



The Bitcoin Hunter: Detecting Bitcoin Traffic over Encrypted Channels

Fatemeh Rezaei¹(✉), Shahrzad Naseri¹, Ittay Eyal², and Amir Houmansadr¹

¹ University of Massachusetts Amherst, Amherst, USA
{frezaei, shnaseri, amir}@cs.umass.edu

² Technion, Haifa, Israel
ittay@technion.ac.il

Abstract. Bitcoin and similar blockchain-based currencies are significant to consumers and industry because of their applications in electronic commerce and other trust-based distributed systems. Therefore, it is of paramount importance to the consumers and industry to maintain reliable access to their Bitcoin assets. In this paper, we investigate the resilience of Bitcoin to blocking by the powerful network entities such as ISPs and governments. By characterizing Bitcoin's communication patterns, we design classifiers that can distinguish (and therefore block) Bitcoin traffic even if it is tunneled through an encrypted channel like Tor and even if Bitcoin traffic is being mixed with background traffic, e.g., due to browsing websites. We perform extensive experiments to demonstrate the reliability of our classifiers in identifying Bitcoin traffic even despite using obfuscation protocols like Tor Pluggable T transports. We conclude that standard obfuscation mechanisms are not enough to ensure blocking-resilient access to Bitcoin (and similar cryptocurrencies), therefore cryptocurrency operators should deploy tailored traffic obfuscation mechanisms.

Keywords: Bitcoin · Traffic analysis · Blocking-resistance

1 Introduction

Bitcoin and similar blockchain-based currencies [35] have seen rapid adoption by consumers and industry because of their many applications in electronic commerce and other trust-based distributed systems. Bitcoin supports \$1–\$4.2B worth of transactions per day, growing steadily. Bitcoin and similar virtual currencies offer significant advantages compared to traditional electronic currencies, which include open access to a global e-commerce infrastructure, lower transaction fees, cryptographically supported contracts [3] and services [32], and transnational operations.

Given this significant importance of electronic currencies, they need to be resistant to embargoes by governments. That is, people investing in cryptocurrencies (by running businesses that rely on such currencies) should be assured that their Internet providers or governments are not able to prevent them from

using their cryptocurrencies if they decide too. For the sake of argument, consider what happens if the Great Firewall of China decides to block all Bitcoin traffic overnight.

In this paper, we investigate the resilience of Bitcoin to blocking by powerful network entities, including ISPs and governments. Note that identifying standard (non-encrypted) Bitcoin traffic is trivial as Bitcoin messages use specific packet contents and formats. Therefore, a trivial countermeasure to prevent an ISP from identifying Bitcoin traffic is to tunnel Bitcoin over an encrypting tool, e.g., VPN, SSH, or Tor. However, previous studies [7, 46] show that encryption is *not enough* to conceal the nature of communications, as the characteristics of encrypted traffic may leak sensitive information. Such attacks are broadly known as *traffic analysis*.

In this paper, we investigate if and how Bitcoin’s traffic can be identified through traffic analysis despite being tunneled through an encrypted channel. First, we characterize Bitcoin’s traffic patterns such as rates, timings, and sizes. Comparing with other protocols, we show that Bitcoin has traffic patterns that are unique, because of the specific types of messages sent by Bitcoin peers. Leveraging such unique features of Bitcoin traffic, we design a toolset of classifiers in order to distinguish Bitcoin traffic over encrypted channels. We perform extensive evaluations of our classifiers by capturing Bitcoin traffic in the wild. Particularly, we use several months of Tor traffic tunneled through Tor [12] and three major Tor pluggable transports [40], namely, FTE [14], meek [31], and obfs4 [39], to evaluate our classifiers. Our experiments show that while such obfuscation mechanisms modify Bitcoin’s traffic by changing packet sizes timings, they are not able to hide the presence of Bitcoin traffic.

In summary we make the following main contributions:

1. We evaluate Bitcoin’s traffic and characterize its patterns such as its packet sizes and traffic shape. We compare Bitcoin’s traffic patterns to other popular protocols showing the uniqueness of Bitcoin traffic.
2. Based on our characterization of Bitcoin traffic, we design a range of classifiers whose goal is to identify Bitcoin traffic despite being tunneled through an encrypted channel (like Tor) and in the presence of background noise (e.g., open browser tabs).
3. Using several months of Bitcoin traffic and other protocols, we perform experiments to evaluate the performance of our classifiers when Bitcoin traffic is tunneled over Tor and three major Tor pluggable transports of FTE [14], meek [31] and obfs4 [39], and in the presence of background noise. Our classifiers are able to identify Bitcoin traffic in all cases with only 10 min of traffic with more than 99% true positive and negligible false positives.

2 Background on Bitcoin Traffic and Its Network Traffic

Bitcoin communications involve various protocol messages that are created by Bitcoin peers. We divide Bitcoin protocol messages into two classes: *synchronization messages*, which are used for propagating user addresses and transactions in the Bitcoin network, and *block-related messages*.

Synchronization Messages: These messages are aimed at keeping Bitcoin peers synchronized with the rest of the Bitcoin network.

addr: Each peer advertises the information and IP addresses of other peers via **addr** message in the network.

inventory(inv): Peers send **inv** to advertise their knowledge about the known objects, like transactions and blocks.

getdata: A peer sends **getdata** message in response to the **inv** to retrieve information about the content of an object, which can be a block or a transaction.

tx: This message describes a transaction in response to a **getdata** message.

Block-Related Messages: Such messages are used to exchange Bitcoin blocks among the peers. The current Bitcoin network is supporting two ways of propagating blocks, full block and compact block propagation.

Full Block Propagation: The sender node first validates the block completely, then it advertises the possession of block by an **inv** message. The receiving peer which does not have the block asks for it by sending an **getdata** message. Finally, the sender node sends the block via a **block** message. A **block** message consists of block version information, previous block hash, and the merkle root value related to this block.

Compact Block Propagation: In 2016, Bitcoin rolled out a new propagation mechanism (version 0.13.0). In this new approach, only a sketch of each block is sent to the peers instead of full blocks. The sketch include 80-byte block's header, the short transactions IDs used for matching already-available transactions and a selection of transactions which sending peer expect that a receiving peer may be missing. The following are the specific messages exchanged to run Bitcoin in the compact block mode.

sendcmpct: This message informs the receiving peer about the mode of communication the sending peer has chosen (low or high bandwidth).

cmpctblock: This message introduced in the compact block relaying and is presenting a sketch of block.

getblocktxn: This message is introduced in compact block relaying and is used to request for the transactions that are missed by sending a list of their indexes.

blocktxn: This message is introduced in compact-block relaying and is used to provide some of the transactions in a block, as requested.

3 Characterizing Bitcoin Traffic

We start by characterizing Bitcoin's traffic patterns.

3.1 Proportion and Distribution of Messages

Bitcoin peers generate various kinds of messages as introduced in the previous section. We show that the distribution and sizes of such messages are quite unique

to the Bitcoin protocol, making Bitcoin traffic easily distinguishable from other protocols.

Distribution of Packet Sizes. Figures 1a to 1h illustrate the packet size histogram of different types of Bitcoin messages in our collected Bitcoin traffic. As can be seen, each type of message has a distinguishing traffic pattern. Note that through our experiments, we find out that `tx` and `inv` are dominating with 43.6% and 27.2% of all packets, respectively. Therefore, the characteristics of these messages will shape the pattern of a Bitcoin peer’s traffic.

Histogram of Packet Sizes in Aggregate Traffic. Figures 2a and 2b show the histogram of packet sizes in the upstream and downstream directions, respectively, in compact block relaying. As mentioned before, `tx` and `inv` dominate the messages sent by a typical Bitcoin peer, therefore their sizes (shown in Figs. 1a and 1h) strongly shape the histogram of Bitcoin traffic, making it uniquely distinguishable from other protocols.

We also show the histogram of Bitcoin traffic in the full block relaying mode in Figs. 2c and 2d. These histograms have a larger spike close to the MTU, unlike the case of compact block relaying. These are because of the larger block sizes (around 1 MB) in the full block relaying.

Comparing to Other Protocols: Figures 2e to 2l show the histogram of other popular protocols, collected as described in Sect. 5. Note that we look at the traffic after going through an encryption tunnel, e.g., a VPN or SSH tunnel, so the histogram includes the (small) TCP ACK packets. As we can see, the packet size distribution of Bitcoin is uniquely different from these other protocols, since a Bitcoin connection is composed of unique messages with specific size distributions shown before. For instance, the large number of `inv` messages shapes the overall distribution of sizes in Bitcoin traffic.

Ratio of Downstream to Upstream. We also measured the ratio of downstream to upstream traffic volumes, which is shown in Fig. 3a. Unlike other protocols like HTTP (shown in Figs. 3b to 3f), Bitcoin traffic has a *symmetric* traffic volume in upstream and downstream. This is due to the fact that Bitcoin peers broadcast most of the bulky protocol messages they receive such as block and transaction announcements.

3.2 Shape of Traffic

Above we showed that the counts and sizes of packets in Bitcoin demonstrate a unique behavior. Here we show that, additionally, the shape of Bitcoin traffic is distinguishable from other protocols.

Full Block Relaying Mode: Figure 4a shows the traffic of a Bitcoin client operating in the full block relaying mode. As can be seen, the small protocol packets, mostly corresponding to `inv` and `tx` messages, appear uniformly over the time. On the other hand, the Bitcoin full blocks appear as large spikes of roughly 1 MB at specific points in time, i.e., once a new block is generated in the network.

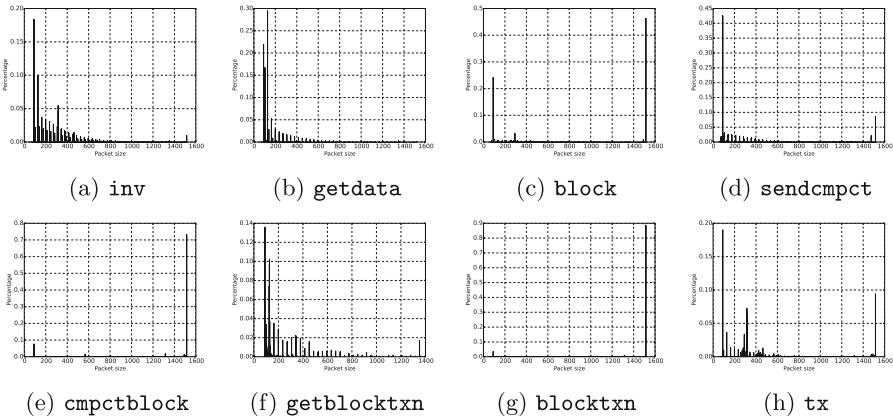


Fig. 1. Packet size distribution of Bitcoin communication messages in compact block relaying

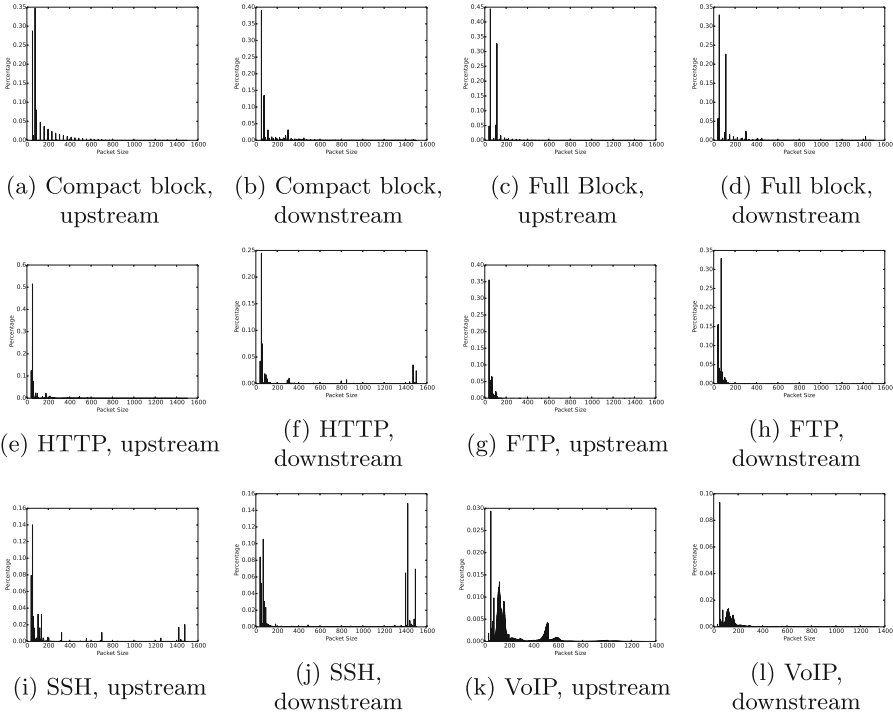


Fig. 2. Upstream and downstream packet size distribution of Bitcoin and several popular protocols

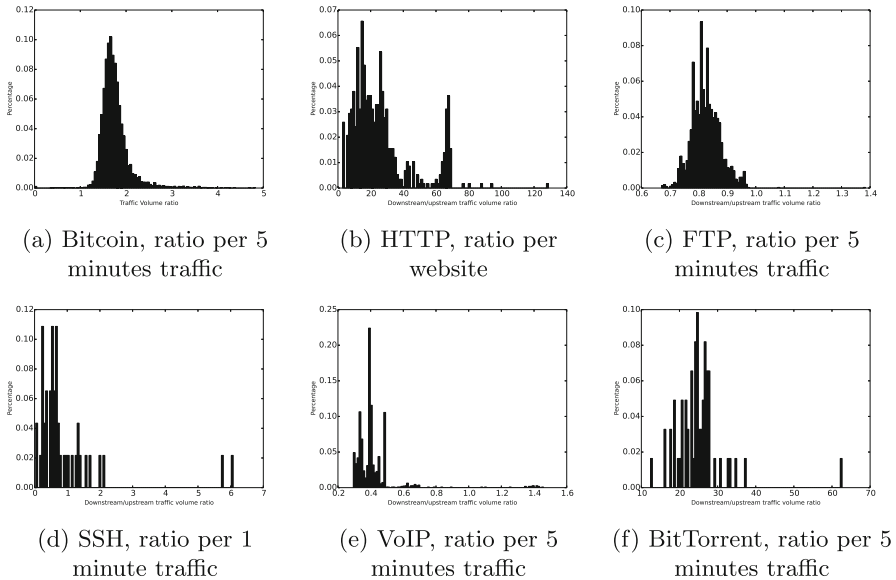


Fig. 3. Ratio of downstream to upstream traffic volume

Compact Block Relaying Mode: In the compact block relaying mode, it is harder to notice the block spikes since only a sketch of the blocks is transmitted. In this mode, transmitting a compact block in the network results in smaller spikes of 100 KB. Spikes of such small sizes may also occur when unverified transactions are transmitted, which will increase the detection’s false positive. Also, a Bitcoin client may operate in the high bandwidth mode, in which the receiver node asks its peers to send new blocks without asking for permissions first. This will lead to more than one peer sending the same block at the same time. This and the large volume of missed transactions result in having spikes with more than 100 KB in the traffic. Figure 4b illustrates when and how compact blocks appear on a peer’s traffic. As can be seen, compact blocks appear at smaller amplitudes than the actual block size, but the behavior is also nondeterministic, since it depends on whether the client has previously received some of the transactions in that block. This intuitively makes detection of compact blocks less reliable than full blocks, as shown later in our experiments.

We also measure the size of compact blocks by measuring the *length* field of `cmpctblock` messages, which is shown in Fig. 5a. As can be seen, most of the compact blocks are as small as 15 KB (in contrast to 1 MB in full blocks).

Finally, we measure the volume of transactions missing from an announced compact block (we do so based on the payload length of `blocktxn` messages). As described earlier, a Bitcoin client operating in the compact block mode will

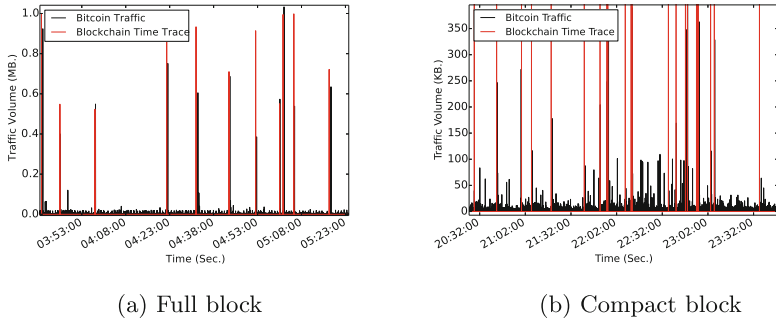


Fig. 4. Comparing time of each block receive with time of blocks in the block chain in a) Full block and b) Compact block modes

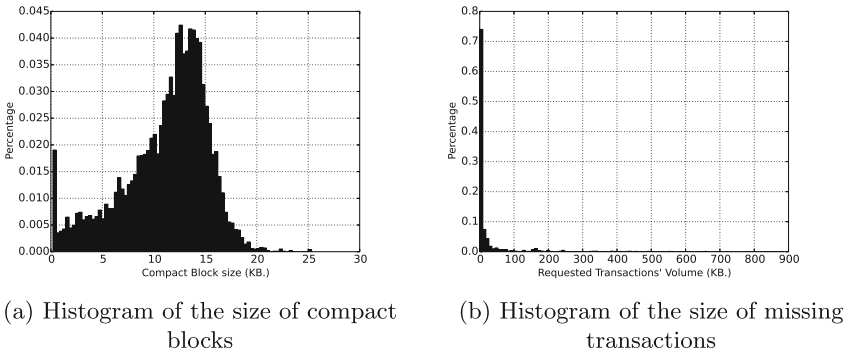


Fig. 5. Histogram of compact block mode components

download such missing transactions. This is shown in Fig. 5b. As we can see most of the transactions have volumes less than 100 KB.

Block Propagation Latencies. The propagation delay in the Bitcoin network is due to transmission delays and block verification by the receiving node at each hop. The transmission delay is the time to exchanging `inv` and `get data` messages, and sending the block via a `block` message. We measure block propagation delay by subtracting the receiving time of the block message and the time stamp in the header of the block message. Figures 6a and 6b show the histogram of propagation delay for 6000 blocks in compact block and full block relaying, respectively. As shown in the figures, we can model this empirical data using a Beta distribution [13].

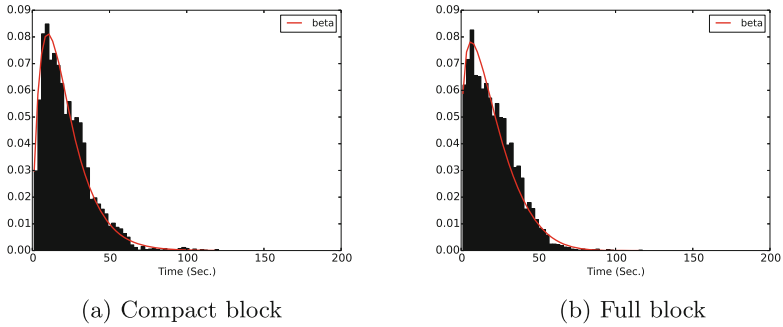


Fig. 6. Histogram of block propagation delays

4 Designing Bitcoin Classifiers

We leverage the features described above to build robust classifiers for Bitcoin traffic. We aim for our classifiers to work even in the presence of encryption and background traffic, e.g., when the machine running Bitcoin is used for web browsing and runs other applications, or when the Bitcoin traffic is tunneled over Tor.

4.1 Size-Based Classifier

As noted in Sect. 3.1, the histogram of Bitcoin packet sizes has a unique pattern. Based on this, we designed classifiers to distinguish Bitcoin traffic from other protocols.

Size Histogram Classifier (SizeHist). This classifier correlates the packet size histogram of a target user’s traffic with the packet size histogram of Bitcoin traffic captured by the adversary. By histogram of packet sizes we mean number of each packet size from 1 to MTU size. To do so, the classifier first divides the given traffic into upstream and downstream directions, then it calculates the histogram of packet sizes in each direction. The baseline (real-time Bitcoin traffic) is also divided in upstream and downstream directions. In the next step, the classifier calculates the cosine similarity between the histogram of the target traffic and several Bitcoin traces in both upstream and downstream direction. If both of the upstream and downstream averaged correlation values are above their thresholds, the target traffic is detected as Bitcoin traffic.

Tor-Specific Classifier (SizeTor): Tor [12] reformats traffic into constant-sized segments called *cells*. However, as the size of a cell is smaller than the MTU of IP packets, depending on the volume of traffic, multiple cells can merge into one single packet. This makes the number of single-cell packets and multiple-cell packets different for different protocols tunneled over Tor. Our **SizeTor**

classifier aims at detecting Bitcoin traffic tunneled over Tor based on the distribution of single-cell packets. As shown in Appendix 3, Bitcoin traffic consists of a large number of small-size packets (due to frequent `inv` messages). Tor will add padding to these small packets to form cells. Therefore, the ratio of single-size packets in Bitcoin over-Tor is larger than regular traffic over Tor, e.g., HTTP-over-Tor. Based on this, our classifier compares the ratio of single-cell packets to all packets; if the ratio is larger than a threshold, the connection will be flagged as a Bitcoin connection. Note that our `SizeTor` classifier can be adjusted for other protocols that similarly change the size of packets.

4.2 Shape-Based Classifier

The main intuition of our shape-based classifier is looking for changes in the traffic volume of a target user around the times of block announcements. Therefore, we assume that the classifier obtains the times and sizes of Bitcoin blocks, e.g., from the public `blockchain.info` repository, or even by running a local Bitcoin client.

For each confirmed Bitcoin block, the algorithm analyzes the volume of the target traffic two time windows with size ω_i around the block time, one before $(t_{block_i} - \omega_i, t_{block_i})$ and one after the block time $(t_{block_i}, t_{block_i} + \omega_i)$. The size of the window depends on the size of the block, the target client's bandwidth, etc., as evaluated later. For an actual Bitcoin traffic with no noise, the difference should be close to the size of the block. If the difference is within the bound (J), the algorithm considers that block to be *detected* in the traffic of the suspected user. The algorithm performs the same for a number of blocks and evaluates the ratio of such "detected" blocks. If the ratio is above a threshold η , the target user is declared to be a Bitcoin client.

Choosing the Threshold η . The threshold should be chosen based on the target user's specific network conditions such as background traffic, network noise, and bandwidth.

To do so, the classifier generates N (e.g., $N = 100$) synthetic block series, which we call *ground falses*. The classifier then correlates the target traffic with each of the N ground false instances using the correlation function. Finally, the threshold η is chosen to be larger than the largest correlation value.

Choosing Other Parameters. The window shape classifier also needs to choose the values of the parameters ω and J . Parameter ω needs to be big enough to contain the most of the traffic of a block during block propagation. Moreover, Parameter J is chosen to take into account that some of the block propagation traffic might not be downloaded in that time window. This parameters needs to be selected based on user's bandwidth and the volume of background traffic, and therefore it needs to be chosen for each client specifically.

4.3 Neural Network-Based Classifier (NN-Based)

As done in similar applications [36], we use a neural network-based classifier to detect Bitcoin in presence of a more complex background noise, e.g., brows-

ing more than one website simultaneously. In following, we explain the feature selection phase, and then describe the design of our neural network.

Feature Selection. To create each sample data, we divide time into intervals and use the volume of traffic in each interval as our features, which is presented in equation $v_{lm} = v_1, v_2, \dots, v_n$. Note that v_I is the volume of traffic in interval I. We choose 10 min as the *sample size*, which is the smallest length to have at least one peak of traffic. Furthermore, to choose the interval length, we try different values of 1, 5, 10 and 20 s. From our experiments, we find out that the interval length of 10 s results in the best performance. Therefore, we choose 10 s as the interval length (l). Since the length of each sample is 10 min (600 s), using equation $n = \text{sample size}/l$, we get an array of length 60 as our feature.

Designing the Model. For our neural network model, we use a combination of convolutional and fully-connected network that consists of an input layer, an output layer, and three hidden layers: one convolutional layer, and two dense layers. The input layer has $n = 60$ number of neurons, which is the size of each sample data, the hidden layers have $n_1 = 64$, $n_2 = 32$ and $n_3 = 16$ number of neurons, respectively, and the output layer has 1 neuron which represents if the sample data contains Bitcoin traffic or not. We use Relu as the activation function of the hidden layers and sigmoid [19] for the output layer. Also, we use binary cross-entropy as our loss function, and Adam optimization [28].

4.4 Combined Classifier

In this classifier, we combine all the attributes used in the above classifiers. More specifically, we are using the size histogram of the packets, downstream to upstream traffic volume ratio, and volume over time. Our final feature set has a length of 1576 (1 for downstream to upstream, 60 for volume over time, and 1515 for size histogram). Note that the feature that we use in sizeTor classifier is a subset of size histogram, and volume over time is capturing the shape of traffic that we utilized in the shape-based classifier. The model that we use has three convolutional layers and one dense layer with sizes: 1024, 128, 64, and 32 for the dense layer. We use the same activation functions and loss function as the NN-based classifier.

5 Experimental Setup

We use Bitcoin Core software¹ to run full node Bitcoin clients on 5 virtual machines on a campus network. Each virtual machine is connected to the Internet with high bandwidth. Before starting the experiments, we leave our Bitcoin clients for a few days to make sure they have downloaded an up-to-date blockchain ledger. The Bitcoin client is passively receiving blocks and participating in the Bitcoin network, but it is not doing any transactions. We capture Bitcoin traffic under different scenarios on a Linux 16.0.4 virtual machine.

¹ <https://bitcoin.org/en/bitcoin-core/>.

5.1 Datasets

Collecting Bitcoin Traffic: We use Bitcoin version 0.12.0 to capture Bitcoin traffic in the full block relaying mode and Bitcoin version 0.14.0 to capture traffic in the compact block relaying mode. We capture Bitcoin traffic for each version for a period of a month.

Bitcoin Tunneled Through Tor: We captured Bitcoin traffic behind Tor [12] for both compact and full block modes. We also captured Bitcoin traffic in the compact block mode behind Tor and popular Tor pluggable transports of obfs4 [47], FTE [14], and Meek-amazon [31].

Bitcoin with Background Traffic: We captured Bitcoin traffic in presence of HTTP background traffic by browsing the top 500 Alexa websites using the Selenium² tool while running Bitcoin software. We also collected Bitcoin traffic with HTTP background for the same set of websites behind Tor and its three pluggable transports using Selenium.

CAIDA Background Traffic: We use CAIDA’s 2018 anonymized traces³ as a dataset for additional background traffic.

HTTP Traffic: We collect top 500 Alexa websites using Selenium tool. Also, we capture these websites over Tor, and three pluggable transports. Moreover, we use the dataset by [36] which has collected the top 50,000 Alexa websites over Tor (Table 1).

Table 1. Traffic class breakdown for CAIDA dataset

Traffic class	Port numbers	Number of connections	% of total
http, https	80, 8080, 443	745262	0.318
dns	53	1073758	0.457
smtp	25	2646	0.001
telnet	23	6958	0.003
ssh/scp	22	4928	0.002
other	–	511700	0.219
all	–	2345252	1.0

5.2 Metrics

We use following metrics to measure the performance of our classifiers:

- **True positive:** True positive shows the proportion of the data which contain Bitcoin, and our model correctly identified as Bitcoin traffic.

² <http://www.seleniumhq.org>.

³ <https://www.caida.org/data/monitors/passive-equinix-nyc.xml>.

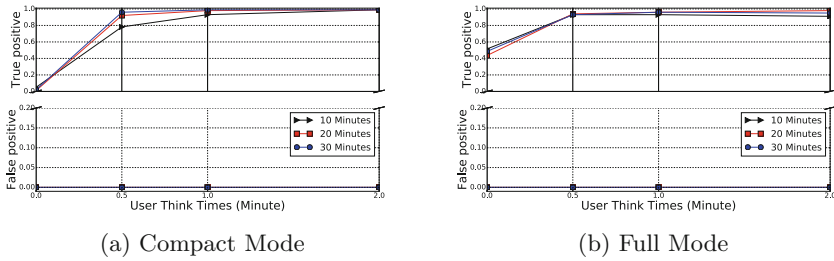


Fig. 7. Result of `sizeHist` classifier on noisy Bitcoin traffic

- **False positive:** False positive shows the proportion of the data which did not contain Bitcoin traffic, and our model incorrectly classified as Bitcoin.
- **Accuracy:** Accuracy shows the proportion of the data which was correctly classified.

5.3 Modeling Normal Users

In this section, we describe four different types of users' profile that we use to evaluate the performance of our Bitcoin classifiers against.

- **Simple User:** A simple user is a Bitcoin client with no background noise and traffic. This type of user does not generate any network traffic except the Bitcoin client application traffic. Therefore, all traffic of the simple user is Bitcoin traffic.
- **Simple Noisy User:** A simple noise user is a Bitcoin client who browses only **one** webpage. To control the background traffic, we introduce a parameter named think time, T , representing the amount of time that the user spends on a particular website.
- **Complex Web (Complicated/Sophisticated) User:** A complex user is a Bitcoin client who browses **multiple** websites simultaneously.
- **Complex CAIDA User:** A complex CAIDA user is a Bitcoin client who is running 1–5 number of CAIDA applications, which is introduced in Table 2 simultaneously in the background.

6 Results

In this section, we implement our classifiers to evaluate their performance on user profiles described in part Sect. 5.3 writing more than a thousands lines of code in Python. First, for each classifier, we declare the user profile(s) that we use for evaluation of its performance and the false data that we use for computing its false positive. Second, we describe the result of each classifier and give a summary and comparison of them at the end of this section.

6.1 User Profiles and False Data

For each classifier, we use a specific user profile and depending on that we choose the false data. As we explained above, false data is the base traffic that we use to compute false positive. Note that, we use same length of traffic for Bitcoin and the false data. In other words, it is the traffic that we compare our Bitcoin traffic with. For example, when we have 10 min of HTTP traffic, we are continuously browsing different websites for 10 min. We make these samples by concatenating the browsing of different websites. In the following, we describe these pairs for each classifier(s).

- For shape-based classifier, we use the simple user profile. Furthermore, we use HTTP which is the typical user behavior as the false data. This experiment evaluates if Bitcoin traffic can be differentiated from browsing an HTTP website.
- For the rest of the binary classifiers in this section, we use simple noisy user profile and attempt to detect the presence of Bitcoin. Note that, similar to the window-based classifier, we use HTTP for the false data. This experiment attempts to evaluate if browsing an HTTP website is enough to hide the Bitcoin traffic.
- For the neural network-based and combined classifiers, we use the complex web and complex CAIDA user for training and testing.

Note that, for these two classifiers, we evaluate our model using 10,000 number of test data, and report the false positive, true positive and accuracy. For rest of the classifiers, we use 500 number of test data for evaluation. Also, the data is balanced, which means we have the same number of data for each category.

6.2 Size-Based Classifiers

SizeHist Classifier. We implement the `sizeHist` classifier on Bitcoin traffic in compact and full block modes using the noisy user model described in Sect. 5.3. Figure 7 shows the performance of this classifier. We control the noise using think time (T). Increasing T decreases the noise and enhances the classifier’s performance. Figure 7 shows that we can reach more than 90% true positive and 0% false positive for both modes when we have 10 min of traffic and set T to 2 min. It is worth stating that we could reach similar results when we set T to 0.5 min and have 20 min of traffic.

6.3 Shape-Based Classifier

To implement this classifier, we set J and ω introduced in Sect. 4.2 to 100 kilobytes and 20 s, respectively. To set η , we compute block detection rate for Bitcoin using ground false shown in Fig. 8. We need to set η to be larger than all the detection rate values for Bitcoin using ground false. Note that each point for Bitcoin using ground false in the figure is the average for 25 different ground falses.

Using this figure, we set η to 0.4. Moreover, Fig. 8 shows the block detection rate for HTTP using ground truth and block detection rate for Bitcoin using ground truth too. Using the η , we can detect all Bitcoin traffic through (August 28–October 5) as Bitcoin. Also, we did not classify any of the HTTP traffic as Bitcoin, which results in 0% false positive. Furthermore, the performance of shape-based classifier quickly diminishes in the presence of a small HTTP background noise. Also, we fail to detect Bitcoin traffic in compact mode because of small block sizes, which makes it impossible to distinguish them.

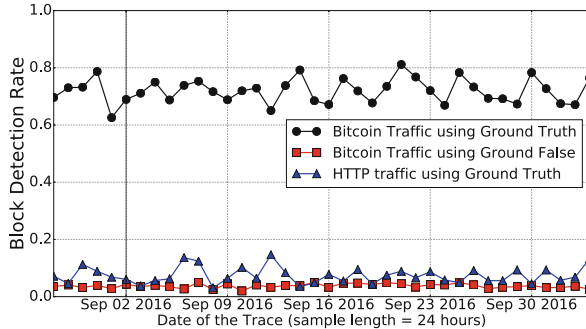


Fig. 8. Block detection rate using shape-based classifier

6.4 Neural Network-Based Classifier

We implement the NN-based classifier using Keras [10] with Tensorflow [1] backend. We use complex web and complex CAIDA user to evaluate its performance. Table 2 shows the result of NN-based classifier for different sizes of training data. As the table indicates, the accuracy of the classifier improves from 62% to 96% when we increase the size of training data from 1000 to 40,000. Note that, true and false positive improves from 44% to 92% and 20% to 2% respectively when we increase the size of training data.

Table 2. Result of neural network classifier.

Training size	False positive (%)	True positive (%)	Accuracy (%)
1000	20	44	62
5000	11	80	85
10,000	6	84	89
40,000	2	92	95

6.5 Combined Classifier

To extend the neural network-based classifier, we defined combined classifier, which uses all of the attributes used in previous classifiers. Using this classifier, we reach 99.84% accuracy with false positive of 0 and true positive of 99.74% having 40,000 of training data and sample size of 10 min. This result is very promising and shows that having enough data and using the attributes that distinguishes Bitcoin from other traffic, we are able to train a neural network model that gives us 0% false positive and more than 99% accuracy.

6.6 Summary and Comparison of the Results

The SizeHist and shape-based classifiers work only when there is a very small background noise or no noise (think time of 2 min). Therefore, they are not useful when Bitcoin traffic has a large amount of background noise. To distinguish Bitcoin traffic in the presence of larger noises, we employ NN-based and combined classifiers. The benefit of these classifiers is that they do not have the training phase required in NN-based techniques. On the other hand, the NN-based classifier and combined classifiers result in a better performance in the presence of higher background noise (complex web and complex CAIDA). The combined classifier outperforms the NN-based by using more features during training. This classifier detects Bitcoin traffic using the complex web and complex CAIDA user explained in Sect. 5.3 with 0% false positive and much higher accuracy (99.84%).

7 Countermeasures

A possible countermeasure against Bitcoin detection is to tunnel Bitcoin traffic over an anonymity system like Tor. We tunnel Bitcoin traffic over Tor to disguise its patterns. We evaluate the invisibility that Tor provides by designing a new classifier (SizeTor) which is tailored for Tor. Also, we use our strongest classifiers (NN-based and combined) to evaluate if Tor succeeds in hiding Bitcoin traffic.

7.1 Bitcoin over Tor

Tor sends traffic in *cells*. If the packets sizes is below the cell size, it will add padding to the packets to reach the fixed cell sizes. This modifies the traffic patterns of Bitcoin traffic, e.g., its packet sizes, therefore it may increase resistance to traffic classification. In the following we evaluate our classifiers against Tor and three state-of-the-art Tor pluggable transports [40] namely *obfs4*, *FTE*, and *Meek* explained in the following.

Pluggable Transports: We evaluated against the following major transports:

- **Obfs4:** *obfs4* is a widely used Tor pluggable transport, which is based on *ScrambleSuit* [44]. It differs with *ScrambleSuit* in public key obfuscation and its protocol for one-way authentication, which makes it faster.




Fig. 9. Histogram of one cell packet ratio for HTTP traffic and Bitcoin traffic

- **FTE:** The FTE transport re-encrypts Tor packets in order to match the regular expressions of a benign protocol like HTTP.
- **meek:** Meek uses *domain fronting* [18] to tunnel traffic through public CDN or cloud platforms.

Bitcoin traffic has a larger ratio of one-cell packets compared to Tor. This is due to a large number of small Bitcoin messages (e.g., `inv`) that are each put into a single-cell packets. Figures 9a and 9b show the histogram of one cell size packets ratio in the upstream and downstream directions.

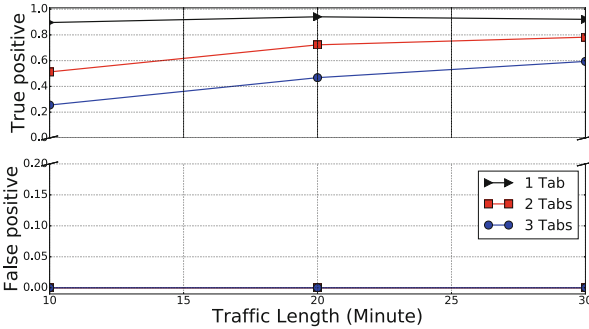
7.2 Evaluating Bitcoin Over Tor

SizeTor Classifier. We implement **SizeTor** classifier on noisy Bitcoin on compact mode for Tor and its three pluggable transports. As it is displayed in Fig. 10, we could detect Bitcoin traffic with high accuracy for the complex user model when the background noise is one open tab. As Fig. 10 indicates, having a traffic size of 10 min is enough to detect Bitcoin traffic with around 90% true positive and 0% of false positive. Moreover, the Figure states that the result of classifier quickly diminishes when we increase the background noise from one to 2 or 3 open tabs. Note that this figure shows the average result of this classifier on Tor and three pluggable transports.

Neural Network-Based Classifier. Table 3 presents the result of NN-based classifier for the complex web user over Tor for 10,000 numbers of test data. As the table suggests, performance of our classifier improves when we increase the size of training data. More specifically, our false positives enhances from 11% to 4% when we increase the training data from 1000 to 40,000. Moreover, when using 40,000 training data, we reach 99% and 4% true and false positive, respectively (97% accuracy). Note that the reason we get better results from this classifier on Tor dataset is that this dataset only contains the complex web user since we did not have CAIDA application over Tor to use for our evaluation.

Table 3. Result of neural network classifier on Tor dataset.

Training size	False positive (%)	True positive (%)	Accuracy (%)
1000	11	98	93
5000	6	98	96
10,000	3	98	97
40,000	4	99	97

**Fig. 10.** Detecting Bitcoin compact mode traffic behind Tor (Tor and its three pluggable transports: meek, obfs, fte) using SizeTor

Combined Classifier Over Tor: We apply the combined classifier on our Tor dataset, which consists of Bitcoin traffic with up to 5 number of open tabs on the background. Our experiments show that having 40,000 of training data is enough to achieve more than 99.72% accuracy with 0.04% false positive and 99.5% true positive. Our combined classifier outperforms the previous classifiers (SizeTor and NN-based) on the Tor dataset as well. This is because we are using a more complex model and a more significant number of features.

8 Related Work

In this section, we overview previous work on classifying different protocols and discuss previous attacks on Bitcoin cryptocurrency.

8.1 Protocol Classification

There is extensive work in literature attempting to classify different applications or protocols in the network. Previously, researchers were focused on classification according to the port numbers [23, 24, 30, 42] and payload [9, 20, 34, 42]. Since many applications use uncertain port numbers [23], or some encrypt their payloads, researchers adopted new methods for protocol classification. Recent techniques use statistics such as packet sizes and timings for classification [5, 8, 11, 15, 16, 25, 48].

In some of the studies, researchers apply machine learning techniques on these statistics to classify different applications. For example, in [15], authors take a semi-supervised approach to classify a variety of applications such as FTP, HTTP, P2P. In [8, 48], authors use SVM classifier to distinguish different applications such as WWW, Mail, and FTP. Also, in [48], authors use SVM to classify a broad application category such as mail, buck traffic, service. Moreover, in [27, 38], authors survey the papers on Internet traffic classification using machine learning techniques.

8.2 Attacks on Bitcoin Cryptocurrency

In this section, we discuss the previous attacks on the Bitcoin network. In [4], the authors discuss routing attacks, their impact, and possible countermeasures. In [21], authors design an eclipse attack in which the attacker isolate a victim from its peers by monopolizing its all incoming and outgoing connections.

In [2, 41], authors use Bitcoin transaction patterns to link users (or link transactions) using some side information. In [6], the authors propose a technique to link the public key of a user to her address or link her transactions. In [26], they analyze the security of using Bitcoin for fast payments and show that the current Bitcoin system is not secure unless they integrate Bitcoin network with some detection mechanism. Also, they study double-spending attacks on fast payments and implement a method to prevent it. Attacks on the blockchain system as a whole [17, 29, 33, 37], and against second-tier payment networks [22, 43, 45] are not directly affected by single participant detection and hacking.

9 Conclusions

The reliable access to Bitcoin and similar cryptocurrencies is of crucial importance due to their consumers and the related industry. In this paper, we investigated the resilience of Bitcoin to blocking by a powerful network entity such as an ISP or a government. By characterizing Bitcoin's communication patterns, we designed various classifiers that could distinguish (and therefore block) Bitcoin traffic even if it is tunneled over an encrypted channel like Tor, and even when it is mixed with background traffic. Through extensive experiments on network traffic, we demonstrated that our classifiers could reliably identify Bitcoin traffic despite using obfuscation protocols like Tor pluggable transports that modify traffic patterns. In order to disguise such patterns, an obfuscating protocol needs to apply significant cover traffic or employ large perturbations, which is undesirable for typical clients.

Acknowledgement. The work was supported by the NSF CAREER grant CNS-1553301 and BSF.

References

1. Abadi, M., et al.: TensorFlow: large-scale machine learning on heterogeneous systems (2015). <http://www.tensorflow.org>

2. Androulaki, E., Karame, G.O., Roeschlin, M., Scherer, T., Capkun, S.: Evaluating user privacy in bitcoin. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 34–51. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39884-1_4
3. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multi-party computations on bitcoin. In: S & P (2014)
4. Apostolaki, M., Zohar, A., Vanbever, L.: Hijacking bitcoin: routing attacks on cryptocurrencies. In: 2017 IEEE Symposium on Security and Privacy
5. Bar-Yanai, R., Langberg, M., Peleg, D., Roditty, L.: Realtime classification for encrypted traffic. In: Festa, P. (ed.) SEA 2010. LNCS, vol. 6049, pp. 373–385. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13193-6_32
6. Biryukov, A., Khovratovich, D., Pustogarov, I.: Deanonymisation of clients in bitcoin P2P network. In: CCS (2014)
7. Cai, X., Zhang, X., Joshi, B., Johnson, R.: Touching from a distance: website fingerprinting attacks and defenses. In: CCS (2012)
8. Cao, J., Fang, Z., Qu, G., Sun, H., Zhang, D.: An accurate traffic classification model based on support vector machines. *Int. J. Network Manage.* **12**, 301 (2017)
9. Choi, T., et al.: Content-aware internet application traffic measurement and analysis. In: Managing Next Generation Convergence Networks and Services, IEEE/IFIP Network Operations and Management Symposium, NOMS (2004)
10. Chollet, F.: keras. <https://github.com/fchollet/keras> (2015)
11. Crotti, M., Dusi, M., Gringoli, F., Salgarelli, L.: Traffic classification through simple statistical fingerprinting. *Comput. Commun. Rev.* **37**(1), 5–16 (2007)
12. Dingledine, R., Mathewson, N., Syverson, P.: Tor: the second-generation onion router. In: USENIX Security (2004)
13. Durrett, R.: Probability: Theory and Examples. Cambridge University Press, Cambridge (2010)
14. Dyer, K., Coull, S., Ristenpart, T., Shrimpton, T.: Protocol misidentification made easy with format-transforming encryption. In: CCS (2013)
15. Erman, J., Mahanti, A., Arlitt, M.F., Cohen, I., Williamson, C.L.: Offline/realtime traffic classification using semi-supervised learning. *Perform. Eval.* **64**, 1194–1213 (2007)
16. Erman, J., Mahanti, A., Arlitt, M.F., Williamson, C.L.: Identifying and discriminating between web and peer-to-peer traffic in the network core. In: WWW (2007)
17. Eyal, I., Sirer, E.G.: Majority is not enough: bitcoin mining is vulnerable. *ACM Commun.* **61**, 95–102 (2018)
18. Fifield, D., Lan, C., Hynes, R., Wegmann, P., Paxson, V.: Blocking-resistant communication through domain fronting. In: PETS (2015)
19. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge (2016)
20. Haffner, P., Sen, S., Spatscheck, O., Wang, D.: ACAS: automated construction of application signatures. In: MineNet (2005)
21. Heilman, E., Kendler, A., Zohar, A., Goldberg, S.: Eclipse attacks on bitcoin's peer-to-peer network. In: USENIX (2015)
22. Jona Harris, A.Z.: Flood & loot: a systemic attack on the lightning network (2020)
23. Karagiannis, T., Broido, A., Brownlee, N., Claffy, K.C., Faloutsos, M.: Is P2P dying or just hiding? [P2P traffic measurement]. In: GLOBECOM (2004)
24. Karagiannis, T., Broido, A., Faloutsos, M., Claffy, K.C.: Transport layer identification of P2P traffic. In: IMC (2004)
25. Karagiannis, T., Papagiannaki, K., Faloutsos, M.: BLINC: multilevel traffic classification in the dark. In: Proceedings of the 2005 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (2005)

26. Karame, G., Androulaki, E., Capkun, S.: Double-spending fast payments in bitcoin. In: CCS (2012)
27. Kim, H., Claffy, K.C., Fomenkov, M., Barman, D., Faloutsos, M., Lee, K.: Internet traffic classification demystified: myths, caveats, and the best practices. In: CoNEXT (2008)
28. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. CoRR (2014)
29. Kwon, Y., Kim, D., Son, Y., Vasserman, E.Y., Kim, Y.: Be selfish and avoid dilemmas: Fork after withholding (FAW) attacks on bitcoin. CoRR (2017)
30. Madhukar, A., Williamson, C.L.: A longitudinal study of P2P traffic classification. In: MASCOTS (2006)
31. Meek Pluggable Transport. <https://trac.torproject.org/projects/tor/wiki/doc/meek>
32. Miller, A., Juels, A., Shi, E., Parno, B., Katz, J.: Permacoin: repurposing bitcoin work for data preservation. In: S&P (2014)
33. Mirkin, M., Ji, Y., Pang, J., Klages-Mundt, A., Eyal, I., Jules, A.: BDoS: blockchain denial of service. arXiv preprint [arXiv:1912.07497](https://arxiv.org/abs/1912.07497) (2019)
34. Moore, A.W., Papagiannaki, K.: Toward the accurate identification of network applications. In: PAM (2005)
35. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system (2008)
36. Nasr, M., Bahramali, A., Houmansadr, A.: Deepcorr: strong flow correlation attacks on tor using deep learning. In: CCS (2018)
37. Negy, K.A., Rizun, P.R., Siner, E.G.: Selfish mining re-examined. In: Bonneau, J., Heninger, N. (eds.) FC 2020. LNCS, vol. 12059, pp. 61–78. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-51280-4_5
38. Nguyen, T.T., Armitage, G.: A survey of techniques for internet traffic classification using machine learning. IEEE Commun. Surv. Tutorials **10**(4), 56–76 (2008)
39. A Simple Obfuscating Proxy. <https://www.torproject.org/projects/obfsproxy.html.en>
40. Tor: Pluggable Transports. <https://www.torproject.org/docs/pluggable-transports.html.en>
41. Reid, F., Harrigan, M.: An analysis of anonymity in the bitcoin system. In: SocialCom/PASSAT (2011)
42. Sen, S., Spatscheck, O., Wang, D.: Accurate, scalable in-network identification of p2p traffic using application signatures. In: WWW (2004)
43. Tsabary, I., Yechieli, M., Eyal, I.: MAD-HTLC: because HTLC is crazy-cheap to attack. arXiv preprint [arXiv:2006.12031](https://arxiv.org/abs/2006.12031) (2020)
44. Winter, P., Pulls, T., Fuss, J.: Scramblesuit: a polymorphic network protocol to circumvent censorship. In: Proceedings of the 12th ACM Workshop on Privacy in the Electronic Society. ACM (2013)
45. Winzer, F., Herd, B., Faust, S.: Temporary censorship attacks in the presence of rational miners. IACR Cryptol. ePrint Arch. (2019)
46. Wright, C., Ballard, L., Monrose, F., Masson, G.: Language identification of encrypted VoIP traffic: Alejandra y Roberto or Alice and Bob? In: USENIX Security (2007)
47. Yawning. Obfsproxy4 (2015). <https://github.com/Yawning/obfs4/blob/master/doc/obfs4-spec.txt>
48. Yuan, R., Li, Z., Guan, X., Xu, L.: An svm-based machine learning method for accurate internet traffic classification. Information Systems Frontiers (2010)