2018-06-12

Cryptocurrency performance analysis of Burstcoin mining

Tobias Larsson **Rasmus Thorsén**

University Diploma Thesis
Network Technology and IT-security, 120 HE credits

University West, Department of Engineering Science

Cryptocurrency performance analysis of Burstcoin mining

Abstract

Burstcoin is a cryptocurrency which differ from the well-known Bitcoin, among several integral component which encompass any cryptocurrency, the mining process is one which differ greatly between the two. This difference is basically that the Bitcoin mining process require constant computational calculation, while Burstcoin mining requires a file known as plot file to be read at intervals. As with many altcoins, Burstcoin have very little analysis surrounding it and as such the efficiency of the mining process in how it is usually configured can be questioned. Therefore, the thesis will focus on factors which might affect performance regarding how efficiently the plot file can be read. To begin with, knowledge about which timespans are interesting for Burstcoin mining will be investigated, more precisely the block forge times for the Burstcoin blockchain will be determined to ascertain if any specific timeframes are important regarding how fast a plot file should be processed. Following, three tests will be made focusing on the performance aspect. The first test focus on whether the file system cluster size matters in how fast a plot file is processed and determines if the processor is bottlenecking the processing speed of the plot file. Secondly, it is investigated if the size of the hard drive where a plot file is located play any role in how fast it is processed. The third test compares how fast a plot file is being read when the hard drive which houses the file operates at the SATA 2 standard compared to the SATA 3 one.

Date:	June 12, 2018
Author:	Tobias Larsson, Rasmus Thorsén
Examiner:	Thomas Lundqvist
Advisor:	Andreas de Blanche
Programme:	Network Technology and IT-security, 120 HE credits
Main field of study:	Computer Engineering
Education level:	First cycle
Course code:	EXM301, 7.5 HE credits
Keywords:	*Cryptocurrency, Proof of Capacity, Optimization, Burstcoin, Performance*
Publisher:	University West, Department of Engineering Science SE-461 86 Trollhättan, Sweden Phone: + 46 520 22 30 00, Fax: + 46 520 22 32 99, www.hv.se

Foreword

We would like to thank both Andreas de Blanche and Thomas Lundqvist as they both made us aware of Burstcoin.

Special thanks go out to Andreas de Blanche serving the role of Advisor, without the guidance of whom we would have had a hard time of finishing our work.

Both Tobias and Rasmus were equally involved in every aspect of the work.

Thanks to the developers of the game Ascension, which provided many well needed breaks.

Table of Content

Nomenclature

Block forging – To generate a block for the Burstcoin blockchain.

Mining – Performing work for the blockchain in the way of computational calculation and supplying the resulting values to the blockchain network.

Nonce – Data groups resulting from the Shabal algorithm.

Scoop – Data stored within nonces, 4096 scoops make up every nonce.

Deadlines – Resulting values from processing of scoops during the mining process.

Plot file – File used for the mining process which contain precomputed nonces.

1 Introduction

Bitcoin is a cryptocurrency; both these terms are becoming more widely known and understood. But to clarify what a cryptocurrency is, it can roughly be defined as three separate things [1]. The first being a blockchain, which is a form of ledger over held assets. Information about said assets are stored in blocks of data, which are continuously generated. The second is a protocol which is used to manage the transactions made over the blockchain through the generation of new blocks. The third is the digital currency which is what holds the value. Since first introduced, one among many effects Bitcoin have had is the origin of altcoins. An altcoin is a derivative of Bitcoin, either based on the same code or implemented using the concept of a blockchain.

One altcoin is Burstcoin which in comparison to Bitcoin alters how the work necessary to make transactions over the blockchain network is performed. This work, also known as mining, is performed through continuous calculations using for example ASICs for cryptocurrencies such as Bitcoin. In contrast, this work is done using hard drives for Burstcoin, also worthy of mention is that Burstcoin is the only cryptocurrency up to date which uses this method.

Since its release in 2014 not much robust analysis of how the Burstcoin protocol functions have been made. There are many aspects involved in the protocol itself, one such is the mining process where a plot file is being read. Lacking deeper insight regarding this process leads to speculation. Which in turn results in statements such as "I read somewhere it should be 30 seconds or below", or "64 kB is optimal" [2] [3]. In turn, such information is what lie as a basis for many miners when setting up their mining operation. This results in a high probability that the processing of the plot file(s) will be suboptimal in that the work takes longer than it needs to. To say the least much is to be desired regarding concrete information surrounding parameters which affect mining performance. Therefore, this thesis will focus on the performance of the mining process and to determine the impact factors selected in this work have.

To begin with, the distribution of block forge times up until 2018-02-07 will be investigated. This establishes how big the difference in full participation of blocks look if the processing is complete within for example 30 or 60 seconds and helps in understanding the impact improvement of mining performance have.

Once time distribution has been established, factors which may or may not affect performance of the mining process itself will be investigated. Where specifically the file system cluster size, size of the hard drive used to store the plot file, and SATA standards will be analysed using two hard drives of sizes 500 GB and 3000 GB.

To test the cluster size, a set of tests will be performed comparing the mining processing speed using three different cluster sizes for both hard drives. The three file system cluster sizes chosen are 4, 64 and 256 Kb. 4 Kb is the default file system cluster size in Windows, 64 Kb was tested because it is a number recommended by other miners and 256 Kb to see if there would be any linear scaling in improvement of processing speed if there exist any based on the file system cluster size.

Determining if hard drive size matters will be investigated through a test where to begin with a plot file is read on the smaller hard drive, the file is subsequently moved to the larger hard drive and subjected to the same test. Following the read times are compared.

Similarly, comparison of how fast the plot file is read will be made when the hard drive containing a plot file is connected to a SATA 2 and 3 port respectively.

2 The blockchain consensus algorithm

A blockchain can provide many services, one of which is to supply an alternative payment platform. It can provide this since it can fill the role of a ledger, or in other words a registry over held value and transactions [4]. The reason for this is because when combined with the protocol and currency, a blockchain solves the distributed trust problem. The distributed trust problem lies in where several decentralised peers with no knowledge about each other need to trust the information provided the other peers. This can be achieved through a consensus algorithm which is a set of rules that need to be followed to participate.

One part of this ruleset is the reliance upon a public and private key pair for every participant of the network. To simplify, the private key is used to sign off on payments, and the public one is the address used when sending currency to someone, the public keys are also registered in the blockchain with the currency they hold. This prevents any participant from spending more currency than is available to them. Attempted transactions which intend to overdraw will just be discarded since the value registered for that key in the blockchain doesn't match the amount specified in the transaction. On the other hand, if transactions are made which conform to the information stored in the blockchain, the miners of the blockchain network accepts these and stores the information within the blockchain as a new block.

The Bitcoin blockchain was the first one to implement this type of consensus protocol, which is known as Proof of Work (PoW) [5]. PoW introduced the concept of miners, and as the name implies requires continuous work to be done. The work is performed by the miners through continuous computational calculation, with the goal of providing a good enough solution to a mathematical problem. This is achieved through a hashing function with the goal of producing a value that satisfies the criterion set by the blockchain algorithm. Through this mining process, the miners prove they are performing work for the blockchain. This result in the generation of new blocks at predefined variable intervals and is where the work provided by the miners come into play. The miner which provide the best solution to the network generates the newest block and gets paid for this. Once the new block is generated the new transactions are verified and registered in the blockchain. While new blocks serve to introduce new transactions into the blockchain they also secure previous ones, because the new blocks include a reference to previous ones. This means that the older a block is, the more reliable it becomes since more blocks can be used to verify it [6].

2.1 Burstcoin consensus

Burstcoin contrary to Bitcoin, bases its blockchain consensus algorithm on Proof of Capacity (PoC) [7]. While PoC also have miners perform work in the form of computational calculations, these are performed prior to providing the solutions. The resulting calculations are stored in a file called plot file, which at a later stage is used to provide solutions. In this sense more storage equals more solutions, meaning a greater chance of providing the best solution.

Compare this to PoW where more computational capacity equals higher probability to provide the best solution.

As such, one big difference between the two consensus algorithms lie in how the solutions are supplied to the blockchain network. PoW require constant computation to generate solutions and is as such very power consuming. For example, the Bitcoin mining operation consumed as much power as Ireland in July of 2014 [8]. Comparatively then, PoC should theoretically be vastly more energy efficient since computational work required to generate solutions is done once, contrary to constant, which is required for a PoW blockchain.

To prevent the mining from being performed on computational hardware such as CPUs, GPUs or ASICs in real time, the algorithm chosen as the basis for the calculational work was Shabal. This algorithm was chosen because it is to slow to be effectively computed in real time. Which ensures that the PoC consensus remain intact and keeps Burstcoin from becoming a PoW cryptocurrency.

2.2 Burstcoin mining

Before the mining process can begin, as mentioned, computational work must be performed to generate candidate solutions. This is known as plotting, which generates the plot file(s) required for mining. A chosen amount of storage is dedicated for the solutions, for example 500 GB. A plotting software is used to run the Shabal hashing function, which generates nonces. The nonces are divided into 4096 smaller places of data called scoops. Each scoop is compounded by 2 hashes, each 32 bytes large. As such every nonce is 524 288 bytes large, or rather 256 KiB [9].

Using our example from earlier, approximately 1 907 000 nonces can be stored in 500 GB space. Exactly how many depend on, among other things the file system cluster size formatting. The resulting nonces from the Shabal algorithm are stored sequentially in the plot file, from nonce 0, 1, 2 etc. Optimization of the plot file can be performed as shown in Figure 1, where the scoops rather than nonces are stored sequentially. This increases the speed of the mining process, because when performed it is searching for a specific scoop number [10]. Meaning less time will be spent searching for the scoop which without optimization would be spread over the hard drive. This is of greater importance when using mechanical drives since search time is higher than for alternatives, such as SSD drives.

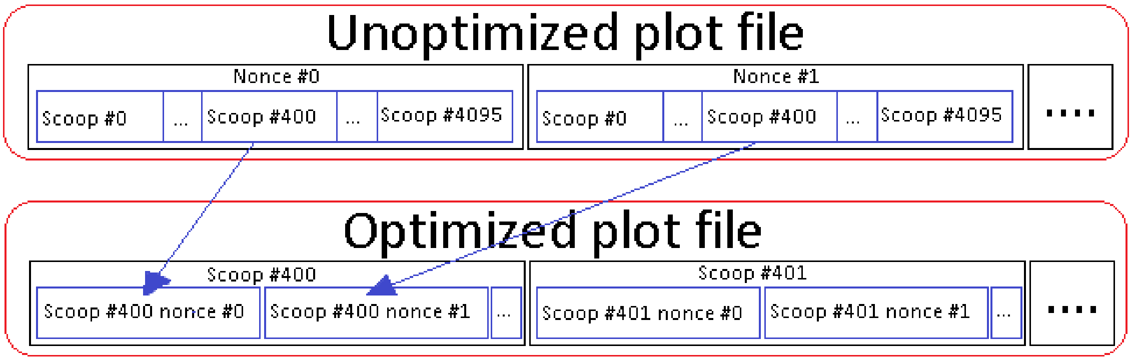

Figure 1 Plot file optimization

Once the plot file(s) are generated they are ready for processing by the mining software. First however the wallet is queried about previous block information by the mining software. The wallet is basically the communicational software layer between the mining software and the blockchain network. Following, the mining software performs a modulus 4096 on the information received from the wallet to know which scoop to look for [10]. When the mining software have this information, it can begin searching the plot file for the specified scoop number as seen in Figure 2. And Since one scoop is inspected per block the total amount which will be read is 1/4096 of the total plot file(s) size. Continuing from the earlier example using the 500 GB hard drive, the total amount of data which would be read from such a plot file is circa 122 MB.

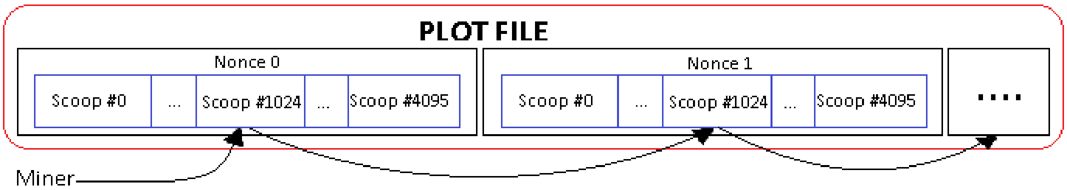

Figure 2 Burstcoin mining process

In the case of Burstcoin, the values which are searched for in the plot file(s) during mining are called deadlines, every time a new block can be forged these are searched for. The deadlines are resulting numerical values once a scoop has been ran through the Shabal algorithm. Each scoop with the number looked for in every nonce result in an individual deadline which represents a time. Once this time threshold is reached the deadline can be used to forge the next block, meaning shorter deadlines are better. However, most deadlines are very long and as such useless, which is why a definable limit can be set to determine which deadlines are interesting. In addition, the miner will ignore any deadlines which are longer than the current best one found [10].

3 Method

To find out which time limits constitute efficient mining and whether performance of the mining operation can be improved, a set of tests have been devised. First the distribution of block forge times is determined. Followed by tests to see if the mining process performance can be improved, either based on the file system cluster size, hard drive size used or SATA standard that the hard drive containing the plot file use. All testing is performed on two hard drives, one 500 GB hard drive and a 3000 GB hard drive.

Determining how many percent of blocks which will be fully participated in if processing is completed within any given time, for example 30 or 60 seconds, requires knowledge about block forge time distribution. This information was gained via inspection of the Burstcoin blockchain through the API of the Cryptoguru block explorer [11]. All block forge times up to block height 455966 (forged 2018-02-07) were gathered and compiled into histograms.

Finding out if changing the file system cluster size have any impact on performance, three identical tests were made using the three chosen file system cluster sizes (4, 64 and 256 Kb) for the two hard drives tested (500 and 3000 GB). All in all, six tests (three per hard drive) were conducted using the same plot file for each hard drive size respectively. To establish how long each test needed to be ran, a long test over 300 blocks was performed to compare different sample sizes. When compared it was determined that 50 blocks would be of sufficient size to yield good data. Once the tests were completed the data was compared for each hard drive to check if there was any difference in how fast the plot file was read and how much the CPU was taxed.

To ascertain if any difference exists regarding read speed of the plot file depending on hard drive size. First a plot file was placed on the 500 GB hard drive and the performance was measured over ten blocks. Following, the same file was moved to the 3000 GB hard drive to measure performance, again over ten blocks. The results are then compared to see if any difference exists in how fast the file is read.

Finally, to test whether there was any difference in processing speed depending on SATA standards 2 and 3. Two tests using the same hard drive, cluster size and plot file was conducted. Where the hard drive was connected to a SATA 2 port for one test and a SATA 3 port for the other. Subsequently both tests are compared.

Prior to gathering data regarding plot file read speed began, a baseline was established. Which consist of Windows 10 Pro x64, with the high-performance power schedule set. An i5 3470 with turbo boost disabled through BIOS, installed on a Hewlett-Packard 339A motherboard with chipset Intel Q75. Eight GB of DDR3 RAM and a system drive where the software required for the tests were located. A separate hard drive was used to store the plot file (the 500 and 3000 GB hard drives).

For the tests we chose two hard drives operating at 7200 RPM, one with 500 GB capacity, a Seagate ST500DM002-1BD142. And another with 3000 GB capacity, a Toshiba

DT01ACA300. To ensure consistency of the tests, only one plot file for each hard drive were generated, the plot file for the 500 GB hard drive was 499 950 551 040 bytes large, and the one for the 3000 GB hard drive was 3 000 175 165 440 bytes large. These two plot files were the only ones used during all testing. To note the sequential read speed of the drives they were benchmarked using HD Tune Pro [12], as seen in Table 1.

Table 1 Sequential read speed hard drive benchmarks

Hard drive	Min read speed	Max read speed	Avg read speed
3000 GB SATA 2	96 MB/s	196,2 MB/s	156,7 MB/s
3000 GB SATA 3	96 MB/s	196,2 MB/s	156,7 MB/s
500 GB SATA 2	68,6 MB/s	130,2 MB/s	107,4 MB/s

With the baseline established, the software tools to perform the tests and to supply performance data were chosen. Qbundle is a software suite which is officially recommended through the Burstcoin homepage [13]. It includes every tool necessary to participate in Burstcoin mining, including XPlotter and BlagoMiner. The plotting software is XPlotter, which generated the plot file(s) and optimise them regarding the scoops. BlagoMiner is the mining software which use a single core of the CPU, there exist several instruction sets which can be used when mining, the one used during testing is AVX.

To measure how long the processing of the plot file took. The Windows tool performance monitor was used. Specifically, the counters "\Process(BlagoMiner_avx)\% Processor Time" which measures how much the mining software taxes the assigned CPU resources (in this case one core) and "\PhysicalDisk("*nr drive-letter*")\Disk Read Bytes/sec" which measures how many bytes of data are read. The smallest polling interval available within performance monitor were chosen, which is 1 second. Additionally, log files automatically generated by BlagoMiner were used to determine how long the plot files were accessed by the BlagoMiner software. These logs were also used to verify which intervals in the performance monitor logs were inspected. When the tests were performed we extracted the interesting data from performance monitor with the help of the Windows command line tool relog. Relog takes the original performance monitor data collector file and outputs the chosen counters into a separate file.

4 Result

This chapter begins by displaying the analysis regarding how fast blocks are forged for the Burstcoin blockchain and which interesting time frames exist. Following three separate sections concerning performance are presented. First file system cluster size is compared to determine if it has any impact on the read speed of the plot file. The second performance factor investigated is whether the read speed of a plot file depends on the size of the hard drive. And finally, the results from an identical test conducted on a hard drive connected to SATA 2 and 3 ports respectively are shown.

4.1 Block forge times

The distribution of block forge time is vital information when determining how fast the system need to process the plot files, or rather how many percent of blocks will be fully participated in if all plot files are read before a certain time.

All forge times for blocks up until block height 455966 (forged 2018-02-07) are shown in three histograms with varying levels of resolution. Each bin in a histogram include values from a range. For example, in Figure 3, bin 60 include every value between 1 and 60 and bin 120 include every value ranging from 61 to 120 etc.

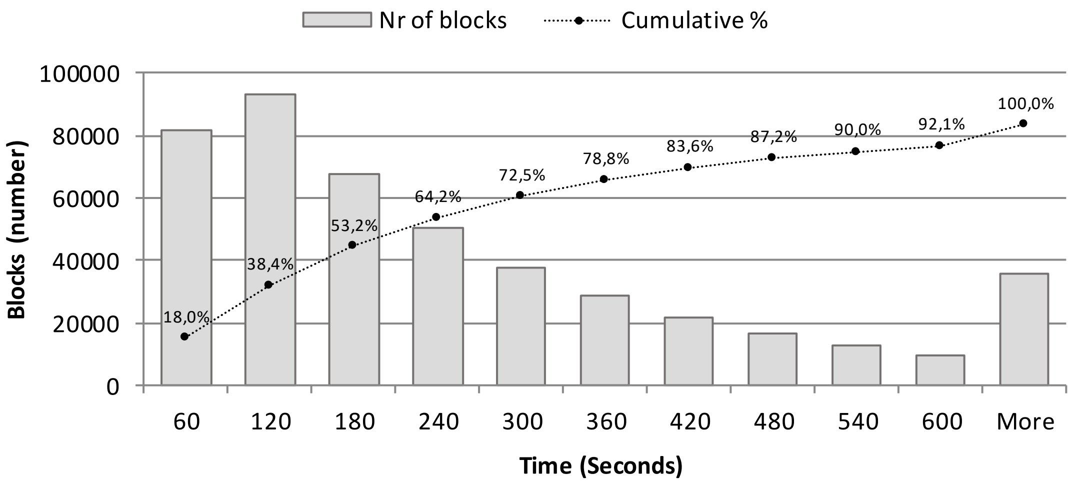

Figure 3 Distribution of block forge times set over a timespan of ten minutes

To find out how the distribution of block forge time look we included a time-span of ten minutes in Figure 3. 7,9 percent of blocks had greater forge times than ten minutes and were sorted into its own bin. The ten minutes demonstrate the timespan during which most blocks are forged. Noteworthy is that 82 percent of all blocks have a forge time which is greater than 60 seconds.

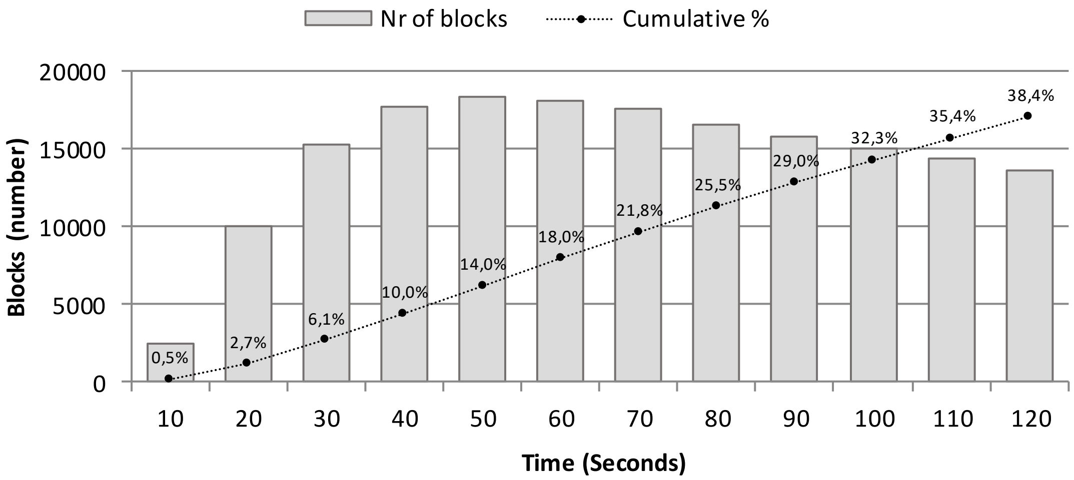

Figure 4 Distribution of block forge times set over a timespan of two minutes

When the resolution is increased a somewhat linear incrementation of between 3-4 percentage points can be noted for bins 20 through 120 as demonstrated in Figure 4. 61,6 percent of all blocks are excluded. This linearity shows that there exist no specific timeframes which are of interest regarding how fast the plot file(s) need to be processed. One point of interest however, is the 50 bin which displays the ten seconds (41-50) during which most blocks are forged. Following this bin (50), subsequent ones decrease in number of blocks forged within each ten second interval in a falling manner.

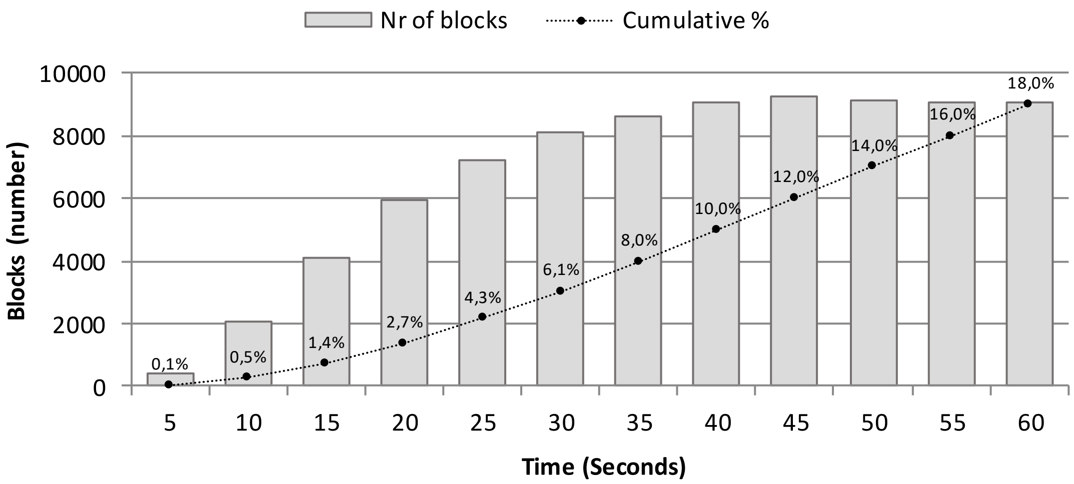

Figure 5 Distribution of block forge times set over a timespan of one minute

To display a more granular picture of the first minute, the resolution is further increased to focus on these blocks. This reveals a completely linear incrementation of 2 percentage points ranging from bins 35 through 60 as shown in Figure 5. 82 percent of all blocks are excluded. To give an example the impact of having a fully processed plot file will have depending on how long it takes to process, a comparison between bin 30 and 40 can be made. If the plot file is fully processed at 30 seconds, full participation will be had in 93,9 percent of all blocks, this can be compared to 90 percent which would be the case if the plot file was fully processed at 40 seconds, a gain of 3,9 percentage points if the plot file is processed in 30 seconds compared to 40.

4.2 File system cluster size

To demonstrate the different processing speed of the plot file using different file system cluster sizes of the hard drives, two sets of diagrams are presented below. The first set focuses on the read speeds and the second investigates how the CPU load look.

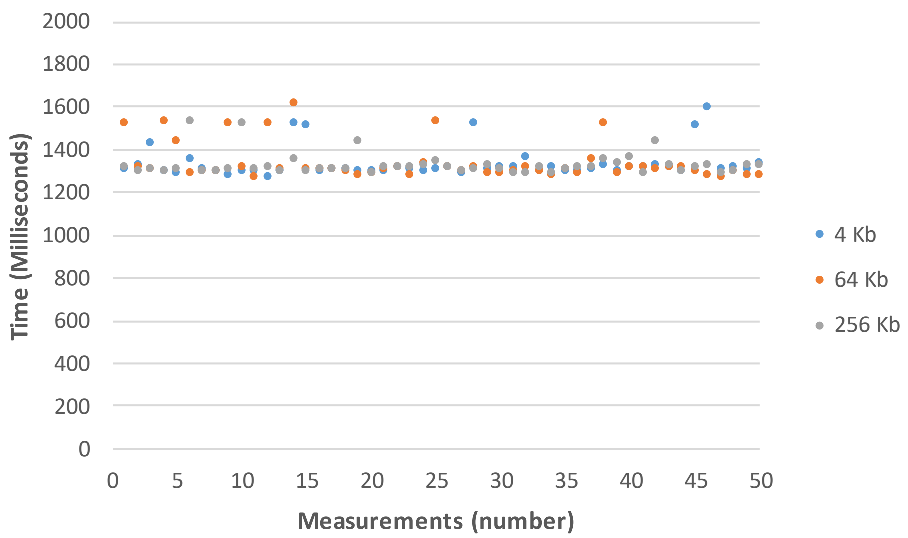

Figure 6 Fifty measurements per cluster size on a 500 GB hard drive

As shown above in Figure 6 there is no discernible difference in how fast the plot file is read on a 500 GB hard drive depending on the file system cluster size. A few outliers are shown however, but the distribution depending on cluster size is even.

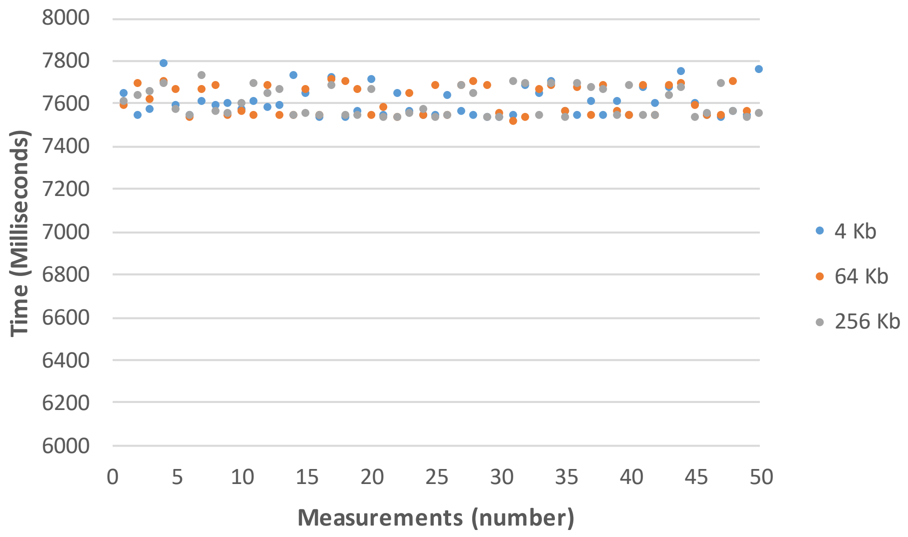

Figure 7 Fifty measurements per cluster size on a 3000 GB hard drive

Just as with the 500 GB hard drive, the 3000 GB hard drive as seen in Figure 7 produces no difference in how fast the plot file is read depending on the file system cluster size.

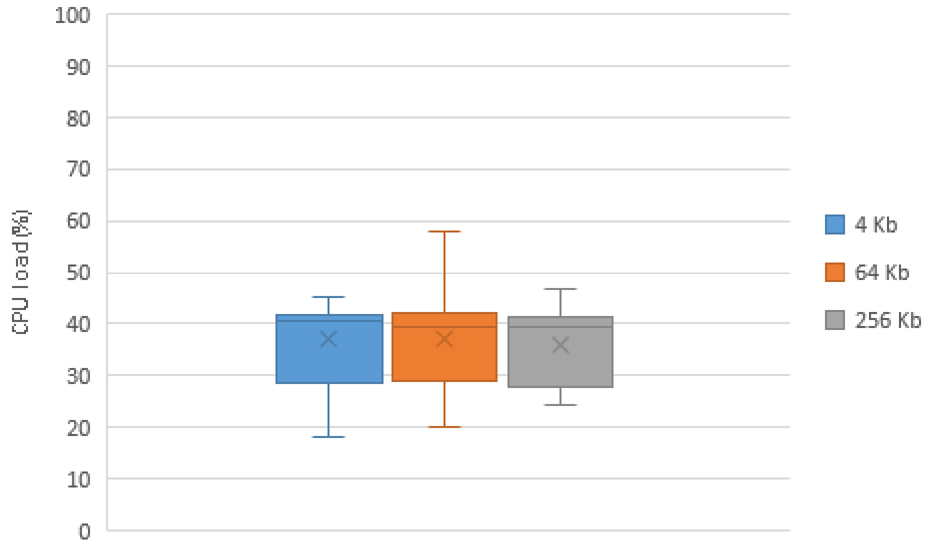

Figure 8 CPU load over fifty measurements per cluster size on a 500 GB hard drive

Using a box diagram in Figure 8 the variance of CPU load for each cluster size is shown. The X signifies the mean value, while the line running through the box represent the median. The upper and lower quartiles are shown as the upper and lower edges of each box and the edge of each line which stretches out from each box display the uppermost and lowermost load value noted. Noticeably there is very little difference in how the CPU load looks regardless of file system cluster size used. An additional note is that the CPU is not even close to maxing out, meaning the read speeds presented earlier is not bottlenecked by the CPU.

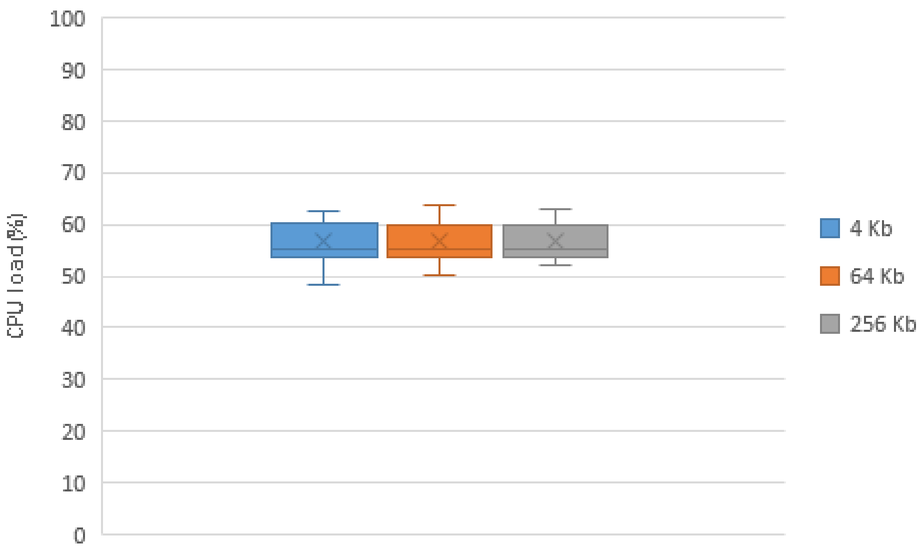

Figure 9 CPU load over fifty measurements per cluster size on a 3000 GB hard drive

In Figure 9 it can be seen just as with the 500 GB hard drive that there is virtually no difference in how the CPU is taxed for each file system cluster size used. The reason for the seemingly higher load is due to how performance monitor presents data in strict seconds. Meaning that what is shown in Figure 8 is the load of the CPU when the plot file is being read on the 500 GB hard drive averaged over three seconds, despite that the operation only really takes 1,3 seconds.

4.3 Comparing hard drives

The results from the tests that determines whether the size of the hard drive which a plot file is placed on matters are displayed below. The data from the blocks used is presented in Table 2 and following displayed as a graphic comparison in Figure 10.

Table 2 Read speed performance, 500 GB and 3000 GB hard drives

Block average		1	2	3	4	5	6	7	8	9	10	Avg
500 GB	**MB/s**	94,1	95,4	94,7	94,1	94,3	95,4	94,7	95,4	94,7	94,7	94,7
3000 GB	**MB/s**	94,7	96	97,3	97,3	96,6	96,6	96	96,7	96,6	96,6	96,5

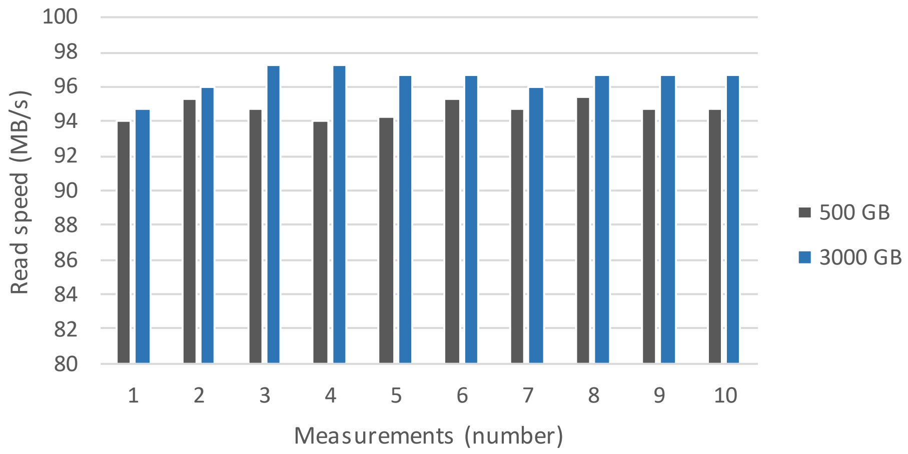

Figure 10 Read speed comparison of the same file on a 500 GB and 3000 GB hard drive

Above in Figure 10 the recorded average read speeds of both hard drives are compared as an individual bin per block. There is no big difference in how fast both drives read the same plot file, but there is however a persistent difference which averages to 1,8 MB/s in favour of the 3000 GB hard drive.

4.4 SATA standards

The difference in how fast the plot files were read differed noticeably when the 3000 GB hard drive was connected to a SATA 3 port contrary to a SATA 2 one. Figure 11 display this difference, the ten blocks compared are shown in Table 3.

Table 3 Plot file processing rate SATA 2 and 3 for a 3000 GB hard drive

Block average		**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Avg**
SATA 2	**MB/s**	96,8	95,5	96,5	95,2	95,8	97,5	96	95,7	97,4	97,2	96,4
SATA 3	**MB/s**	114,1	111,6	114,6	114	113,9	113,9	113,8	113,8	113,7	111,8	113,5

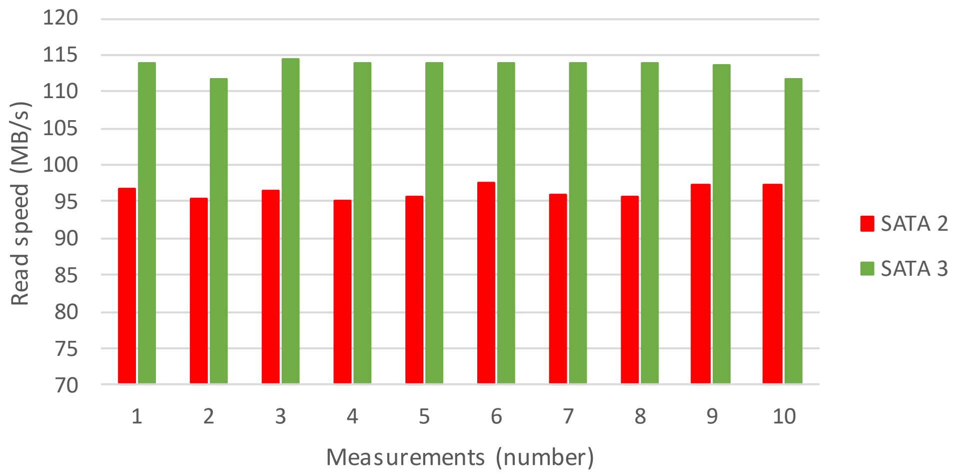

Figure 11 Average read speed on the 3000 GB hard drive using SATA 2 and 3 standards

To visualize the difference between the two standards, Figure 11 present the comparison. For every block measured the first and last seconds during processing were excluded to display a value representative of the maximum read speed of the hard drive. When averaging the read speeds for all blocks per standard, SATA 2 speeds reached 96,4 MB/s while SATA 3 reached 113,5 MB/s. This totals a 17,1 MB/s difference, this even though the benchmarks of the hard drives displayed in Table 1 showed that the hard drive provide equal sequential read performance regardless SATA standard used.

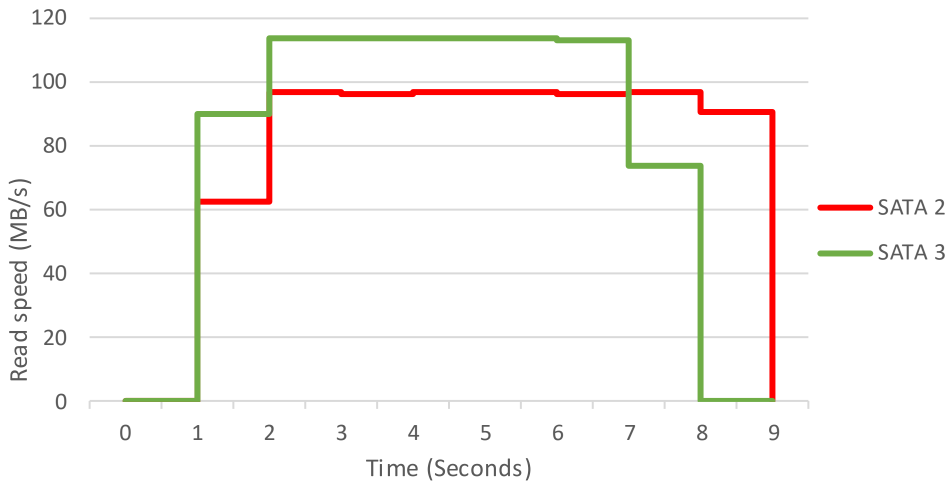

Figure 12 SATA read speed difference, single block

To display a more detailed overview regarding the difference between read speeds using SATA 2 and 3, one block is inspected per standard as seen in Figure 12.

5 Analysis

The analysis of our results regarding block forge time distribution and performance impact the chosen factors have are described below. Beginning with the time distribution of forged blocks. Following the performance results are analysed starting with file system cluster size, secondly hard drive size impact is determined, and finally SATA standards 2 and 3 are compared.

To begin with, block forge time distribution was established to identify which timeframes exist and are interesting. In the section Block forge times, all blocks up until block height 455966 which were forged 2018-02-07 are presented in varying timespans. More specifically, it can be seen in Figure 3 that 38,4 percent of all block forge times occur under the first two minutes, following the number of blocks forged start to decrease and only another 34,1 percent of all blocks forged are forged between two and five minutes. Finally, from five minutes to ten, only another 19,6 percent of all blocks are forged. When increasing the resolution to focus on two and one minutes in Figure 4 and Figure 5 the percent gain of full participation in blocks is quite linear regardless of how fast the plot file is processed. This linearity contrasts quite much with information that are sometimes given by other miners [2], which specify that it is of some importance to have finished processing of a plot file within 30 seconds.

It is also seen in Figure 5 that the percental gain is roughly two percent per five seconds ranging bins 25 to 60, meaning if the plot file processing is finished five seconds faster an additional two percentage points of full participation in all blocks being forged will be had. If comparing bin 25 to 60 it can be noted that full participation in blocks forged are 95,7 percent compared to 82 percent at 60 seconds, this is a difference of 13,7 percentage points or roughly two percentage points per five seconds. One fact which is important to consider when comparing numbers in this way is that blocks not fully participated in due to the plot file not being completely processed when the block is forged, is that some percent of the plot file will be read, meaning it is not a complete loss in participation of that block if the plot file is not fully processed when the block is forged.

In the section File system cluster size, the impact file system cluster size has on performance is investigated. As shown in Figure 6 and Figure 7 there is no difference in how fast a plot file is processed. However, as can be seen in Figure 6, some outliers exist which are likely a result depending on where the scoop being read is located on the hard drive, since files located on tracks further in on a hard drive takes longer to read than those located further to the edge. The result that there is no difference is further verified in Figure 8 and Figure 9 which shows that the CPU is not bottlenecking the read performance of the hard drives.

The section Comparing hard drives in the result chapter shows that there is a slight gain in performance depending on the hard drive size which a plot file is placed on. This increase is seen in Figure 10 and amounts to an average increase of 1,8 MB/s in favour of the larger hard drive of the two tested.

Finally, in SATA standards, a comparison of SATA 2 and 3 standards are made. This comparison shows a big difference in how fast a plot file is being read. An average difference of 17,1 MB/s is gained from using SATA 3 over the 2 standard. As shown in Figure 11 ten blocks are compared.

6 Discussion

When the subject of this thesis was chosen, neither of us had any prior knowledge about Burstcoin. And since it differs greatly in how the mining process is conducted in comparison to more well-known cryptocurrencies such as Bitcoin, it attracted our attention. The big difference regarding the mining process lie in how candidate solutions which enables transactions to be made over a blockchain are calculated and presented to the network. Where a PoW cryptocurrency such as Bitcoin have the miners continuously generate said solutions, the Burstcoin PoC cryptocurrency have the miners perform calculations once, store the solutions in plot file(s) which are later used to supply the solutions to the blockchain network.

Prior to this work, as far as we know, very little, if any, have been done to analyse the performance of the Burstcoin mining operation. This results in a high probability that the method generally used when mining Burstcoin is suboptimal because the mining operation is configured based on speculation. For this reason, the thesis has focused on mining performance, to see whether any improvements can be made regarding efficiency. To begin with the time distribution of forged blocks were investigated. Once this information was known, three factors which might affect performance were examined. The first one was file system cluster size, secondly it was determined whether the disk size matter in how fast a plot file is read. And finally, SATA standards 2 and 3 were compared.

To gather data regarding performance of the mining process, the Windows tool performance monitor were chosen. Early on problems were encountered regarding interpretation of performance data which was outputted by performance monitor. The tool is very powerful and can display many details regarding the operation of a Windows PC in a granular manner. Performance monitor sorts data through counters, so many exist that several reports data about the same thing, but in a slightly different manner. As a result, at first the wrong counters were chosen, after many hours analysing the resulting files outputted by performance monitor it was possible to determine which should be interpreted regarding performance. Another problem experienced with the tool was how the data was presented natively as graphs and not numbers. This was solved with the Windows command line tool relog, which can convert performance monitor files into csv files for use in Excel. One more factor which caused some difficulty concerning interpretation of the data, was that performance monitor polling granularity is limited at one second. All these things combined, lead to that the biggest challenge of this thesis by far was interpretation of the data outputted by performance monitor and is why the work might have been more streamlined if another tool were chosen.

Of all results, the most surprising one was how fast the plot file were processed when the SATA 3 standard were used contrary to SATA 2. Prior to testing it was reasoned that there might exist a small difference in how fast the file might be read, even though that the SATA 2 standard have sufficient throughput to handle anything a hard drive can throw at it, but not the big increase as seen in Figure 11. Even more interesting is the fact that the hard drive had identical performance using both SATA 2 and 3 standards when it was benchmarked as seen in Table 1. Its uncertain as to why this big difference exist, but one possible reason might lie in how BlagoMiner is programmed.

Another area of focus in this work orbited around how the time distribution of forged blocks look with the goal of establishing which interesting timeframes exist when setting up a Burstcoin mining operation. The resulting information gained verified one expectation, that the mining rig used during testing most likely would participate fully in most blocks. As seen in Figure 5, only 0,1 percent of all blocks weren't fully participated in when tests ran using the 3000 GB hard drive as can be seen in Figure 7 where the plot file processing time were around 7,6 seconds. One point of interest following the results however, is that there is no reason for having fully processed plot file(s) within any specific timeframe, for example before 30 seconds which is stated by some miners [2].

When investigation went into whether the file system cluster size have any impact on how fast a plot file is processed, it was uncertain if any difference would exist. However, statements such as [3] say the file system cluster size of 64 Kb is optimal and as such was worthy of investigation. Though once the tests were finished it was determined that there exists no discernible difference, as can be shown in Figure 6 and Figure 7. The results nevertheless show that there exist outliers, most easily visible in Figure 6. These outliers are likely a result from where the scoop being read for a block exist on a hard drive, as files located on tracks further in on a hard drive take longer to read. However, one benefit which were observed was that increasing file system cluster size resulted in that more nonces would fit on the same storage size. Between testing, when moving a plot file between hard drives when changing the file system cluster size, it was noted that more free space became available.

There are several aspects regarding mining performance which were not considered in the testing conducted during this work, for example if the instruction set used by the mining software (BlagoMiner) have any impact on how fast a plot file is processed. For example, it was not possible to determine if the newest instruction set available for Burstcoin mining (AVX2) would improve the speeds of plot file processing, since the CPU used in testing do not support it. Another thesis which explores this area, among others is "Utforskande studie om prestanda i utvinning av kryptovalutan Burstcoin" [14].

7 Conclusion

As shown in the section Block forge times it is not vital to have fully processed plot files within sixty seconds where only 18 percent of blocks have shorter forge times. This is further explained in the section Analysis, where the result is dissected and points to the fact that the block forge time is very linear regarding block forge time distribution. Also worthy of consideration is that even if the plot file is not completely processed, some percentage of it will be, meaning the blocks not fully participated in will not yield a complete exclusion for those blocks.

Another factor which were proved to be of no value regarding performance was file system cluster size. As can be seen in the section File system cluster size, no difference in how fast a plot file is processed exist. In the section Comparing hard drives, it was established that the size of the hard drive where a plot file is stored matters to some degree, in Figure 10 the difference averages to 1,8 MB/s. One factor which made a greater difference was the SATA standard which the hard drive was operating at. In the tests displayed in the section SATA standards, there was an average of 17,1 MB/s difference in how fast the plot file was read when comparing SATA 3 to SATA 2, with SATA 3 being faster. This was somewhat of a surprise since as noted in the benchmark as seen in Table 1 there was no difference in how fast the hard drive could read sequentially regardless of which SATA standard it was operating at.

8 Future work

During the work conducted for this thesis, several new points of interest regarding performance of the mining process used in Burstcoin became apparent and are listed below.

- Determine why SATA 3 improves read speed processing compared to SATA 2 when using BlagoMiner.
- Investigation where other performance bottlenecks exist, for example RAM or caching.
- Confirming how big an impact the location of scoops on a hard drive have on performance.
- Scaling up the tests conducted in this thesis through inclusion of much larger plot files to determine if performance is linear regarding plot file size.

9 References

[1] M. Swan, ”Technology Stack: Blockchain, Protocol, Currency,” i *Blockchain: Blueprint for a New Economy*, O'Reilly, 2015, pp. 1-2.

[2] SvenE, ”burstforum.net,” 2017. [Online]. Available: https://burstforum.net/topic/4970/question-about-miner-read-speeds. [Accessed 3 May 2018].

[3] haitch, “burst-team.us,” May 2017. [Online]. Available: http://burst-team.us/topic/5059/noob-question/14. [Accessed 3 May 2018].

[4] M. Swan, ”The Double-Spend and Byzantine Generals’ Computing,” i *Blockchain: Blueprint for a New Economy*, O'Reilly, 2015, p. 2.

[5] S. Nakamoto, “bitcoin.org,” 3 Januari 2009. [Online]. Available: https://bitcoin.org/bitcoin.pdf. [Accessed 3 May 2018].

[6] A. Narayanan, J. Bonneau, E. Felten, A. Miller and S. Goldfeder, Bitcoin and Cryptocurrency technologies, Princeton University Press, 2016.

[7] S. Gauld, F. v. Ancoina and R. Stadler, ”https://burstwiki.org/wiki/Whitepaper:Burst,” 27 December 2017. [Online]. Available: https://dymaxion.burst.cryptoguru.org/The-Burst-Dymaxion-1.00.pdf. [Accessed 3 May 2018].

[8] D. Malone and K. J. O’Dwyer, ”Bitcoin Mining and its Energy Footprint,” Hamilton Institute National University of Ireland Maynooth, 2014.

[9] Quibus, ”https://burstwiki.org/wiki/Main_Page,” 13 October 2017. [Online]. Available: https://burstwiki.org/wiki/Technical_information_to_create_plot_files. [Accessed 3 May 2018].

[10] Quibus, ”https://burstwiki.org/wiki/Main_Page,” 21 October 2017. [Online]. Available: https://burstwiki.org/wiki/Technical_information_about_mining_and_block_forging. [Accessed 3 May 2018].

[11] ”https://wallet.burst.cryptoguru.org:8125/index.html,” [Online]. Available: https://wallet.burst.cryptoguru.org:8125/test. [Accessed Februari 2018].

[12] ”HD Tune Website,” http://www.hdtune.com, [Online]. Available: http://www.hdtune.com.

[13] ”http://www.burst-coin.org,” [Online]. Available: http://www.burst-coin.org/download-wallet. [Accessed 2018].

[14] M. Hilmerson and J. Carlberg, ”Utforskande studie om prestanda i utvinning av kryptovalutan Burstcoin,” University West, Trollhättan, 2018.