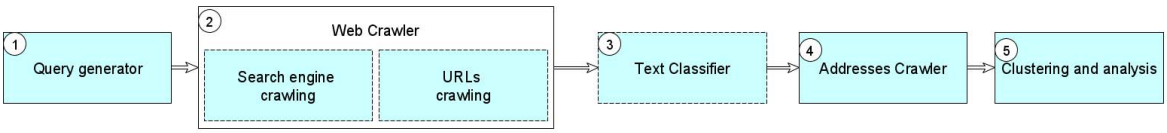

Figure 2. Bitcoin Generator Scam detection and analysis model.

and extract search queries and use it to search for more scam instances using different search engines. We used an array of techniques to collect this initial dataset:

1) **Google search:** We started our data collection by manually searching for BGS instances on the Google search engine. We have used several search queries related to the scam, such as "free bitcoin generator", "BTC generator tool", and "earn free bitcoins instantly", and identified an initial set of 52 BGS instances this way. We also obtained 30 new search queries using Google's automatic "related search" suggestions while doing this initial collection. This gave us our initial set of queries for starting our automated web crawl.
2) We also used the site **Bitcoin.fr [17]:** it contains a list of Bitcoin and cryptocurrency scam domains. That list was collected using the testimonies of the site users, as well as a collection of several other scam lists, including CryptoFR, the House of Bitcoin, adcfrance.fr, scambitcoin.com, and badbitcoin.org. At the time of crawling, the list contained 6,230 domains[3].
3) Another site that we used is **CuteStat.com:** it is a web site that collects information related to domains, websites, usage reports, IP address, host, etc. One service that this website provides is a list of up to 100 domains that have content "related" to the search you have performed. We have collected this way 610 new domains that are related to the searches we have done using the queries collected in step 1.
4) Finally, we have used the **Internet Archive [18]:** this is a digital library that provides a large collection of readily available digitized materials, including internet sites, games, music, and public-domain books for free. We have used the Archive to collect thousands of snapshots for the set of domains we collected from *Bitcoin.fr* and *CuteStat.com*.
5) **Identifying BGS Instances:** Since the number of snapshots returned by the Archive is in thousands, and we could not manually check every one of them, we limited our search to the snapshots in which we could detect a bitcoin address right inside the HTML. This reduced the number of possible BGS domains to only 307, a number we could handle manually. We inspected these domain snapshots one by one. In the end, 252 of these domains were indeed BGS instances, while the other 55 other domains where different types of scams such as bogus charity and HYIP.

Following these steps, we have collected our initial set of BGS instances, containing 304 pages. In order to create a benign dataset for training our model, we then manually inspected 400 randomly selected pages that we had collected but not flagged in the process during the first week of operation. This gave us 374 benign pages but also 26 new BGS instances. Therefore, our final dataset consists of 330 BGS pages, complemented with 330 benign pages selected from the set of 374 pages we had.

3. We are planning on providing to *Bitcoin.fr* the list of domains that we have collected during this study.

2.2. Search Query Generator

Finding good search queries that have a high likelihood of leading to the scam pages is an important task. Kharraz *et al.* [19] used Google Trends service to generate such queries, Srinivasan *et al.* [20] used a probabilistic analysis, and Badawi *et al.* [21], [22] used both techniques. In our work, we first collected Google's automatic search suggestions as we manually searched for BGS. We then used these suggestions to create the first set of queries and perform an initial web crawl.

As described in Section 2.1, we were able to collect and manually verify 330 BGS instances from our initial web crawling, as well as from a list of blacklisted domains [17], from the site cutestat.com, and from the Internet archive [18]. Finally, the contents of the "Keywords" meta tag were extracted from these instances to augment our original queries. The "Keywords" meta tag represents a set of a comma-separated list of keywords that are relevant to the web page, and used to inform the search engines about its content [23]. Our final query list contains 539 search queries[4].

2.3. Web Crawler

Our crawler uses the previously identified queries to search daily for BGS pages using *Google.com*, *Bing.com*, *search.yahoo.com*, and *search.1and1.com*. For each query, we only consider the first and second pages returned by each search-engine (that is, 20 search results). The crawler is based on ChromeDriver[5] and Python Selenium[6]. Using Python beautifulsoup[7] and the CSS selectors, we extract and crawl the URLs resulting from our searches. For the crawling process, we use a lightweight scripted headless browser built using python by integrating Selenium, ChromeDriver, and BeautifulSoup. We automatically collect data about the URLs that we crawl, including any URL redirections, the HTML content, a screen-shot of the landing page, and its resources (scripts, CSS files, etc.).

2.4. Classification Module

In our crawling process, the majority of the URLs returned by the search engines or crawled by our system are either benign pages having nothing to do with BGS, or are non-BGS pages that have a link to one or more BGS instances. Since we are building an automated system, we need a classifier to automatically decide which ones of the URLs we visit are genuine BGS instances. To that extent, we have developed a text-based classification model. Our classification model achieved good accuracy, with a True Positive Rate (TPR) above 98.7% and False Positive Rate (FPR) lower than 1%.

In our analysis, True Positives (TP) are BGS pages actually classified as BGS, True Negative (TN) are benign pages classified as benign, and thus False Positive (FP) are benign pages wrongly classified as BGS and False

4. The complete list is available at https://bit.ly/bgsieeesb2020.
5. http://chromedriver.chromium.org/
6. https://selenium-python.readthedocs.io
7. https://pypi.org/project/beautifulsoup4/

classifier	page Type	classified Clean	classified BGS	F1
SVC	clean	327	3	98.92
	gen	4	326	
MLP	clean	327	3	98.92
	gen	4	326	
RF	clean	329	1	95.9
	gen	25	305	
NB	clean	327	3	96.58
	gen	19	311	
KNN	clean	319	11	97.9
	gen	3	327	

TABLE 1. RESULTS OF A 10-FOLD CROSS-VALIDATION WITH FIVE CLASSIFIERS.

	Actually Clean	Actually BGS pages
Classified clean	99	1
Classified BGS	2	98

TABLE 2. CLASSIFIER ACCURACY ON PAGES THAT HAVE NOT BEEN OBSERVED IN THE TRAINING PHASE.

Negative (FN) are GCS instances wrongly classified as benign. As usual, the F1 score is derived as follows:

$$F1 = 2 * (Precision * Recall)/(Precision + Recall)$$

where, $Precision = TP/(TP + FP)$ and $Recall = TP/(TP + FN)$. The higher $F1$, the better.

Text Classifier We tested five different classifiers from Scikit-learn python library [24] on our training set: Support Vector Classifier (SVC), Naive Bayes(NB), k-nearest neighbors (KNN), Random Forest (RF), and Multi-layer Perceptron (MLP).

Features Processing: We trained our models using the text presented to the users on the web pages, more specifically the term frequency-inverse document frequency (TF-IDF) of the words displayed to the users. We used Scikit-learn TfidfVectorizer [24], which creates the TF-IDF matrix from a collection of raw documents. We use TF-IDF to scale down the impact of less informative tokens that occur very frequently in our dataset.

To evaluate each of these classifiers, we used 10-fold cross-validation on the labeled dataset we prepared in section 2.1. We show the results in Table 1. As can be seen, SVC and MLP achieved the highest F1 score, 98.92, followed by KNN at 97.9. The other classifiers also performed fairly well with RF having the lowest F1 score. Based on these results, we used the SVC classifier throughout our experiments.

After using our classifier on newly found pages for a few days, we randomly selected 100 pages classified as benign and 100 pages classified as BGS instances, for manual verification. Our model correctly classified 197 of these 200 pages. Two benign pages were misclassified as BGS which yields a true positive rate of 98%, and a one BGS instance was misclassified as benign which yields a true negative rate of 99%.

2.5. Interacting with BGS instances

To collect the cryptocurrencies addresses that the scams are using, we need to interact with the BGS instances, provide the expected inputs, and follow the specific instructions in order to reach the final stage (the fourth image of Figure 1), where the scam address is provided: this is the address to which the scam asks the victim to transfer money. Usually, this process requires 5 to 10 minutes on average. During this time, the attacker typically displays a detailed "log" of the hacking process, which is supposed to be taking place in real-time (see, for example, the third image of Figure 1). This log displays bogus proxy servers names, server IP addresses allegedly being hacked, the ledger's block in which the transaction is supposed to be added, etc[8]. Furthermore, some BGS also display a fake chat box, and a pop-up showing claimed current users and the number of mined bitcoins they supposedly gained. The goal of this "live" fake data is presumably to help persuade the victim that the site is effective and encourage them to transfer the fee. In some cases, however, the scam address can be found immediately in the HTML of the BGS instance. For these pages, we collect the scam address without further interaction with the BGS instance. Furthermore, in addition to the "live" crawling explained above, we also crawl the Internet archive [18] to collect addresses that have been used by the instance in the past.

2.6. Crawler Effectiveness

In this section, we discuss the ability of our crawler to detect scam addresses before anyone is victimized. Overall, by crawling the online instances and the snapshots taken by the Internet archive, we have detected 1,302 BGS Bitcoin addresses with at least one transaction. 480 of these addresses (36.86% of the total) are not in the Internet archive. 776 of these addresses (59.6% of the total) have been extracted from the Internet Archive but were never found by our live crawler. Finally, the remaining 46 addresses (3.53% of the total) are been found both on the internet archive and by our live crawler. 290 of the 526 addresses found by our live crawler were found before they had any transaction; the transactions started after the addresses had been detected by our system. That is one of the unique strengths of our model, the ability to detect a problematic address before it is recorded on the blockchain.

By cross-checking all the addresses that have at least one transaction with the addresses in WalletExplorer.com, we find that 54 of the 1,302 bitcoin addresses (4.14%) belong to online wallet services. Figure 3 shows the breakdown of BGS addresses per online wallet services: most of them use CoinJar.com and MoonBit.co.in. These results are similar to the one reported by *Toyoda et al.* in [5]: in their study, online wallet services controlled 4.6% of addresses of HYIP operators they collected.

3. Scam Collection and Measurement

Our experiments were run on our university's server as well as on dedicated servers provided by Compute Canada [9]The results reported in this paper come from data collected from November 2019 to February 2020. In this section, we present some basic numbers obtained directly from our crawler and classifier.

8. A complete example of one such log is presented in our public data repository.

9. https://www.computecanada.ca/research-portal/

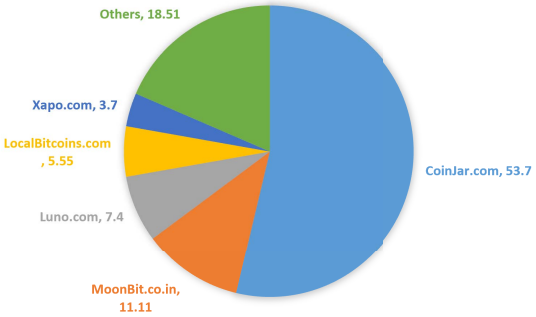

Figure 3. Breakdown of online wallets addresses used by BGS.

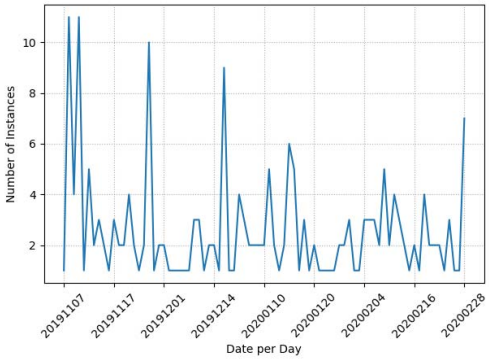

Figure 4. Number of BGS URLs detected per day.

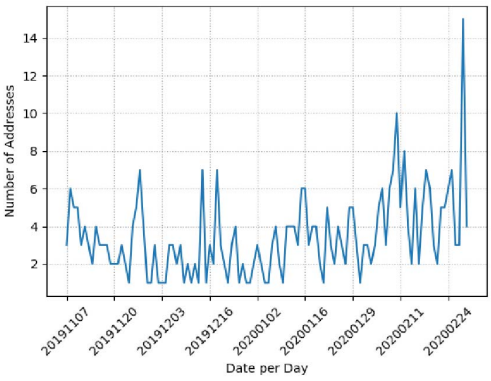

Figure 5. Number of Bitcoin addresses detected per day.

Over the course of our experiments, our system identified more than 1,300 cryptocurrencies addresses (almost all bitcoins) used by BGS instances and that have at least one transaction. These addresses have been found on 504 unique scam domain names[10]. On 55,4% of these domains (280 of them), we only find one address. At the other extreme, 27 of these domains (5.34%) are associated with ten addresses or more. Also, about 99% of the addresses that we have found are bitcoin addresses: only 14 addresses are for other cryptocurrencies. Seven are Ethereum addresses, four are Bitcoin Cash (BCH) addresses, two are Ripple addresses, and one is a Litecoin (LTC) address. Therefore, the rest of our analysis only uses bitcoin addresses. Finally, our analysis also showed that none of the Alexa top 1K domains[11] contain actual BGS instances. Therefore, we only report results for URLs hosted on domains outside Alexa top 1k.

In Figures 4 and 5, we present the number of BGS URLs and addresses detected per day, excluding URLs and addresses found in the internet archive. That is, these figures only include newly discovered and currently active BGS instances.

On average, our model detected about 2 new BGS instances and 3.3 new bitcoin addresses every day during the period of November 2019 to Feb 2020 (note that a new BGS instance does not necessarily mean a new address since there are some addressed that are shared among instances). These numbers are relatively stable throughout the period. Therefore, we can extrapolate that our system will identify more than 500 new BGS instances and more than 1,000 bitcoin addresses within a year of crawling if the situation does not evolve.

We use the BGS instance identified during our crawls by our classifier to find bitcoin addresses used by the scammers. We have two ways of doing this: first, we keep interacting with the known BGS instance daily once it has been added to our database. This way, if an instance publishes new addresses, our system will pick them up within 24 hours. We also look back in time, thanks to the Internet archive [18]. We thus collect some of the addresses that have been used by the instance in the past, before we discovered it. Our database is therefore a mix of currently active addresses and addresses that have been active months or years ago.

Overall, our model was able to collect 6,395 bitcoin addresses involved in BGS. 1,302 of these addresses have at least one transaction. However, one particular BGS instance is responsible for most of the transaction-less instances: the domain *bitmake.io* has a hard-coded list of 5,001 addresses, and one of these addresses is selected randomly when a payment is made. At the time of writing, on that particular BGS instance, only 30 of the 5,001 addresses have transactions, so that site alone is the source of 4,971 of the 5,093 transaction-less addresses in our database (that is 97.6% of them). Without that site, over 91% of the addresses have transactions.

4. Analysis

In this section, we use the data collected in section 3 to estimate how much money was made through the BGS. We also look at the few cases in which bitcoins were actually transferred back to the initial address. Finally, we discuss a couple of techniques used by scammers that are making systems like ours less effective.

10. In general, we only consider second-level domain names when comparing scams URLs, excepted for hosting services, for which we consider the third-level domain name. So *btctool.epizy.com* and *generatorbitcoin.epizy.com* are counted as two separate attacks even though they are on the same second-level domain name because they are both using the hosting service *epizy.com*.

11. https://www.alexa.com/

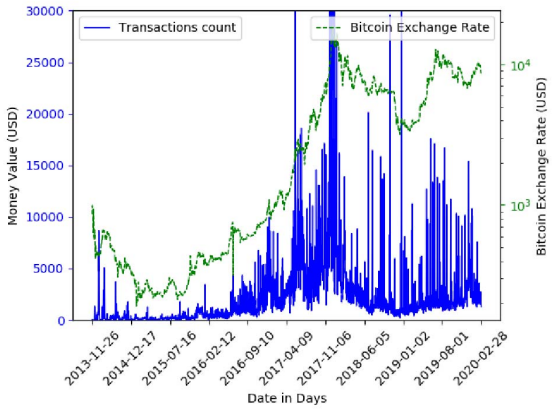

Figure 6. Daily deposited money to BGS addresses.

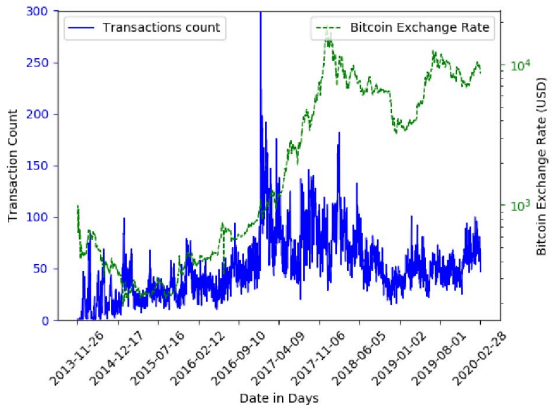

Figure 7. Daily incoming transactions to BGS addresses.

4.1. Bitcoin Addresses Payment Analysis

We now focus on the scale of BGS by analyzing the transactions recorded on the addresses we have collected. As previously mentioned, there are 1,302 addresses with at least one transaction. Overall, these 1,302 addresses have received 106,652 transfers from 182,715 unique addresses. In total, these addresses have received 2,329.94 bitcoins, with an average of 0.021846 bitcoin per transaction.

To convert the amounts of the transactions to USD, we use the average exchange rate of the day of the transaction, obtained from *bitcoincharts.com*. Overall, the BGS addresses in our collection gathered 5,098,178 US dollars from November 2013 until February 2020. Figures 6 and 7 depict respectively the total number of transactions and their corresponding total value in USD, compared to the exchange rate of bitcoin. As can be seen in the figure, there is a clear apparent correlation between the market value of bitcoin and the success of BGS, which certainly is not surprising. Although currently less active than it was during peak bitcoin value, BGS is still going steady and continues receiving money on a daily basis. We also see in Figure 8 that scammers tend to remove the funds from the receiving address without delay.

If we try to estimate the accuracy of our numbers: the total number of addresses and of instances are reported as is, without extrapolation (e.g. using techniques such as multi-input heuristic clustering algorithm [25]). It is thus an underestimate of the actual total number of addresses and instances, since we certainly have not detected a number of them. In terms of future addresses and instances to expect, one should not use the total numbers that we have found over 4 months as an indication of future growth since they contains historical data. Instead, one should use the numbers discussed in Section 3 which excludes these historical figures. Finally, for the total amount of dollars, our figure might be overestimated since we cannot distinguish between the payments made to the scam and the payments made to the same address for some other reasons.

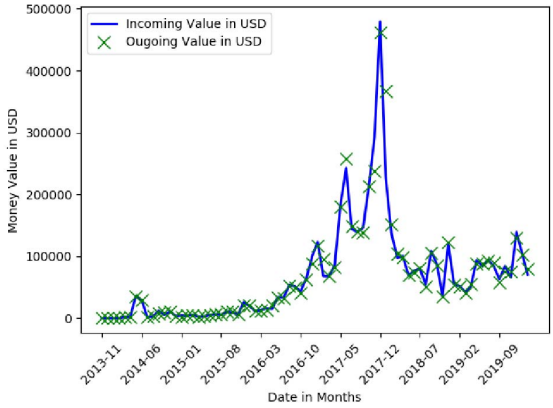

Figure 8. A comparison between the Incoming and Outgoing deposits in BGS addresses.

4.2. Payback Analysis

We also looked at the possibility of actual payback in the scam; in the spirit of a Ponzi scheme for example. With the anonymity of bitcoin, it is unclear that paying back some of the potential victims in order to attract more would be very effective, and indeed our analysis does not show any evidence that this is happening.

We look at the subset of the BGS addresses that received bitcoins from a specific address[12] and sent anything back to the same address, the way it would be if the scam was actually working as advertised. Out of the 1,302 scam addresses in our database, 56 (4.3%) have sent bitcoins back to an address from which they had received bitcoins before

Naturally, we would like to understand if those 56 addresses did indeed send money back to victims, or if both addresses belong to the scammer. To do that, we look at the transaction history of these 56 addresses and

12. A limitation of this analysis is that we rely on the payback to be made to the same address from which the initial payment was made. This is not necessarily the case in a Bitcoin transaction [26], and we would miss the hypothetical transactions for which this is not the case. On the other hand, most of the scams examples that we surveyed did not include any other way to get a payment back.

	Total scam	Addresses that received back any payment		Addresses that did not receive back anything
		Scam	Normal	
#addresses	182,715	19	584	182,112
#inTransac	106,652	191	1,317	104,927
#outTransac	42,440	73	1,479	40,888
#inBTC	2,329.9	10.89	101.86	2,217.15
#outBTC	2,321.57	7.6	113.59	2,200.38
#inUSD	$5,098,178	$7,396	$53,639	$5,037,143
#outUSD	$5,186,658	$9,146	$58,800	$5,118,712

TABLE 3. DETAILED ANALYSIS FOR THE SCAM ADDRESSES PAYBACK (TRANSAC REFERS TO TRANSACTIONS).

discover 603 addresses that have both sent and received bitcoins from at least one of these 56 addresses. We further divide these 603 addresses into two groups:

Scam Addresses: these are the addresses that were already identified as scam addresses by our system. A total of 19 of the 603 addresses are in that category. These do not represent payback, but simply internal transactions in the scam.

Normal Addresses: these are the addresses for which we have no concrete evidence that they are under the control of the scammers (which certainly does not mean that they are not). We have a total of 584 addresses in that group (out of a total of 182,715 addresses). The balance for these addresses is almost on part: collectively, they have sent 101.86 bitcoins and received 113.59 bitcoins. Overall, these 584 addresses sent 53,639 USD and received only 58,800 USD. However, one of these addresses has sent 0.001 bitcoin to a single scam address and received back 15 bitcoins, so excluding that outlier, on average, these addresses have received fewer bitcoins than they have paid. We cannot reach a conclusion about these addresses at this time, but it is clear that overall, they do not impact or change our results, as illustrated in Table 3.

4.3. Evasion Techniques

In this section, we look at two techniques that we have seen being used and it makes our analysis more difficult. First, many of the domains are regularly changing bitcoin addresses during their lifetime. Second, in some cases, the address is not provided directly, and a URL-shortener service is used to deliver the information instead.

The underlying intent of these techniques is difficult to ascertain. As a side effect, it does make it more difficult for automated systems like ours to find the BGS domains and their addresses. It may, in fact, be the reason why scammers do these things to prevent detection and extend the lifespan of their attacks.

Regularly Changing the Bitcoin Address. We have detected the use of at least two different addresses in 217 of the BGS domains that we have found (that is 42% of the domains). In fact, for 27 of these domains (5.3% of the total) we have detected at least 10 bitcoin addresses. The worst offenders have, respectively, 5,001 addresses, 133 addresses, 42 addresses, and 37 addresses that we know of. Of course, using many addresses reduces the number of transactions per address, which in turn makes detection techniques based on transaction history more difficult.

Figure 9 shows an example of a site in which the address presented to the victim is selected randomly from an array of static choices. As already explained, the most extreme case that we have seen is *bitmake.io*, which contains a list of 5,001 addresses, but only 30 of these addresses have any transactions.

Using Shortener Services. Some BGS instances do not provide the payment address directly. Instead, what is provided is a link through a URL shortening service such as Bitly.com. One consequence of this is that automated systems that are relying on detecting the payment address in the page HTML are going to fail on these instances.

As an example, the BGS instance *generatebitcoin.xyz* uses the link *bit.ly/BitcoinMiningFee* to direct victims to the payment address. This technique is not widespread at the moment, and we only have found 6 BGS instances doing this so far.

5. Related Work

Bitcoin's pseudo-anonymity and its ability to use new addresses for each transactions have been some of the reasons why miscreants have been using it for their attacks, e.g., for money laundering [9], [10], ransomware [12]–[14], and, perhaps most related fraud to our work, High Yield Investment Programs (HYIP) [4], [5], [7], [8]. On the other hand, researchers have leveraged the publicly readable blockchain transaction records to analyze these fraudulent activities [4], [5], [7], [8].

State of the art in academic work on bitcoin scam detection is usually based on some manual collection of addresses involved in the scam. The starting point could be a manual search on a forum in which the attack is being discussed, e.g., bitcointalk.org [8], or it could by a semi-automated crawl of that same forum, followed by manual addresses collection [4]–[6]. One scam addresses have been collected, their transaction histories are used to extract distinguishing features and tell benign addresses apart from scam addresses [4]–[6], [8], [27]. These features are then used to train a classifier [5], [8].

Toyoda *et al.* [27] proposed a binary classifier to identify the bitcoin addresses of HYIP operators. They manually scrapped over 1,500 addresses related to the topic of HYIP on BitcoinTalk and blockchain.info. They then analyzed these addresses transaction history and identified a set of features to learn from. They trained an RF classifier on these features, showing that it could identify HYIP addresses with a true positive rate of 83% and a false positive rate of 4.4%.

In a follow-up work [6], the authors extended the features and proposed a multi-class bitcoin-enabled service identification to classify bitcoin addresses into one of seven major "services" such as exchange, gambling, mixer, and scam. Furthermore, the authors proposed to extract the features based on the address or based on the owner. In the latter situation, the features are extracted from a group of addresses that are deemed to be owned by the same person. The authors used the multi-input heuristic clustering algorithm. The owner-based classification achieved a higher accuracy of 72% compared to 70% to the address-based classification.

```
// ============================================
// Bitcoin Crypto Finder Address Database
// ============================================
var Address=new Array() // do not change this
Address[0] = "<input type='text' class='form-control' size='38' onClick='this.select();' class='depositAddy'
value='1GjBvUuAqchiurAEwiPqFjXNx7FEyk8uNQ' readonly>";
Address[1] = "<input type='text' class='form-control' size='38' onClick='this.select();' class='depositAddy'
value='1N3h3bqb7srRnEssFTMR9UntYNwAnnQs4u' readonly>";
Address[2] = "<input type='text' class='form-control' size='38' onClick='this.select();' class='depositAddy'
value='1Ctf2oqrfRb7yb81ohgp8BFcW1zTB4rGQz' readonly>";
Address[3] = "<input type='text' class='form-control' size='38' onClick='this.select();' class='depositAddy'
value='1CxSoZvgra4J5eeKjLpruUjvzWFSPtGQsx' readonly>";
Address[4] = "<input type='text' class='form-control' size='38' onClick='this.select();' class='depositAddy'
value='1PjGqm9FLepiDqrpKJvh6b5e5DyZNCjG9Q' readonly>";
// ====================================
// Do not change anything below this line
// ====================================
var Q = Address.length;
var whichAddress=Math.round(Math.random()*(Q-1));
function showAddress(){document.write(Address[whichAddress]);}
showAddress();
```

Figure 9. A real world example of a BGS instance in which the payment address is selected randomly from a list.

In [5], Toyoda *et al.* proposed an improvement to their previous work [27]. In the new, improved model, the authors proposed an automatic crawler to collect possible posts that advertise HYIP schemes. They then manually inspected these posts and extracted more than 2,000 addresses of HYIP operators. Their new model achieved a true positive rate of 95% and a false positive rate of less than 5%. Furthermore, they stated that limiting the number of transactions taken into account per address to only a 100 transaction reduces the features computation time while maintaining good detection accuracy.

In [28], a model based on time series analysis to detect anomalies in the transactions history was proposed. The authors suggested that extracting numeric features, such as the frequency of in/out transactions by sliding windows can be used to detect anomalies. They have shown that their model was able to detect the major anomalies in the Pirate@40 HYIP scheme. However, as stated by the authors, this model is more useful in real-time sales inference of marketplaces and digital forensics.

Bartoletti *et al.* [8] manually collected 32 addresses of HYIP operators using old posts on Reddit and bitcointalk.org. They have used these addresses along with 6,400 non-HYIP addresses to train a binary classifier based on features extracted from the addresses transaction history. Their analysis showed that using a random forest classifier along with a "cost-sensitive" approach that gives a higher penalty for misclassifying an HYIP address than for misclassifying a legitimate address achieves the best detection accuracy. Additionally, the authors used a multi-input heuristic clustering algorithm [25] to collect additional addresses controlled by the same scammers. Their analysis showed that more than 50% of the scammers use more than one bitcoin address, for a total of 1,211 addresses. Overall, the scammers received 10 million dollars to the addresses they control.

Our approach is different from all of these previous ones. We do not base our analysis only on previously reported campaigns. Instead, we actually search for new, previously unreported instances. What is more, at this stage, we do not use existing transactions in the detection phase, which allows us to find addresses that do not have any payment yet.

Although recent papers have provided important insights into different types of fraudulent activities related to bitcoin, we are not aware of any study focusing on what we have called the "Bitcoin Generator Scam" before this one.

6. Limitations and Future Work

One of the main limitations of our study is that we only look for BGS instances based on the ones we have already found. Thus, some of our current results may be biased by the type of BGS instances we are looking for, and a more systematic search would shed new light on the situation. For example, by improving our search queries, new and different BGS instances might come to light.

Another limitation is that we have trained our classifier on pages with English text. Thus, our crawler and our results only deal with English instances of BGS. That certainly doesn't mean that the scam is not active in other languages, and we would have overlooked these instances in that case.

Finally, we depend on text classification to detect BGS instances. However, this type of classification can be evaded relatively easily. We could enhance our feature set to be less dependant on the text that is being presented to the user. In our future work, we will build a more accurate text classification model and we will add some non-text-based feature. Additionally, we would like to further explore the relations between the scam addressees. For this purpose, clustering techniques such as multi-input heuristic clustering can be used [5], [8].

Furthermore, we would like to bring in some of the detection techniques based on transaction history described in the literature to try to add new addresses and find new transactions related to BGS. This will also help in our analysis of the links between different BGS instances. Moreover, our analysis have showed that some attackers use the same address in different BGS domains to carry out their fraud (we have already 85 addresses that have been detected on more than one BGS domain). We would like to also look at addresses reuse in other types of scams as well. We have already crosschecked our addresses with a public dataset maintained by the authors of [8]. We have

7 addresses in common with this dataset. 3 are flagged as being used in Ponzi schemes, while 4 are not classified. In our future work, we will analyze the relationship between these domains and their addresses in more details. We will also do further analysis to distinguish between BGS instances that promise to generate bitcoins by hacking blockchain and doubler instances with characteristics similar to Ponzi schemes.

7. Conclusion

In this paper, we investigated what we call the "Bitcoin Generator Scam", a scam in which victims are looking for a quick and easy way to make money on bitcoin are tricked into thinking that some "hacks" can be used to somehow generate bitcoins cheaply. The scam is being advertised through web pages. We have created a system that proactively looks for these web pages and monitor their evolution over time. Because we are able to go straight to the source of the scam, we can avoid trying to detect evidence of the scam on the blockchain, which is typically difficult, error-prone, and only works for addresses on which payments have already occurred. We also innovate with the source of information we use: in addition to using traditional search engines, we showed that services such as the Internet Archive and CuteStat.com can be used to increase the number of instances found significantly.

Our data collection spanned four months; in that time, we collected 6,395 Bitcoin addresses associated with more than 500 unique scam domains. These addresses have received 5,098,178 US dollars, with an average of 47.3 dollars per transaction.

Our study is, as far as we know, the first in-depth look at the BGS. Nevertheless, we believe that our main contribution is our automated approach, which can be used on different scams that also have a significant web presence. By actively looking for instances and using machine learning to classify the results, we were able to discover 6,395 addresses directly advertised by scam instance. This is a much greater number of addresses than usually found in state-of-the-art research, where typically direct evidence is manually collected, and the bulk of the addresses come from "multiplier" techniques such as the multi-input heuristic clustering algorithm [25].

All the data used in our study is freely available at https://bit.ly/bgsieeesb2020.

Acknowledgments. This work was supported in part by Canada's Natural Sciences and Engineering Research Council. The authors would like to thank Marie Vasek and the anonymous referees for their valuable suggestions.

References

[1] J. Kamps and B. Kleinberg, "To the moon: defining and detecting cryptocurrency pump-and-dumps," *Crime Science*, vol. 7, no. 1, p. 18, 2018.

[2] CoinMarketCap, "Cryptocurrency market capitalizations," https://coinmarketcap.com/, 2020.

[3] S. Nakamoto and A. Bitcoin, "A peer-to-peer electronic cash system," *Bitcoin.–URL: https://bitcoin. org/bitcoin. pdf*, 2008.

[4] M. Vasek and T. Moore, "Analyzing the bitcoin ponzi scheme ecosystem," in *International Conference on Financial Cryptography and Data Security*. Springer, 2018, pp. 101–112.

[5] K. Toyoda, P. T. Mathiopoulos, and T. Ohtsuki, "A novel methodology for hyip operators' bitcoin addresses identification," *IEEE Access*, vol. 7, pp. 74 835–74 848, 2019.

[6] K. Toyoda, T. Ohtsuki, and P. T. Mathiopoulos, "Multi-class bitcoin-enabled service identification based on transaction history summarization," in *iThings/ GreenCom/ CPSCom/ SmartData/ Blockchain/ CIT/ Cybermatics 2018*. IEEE, 2018, pp. 1153–1160.

[7] M. Vasek and T. Moore, "There's no free lunch, even using bitcoin: Tracking the popularity and profits of virtual currency scams," in *International conference on financial cryptography and data security*. Springer, 2015, pp. 44–61.

[8] M. Bartoletti, B. Pes, and S. Serusi, "Data mining for detecting bitcoin ponzi schemes," in *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*. IEEE, 2018, pp. 75–84.

[9] C. Brenig, R. Accorsi, and G. Müller, "Economic analysis of cryptocurrency backed money laundering." in *ECIS*, 2015.

[10] M. Möser, R. Böhme, and D. Breuker, "An inquiry into money laundering tools in the bitcoin ecosystem," in *2013 APWG eCrime Researchers Summit*. Ieee, 2013, pp. 1–14.

[11] M. Spagnuolo, F. Maggi, and S. Zanero, "Bitiodine: Extracting intelligence from the bitcoin network," in *International Conference on Financial Cryptography and Data Security*. Springer, 2014, pp. 457–468.

[12] K. Liao, Z. Zhao, A. Doupé, and G.-J. Ahn, "Behind closed doors: measurement and analysis of cryptolocker ransoms in bitcoin," in *2016 APWG eCrime*. IEEE, 2016, pp. 1–13.

[13] S. Bistarelli, M. Parroccini, and F. Santini, "Visualizing bitcoin flows of ransomware: Wannacry one week later." in *ITASEC*, 2018.

[14] K. Chapman, "Riviera beach commissioners vote to pay ransom to hacker who shut down city computers." https://bit.ly/2TTulE0, June 19th 2019.

[15] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia, "Dissecting ponzi schemes on ethereum: identification, analysis, and impact," *Future Generation Computer Systems*, vol. 102, pp. 259–277, 2020.

[16] T. Moore, J. Han, and R. Clayton, "The postmodern ponzi scheme: Empirical analysis of high-yield investment programs," in *Financial Cryptography and Data Security*, A. D. Keromytis, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 41–56.

[17] Jean-Luc, "Liste d'escroqueries liées à bitcoin et aux cryptomonnaies - bitcoin.fr," http://bit.ly/2Pi5YN7, 2020.

[18] WaybackMachine, "Wayback machine," https://web.archive.org/.

[19] A. Kharraz, W. Robertson, and E. Kirda, "Surveylance: Automatically detecting online survey scams," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 70–86.

[20] B. Srinivasan, A. Kountouras, N. Miramirkhani, M. Alam, N. Nikiforakis, M. Antonakakis, and M. Ahamad, "Exposing search and advertisement abuse tactics and infrastructure of technical support scammers," in *WWW'18*, 2018, pp. 319–328.

[21] E. Badawi, G.-V. Jourdan, G. Bochmann, I.-V. Onut, and J. Flood, "The "game hack" scam," in *ICWE 2019. Springer LNCS 11496*, 2019, pp. 280–295.

[22] E. Badawi, G.-V. Jourdan, G. Bochmann, and I.-V. Onut, "Automatic Detection and Analysis of the "Game Hack" Scam," *Journal of Web Engineering*, vol. 18, no. 8, 2020.

[23] w3schools, "Html meta name attribute," https://bit.ly/2I87xcw.

[24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[25] F. Reid and M. Harrigan, "An analysis of anonymity in the bitcoin system," in *Security and privacy in social networks*. Springer, 2013, pp. 197–223.

[26] "Bitcoin wiki," https://en.bitcoin.it/wiki/From_address.

[27] K. Toyoda, T. Ohtsuki, and P. T. Mathiopoulos, "Identification of high yielding investment programs in bitcoin via transactions pattern analysis," in *GLOBECOM 2017*. IEEE, 2017, pp. 1–6.

[28] K. Toyoda, T. Ohtsuki, and P. Mathiopoulos, "Time series analysis for bitcoin transactions: The case of pirate@ 40's hyip scheme," in *IEEE ICDMW'18*. IEEE, 2018, pp. 151–155.