

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/332810286>

# A Lyra2 FPGA Core for Lyra2REv2-Based Cryptocurrencies

Conference Paper · May 2019

DOI: 10.1109/ISCAS.2019.8702498

CITATIONS

9

READS

148

4 authors, including:



[Michiel Van Beirendonck](#)

KU Leuven

20 PUBLICATIONS 194 CITATIONS

[SEE PROFILE](#)



[Pascal Giard](#)

École de technologie supérieure, Montréal, Canada

62 PUBLICATIONS 1,439 CITATIONS

[SEE PROFILE](#)



[Alexios Balatsoukas-Stimming](#)

Eindhoven University of Technology

127 PUBLICATIONS 3,028 CITATIONS

[SEE PROFILE](#)

# A Lyra2 FPGA Core for Lyra2REv2-Based Cryptocurrencies

Michiel van Beirendonck\*, Louis-Charles Trudeau†, Pascal Giard†, and Alexios Balatsoukas-Stimming\*

\*Telecommunications Circuits Laboratory, École polytechnique fédérale de Lausanne (EPFL), Lausanne, Switzerland

†Electrical Engineering Department, École de technologie supérieure (ÉTS), Montréal, Canada

**Abstract**—Lyra2REv2 is a hashing algorithm that consists of a chain of individual hashing algorithms and it is used as a proof-of-work function in several cryptocurrencies that aim to be ASIC-resistant. The most crucial hashing algorithm in the Lyra2REv2 chain is a specific instance of the general Lyra2 algorithm. In this work we present the first FPGA implementation of the aforementioned instance of Lyra2 and we explain how several properties of the algorithm can be exploited in order to optimize the design.

## I. INTRODUCTION

Recently, there has been a surge in the popularity of cryptocurrencies, which are digital currencies that enable transactions through a decentralized consensus mechanism. Most cryptocurrencies are based on a *blockchain*, which is an ever-growing list of transactions that are grouped in blocks. Individual blocks in the chain are linked together using a cryptographic hash of the previous block, which ensures resistance against modifications, and every transaction is digitally signed. A blockchain needs to be protected from the double spending problem (i.e., an attacker spending the same digital money twice) and this is generally achieved by using a proof-of-work (PoW) system. This system requires that new blocks provide proof that a certain amount of processing power went into constructing them before they get accepted in the chain. For cryptocurrencies, this is typically achieved by appending random numbers to a block until its cryptographic hash meets a certain condition. The chain with the most cumulative proof-of-work (PoW) is accepted as the correct one, so that an attacker must control more than half of the active processing power to perform a double-spend attack. Processing nodes that help to compute the hashes of new blocks (called *miners*) are rewarded with a fraction of the cryptocurrency.

The first cryptocurrency, i.e., Bitcoin [1], was initially mined using desktop CPUs. Then, GPUs were used to significantly increase the hashing speed. Eventually, GPU mining was outpaced by FPGA miners, which were in turn surpassed by ASIC miners. Nowadays, the majority of the computing power on the Bitcoin network is found in large ASIC farms, each operated by a single entity, which makes the decentralized nature of Bitcoin debatable. To solve this issue, new PoW algorithms have been proposed that aim to be ASIC-resistant. ASIC resistance is achieved by using hashing algorithms that are highly serial, memory intensive, and parameterizable so that a manufactured ASIC can be easily made obsolete by simply changing some of the parameters, meaning that GPU mining is




Fig. 1. The Lyra2REv2 chained hashing algorithm.

much more cost-effective. However, since GPUs are generally much less energy-efficient than ASICs, a massive adoption of ASIC-resistant cryptocurrencies would significantly increase the (already very high) energy consumption of cryptocurrency mining. FPGA-based miners, on the other hand, are flexible, energy efficient, and readily available to the general public at reasonable prices. Thus, they are an attractive platform for ASIC-resistant cryptocurrencies.

A prime example of an ASIC-resistant hashing algorithm is Lyra2REv2 (used by Vertcoin [2], MonaCoin [3], and other cryptocurrencies), whose chained structure is shown in Fig. 1. The BLAKE, Keccak, Skein, BMW, and CubeHash hashing algorithms are well-known and have been studied heavily, both from a theoretical and from a hardware implementation perspective, as they were all candidates in the SHA-3 competition. On the other hand, to the best of our knowledge, no hardware implementation of the version of Lyra2 [4], [5] that is used in the Lyra2REv2 algorithm has been reported in the literature.

**Contributions:** In this paper, we present the first hardware implementation of the version of Lyra2 used in Lyra2REv2 as a stepping stone towards the implementation of an energy-efficient FPGA miner for Lyra2REv2 cryptocurrencies. Post-layout results for two Xilinx FPGAs show that our proposed Lyra2 hardware architecture consumes very few FPGA resources to achieve a hashing throughput between 2.6 MHash/s and 3.7 MHash/s with an energy efficiency between 432 nJ/Hash and 323 nJ/Hash.

## II. BACKGROUND

In this section, we provide the necessary background on some components of the Keccak and BLAKE2 hashing algorithms, since they are also used in Lyra2.

### A. The Keccak Duplex

Keccak is a family of hashing algorithms based on a cryptographic sponge [6], [7]. A cryptographic sponge is a function that takes an arbitrary-length input to produce an arbitrary-length hashed output. Lyra2 uses a specific implementation

**Algorithm 1** The G-function of BLAKE2b as used in Lyra2

---

```

1: INPUTS:  $a, b, c, d$ 
2: OUTPUTS:  $a', b', c', d'$ 
3:  $a' \leftarrow a + b$ 
4:  $d' \leftarrow (d \oplus a') \ggg 32$ 
5:  $c' \leftarrow c + d'$ 
6:  $b' \leftarrow (b \oplus c') \ggg 24$ 
7:  $a' \leftarrow a' + b'$ 
8:  $d' \leftarrow (d' \oplus a') \ggg 16$ 
9:  $c' \leftarrow c' + d'$ 
10:  $b' \leftarrow (b' \oplus c') \ggg 63$ 

```

---

of the sponge, called the *duplex* construction, which has a state that is preserved across different inputs. The duplex construction with naming conventions as adopted in Lyra2 can be found in [8, Fig. 2]. It consists of a permutation function  $f$  that operates on a  $w$ -bit state vector, where  $w = b + c$  and the parameters  $b$  and  $c$  are called the *bitrate* and the *capacity* of the sponge, respectively, as well as a padding rule  $\text{pad}$ . We note that the permutation  $f$  is iterative and performs a pre-defined number of iterations, also called *rounds*.

A call to the duplex construction proceeds as follows. An input string  $M$  is first fed into the duplex. Then, it is padded to length  $b$  and XOR'd into the lower  $b$  bits of the state. The state is then fed through the permutation  $f$ . The output of  $f$  is the new state of the duplex, while its lower  $l$  bits are the output hash, where  $l < b$ . If we consider the duplex construction as an object  $H$ , then the aforementioned procedure is referred to as a method  $H.\text{duplex}(M, l)$ . The following two auxiliary methods are useful to simplify the notation:  $H.\text{absorb}(M)$  updates the state using the input  $M$  but discards the output (equivalent to  $H.\text{duplex}(M, 0)$ ), while  $H.\text{squeeze}(l)$  reads  $l$  output bits and then calls  $H.\text{absorb}(\emptyset)$ , where  $\emptyset$  denotes an empty input string.

### B. The BLAKE2b Round Function

BLAKE2 [9] is a family of hash functions designed for fast software implementations. It is the successor of BLAKE as submitted to the SHA-3 competition [10]. The Lyra2 algorithm heavily draws from the round function of BLAKE2b, the 64-bit variant of BLAKE2. The round function consists of an arrangement of blocks that apply a so-called G-function to a 16-word state, where one G-function operates on 4 different state words. For BLAKE2b a word has 64 bits meaning that 16 state words amount to 1024 bits. The total round transforms these 1024 bits using four G-blocks, rearranges the output, and then does a four G-block transformation again. Algorithm 1 describes the modified BLAKE2b G-function as used in Lyra2.

## III. THE SIMPLIFIED LYRA2 ALGORITHM OF LYRA2REV2

Lyra2 was initially created as a password hashing scheme (PHS) for secure storage [4], [5].

Lyra2 uses the duplex construction from Keccak, where the permutation function  $f$  is the round function from BLAKE2b. In the remainder of the text, calls to a full-round (i.e., 12 iterations) duplex will be denoted as calls to  $H$ , while reduced-round duplexing as calls to  $H_\rho$ , where  $\rho$  denotes the reduced number of rounds. Because the G-functions are specified to

**Algorithm 2** The Lyra2 algorithm as specified in Lyra2Rev2.

---

```

1: PARAMS:  $H, \rho, \omega, T, R, C, k, b$  as  $H.b$ 
2: INPUT:  $pwd$ 
3: OUTPUT:  $K$ 
4:  $\triangleright$  Bootstrapping Phase
5:  $params \leftarrow \text{len}(K) || \text{len}(pwd) || \text{len}(pwd) || T || R || C$ 
6:  $H.\text{absorb}(\text{pad}(pwd || pwd || params))$ 
7:  $\triangleright$  Setup Phase
8: for  $col \leftarrow 0$  to  $C - 1$  do
9:    $M[0][C - 1 - col] \leftarrow H_\rho.\text{squeeze}(b)$ 
10: end for
11: for  $col \leftarrow 0$  to  $C - 1$  do
12:    $M[1][C - 1 - col] \leftarrow M[0][col] \oplus H_\rho.\text{duplex}(M[0][col], b)$ 
13: end for
14: for  $row^0 \leftarrow 2$  to  $R - 1$  do
15:    $prev^0 \leftarrow row^0 - 1$ 
16:    $row^1 \leftarrow row^0 - 2$ 
17:   for  $col \leftarrow 0$  to  $C - 1$  do
18:      $rand \leftarrow H_\rho.\text{duplex}(M[row^1][col] \oplus M[prev^0][col], b)$ 
19:      $M[row^0][C - 1 - col] \leftarrow M[prev^0][col] \oplus rand$ 
20:      $M[row^1][col] \leftarrow M[row^1][col] \oplus (rand \ll \omega)$ 
21:   end for
22: end for
23:  $\triangleright$  Wandering Phase
24: for  $row^0 \leftarrow 0$  to  $R \cdot T - 1$  do
25:    $prev^0 \leftarrow row^0 - 1$ 
26:    $row^1 \leftarrow \text{lsb}(rand) \bmod R$ 
27:   for  $col \leftarrow 0$  to  $C - 1$  do
28:      $rand \leftarrow H_\rho.\text{duplex}(M[row^1][col] \oplus M[prev^0][col], b)$ 
29:      $M[row^0][col] \leftarrow M[row^0][col] \oplus rand$ 
30:      $M[row^1][col] \leftarrow M[row^1][col] \oplus (rand \ll \omega)$ 
31:   end for
32: end for
33:  $\triangleright$  Wrap-up Phase
34:  $H.\text{absorb}(M[row^1][0])$ 
35:  $K \leftarrow H.\text{squeeze}(k)$ 

```

---

operate on an array of 16 64-bit words, Lyra2 uses a duplex with a width of  $w = 16 \cdot 64 = 1024$  bits. Pseudocode for the simplified version of Lyra2 that is used specifically in Lyra2Rev2 is given in Algorithm 2 and can be compared to the original Lyra2 pseudocode available in [4]. In the following sections, we explain each phase of the simplified Lyra2 algorithm in more detail.

### A. Bootstrapping Phase

In the bootstrapping phase, the duplex is initialized with a state that depends on the password input  $pwd$ , a salt (which in Lyra2Rev2 is set to be equal to  $pwd$ ), and the parameters  $T$ ,  $R$ , and  $C$  by using a full-round  $\text{absorb}$ . The duplex  $H$  in Algorithm 2 internally uses a bitrate  $b = 768$  bits and a capacity  $c = 256$  bits. The  $H.\text{absorb}(\cdot)$  call on line 5, however, considers only inputs of 512 bits instead of  $b$  bits, so as to not overwrite the upper part of the initialization state, i.e., the 512-bit initialization value  $IV$  specified by BLAKE2b. This results in two full-round  $\text{absorb}$ s, where the first and second  $\text{absorb}$ s process  $(pwd || pwd)$  and  $\text{pad}(params)$ , respectively.

### B. Setup Phase

During the setup phase, an  $R \times C \times b$  memory matrix  $M$  is initialized using the single-round duplex  $H_1$ . During setup, rows are initialized from first to last, while columns within a row are initialized from last to first. From the second row onward a previous row is re-read, making it impractical to

only store parts of the memory matrix. Also, from the third row onward, in addition to the previous row, i.e.,  $prev^0$ , a specific pre-initialized row, i.e.,  $row^1$ , is revisited (i.e., read and updated) in a deterministic manner. Rows are re-read or revisited from the first to the last column. Revisited rows use a rotated version of the duplex output, where the rotation number is chosen as  $\omega = 64$  in Lyra2Rev2.

### C. Wandering Phase

The wandering phase is generally the most time-consuming phase and it proceeds similarly to the setup phase. Specifically, it revisits two rows  $row^0$  and  $row^1$ , where  $row^0$  is chosen deterministically but  $row^1$  is chosen in a pseudorandom fashion by using the least significant part of the duplex output. We note that the pseudorandom and deterministic row can collide, resulting in the operations on line 26 and 27 to sequentially read from and then write to the same matrix cell.

#### D. Wrap-up Phase

The wrap-up phase consists of a full-round absorb of a specific cell of  $M$  followed by a squeeze of the hashed output  $K$ . This specific cell is likewise pseudorandom, as it is selected as the first cell of the lastly revisited pseudorandom row. The requested squeeze length  $k = 256$  is lower than the bitrate  $b = 768$ , which means that the final output is provided directly from the duplex state without a permutation  $f$ .

#### IV. FPGA IMPLEMENTATION OF SIMPLIFIED LYRA2

In the current instance of Lyra2 as used in Lyra2REv2, the timecost parameter is  $T = 1$ , the number of rows in the memory matrix is  $R = 4$ , the number of columns in the memory matrix is  $C = 4$ , and the desired hashing output length is  $k = 256$ . We note that our architecture is optimized for these parameter values, but it can be modified relatively easily to accommodate potential parameter changes if a hard fork is decided. Moreover, for the aforementioned parameters, the memory matrix  $M$  is 1.5 kB in size, which is clearly not prohibitively large to be implemented either on an FPGA or on an ASIC. The claimed ASIC-resistance of the Lyra2REv2 algorithm comes from the fact that  $T$ ,  $C$ , and  $R$  can be increased easily if necessary.

The datapath of our proposed FPGA implementation of the simplified Lyra2 algorithm used in Lyra2REv2 is shown in Fig. 2, where the duplex construction with its state, round, and XOR input block can be clearly distinguished. The memory matrix  $M$  is mapped to a block RAM (BRAM). To reduce the complexity of the multiplexer (MUX) at the input of the duplex, the BRAM also contains constant vectors of  $b$  bits used during the bootstrapping and setup phases: an all-zero vector and the  $\text{pad}(\text{params})$  vector. We first describe a version of the hardware architecture where each round of the  $f$  function is executed in a single clock cycle (CC). We then describe how this basic architecture can be improved through pipelining.

### A. Basic Iterative Architecture

Our basic iterative Lyra2 architecture requires 68 CCs per hash: 24 for the bootstrapping phase, 16 for the setup and wandering phases, and 12 for the wrap-up phase.




Fig. 2. Datapath of our proposed Lyra2 FPGA architecture.

1) *Bootstrapping Phase:* During the bootstrapping phase, the duplex processes two 512-bit input blocks from  $\text{pad}(pwd \parallel pwd \parallel params)$  using a full-round absorb. In Lyra2Rev2,  $pwd = cube_{\text{out}}$ , with  $cube_{\text{out}}$  the output from CubeHash, the previous algorithm in the chain. Thus, as shown in Fig. 2, the  $(pwd \parallel pwd)$  vector is one of the inputs to the duplex’s MUX. On the other hand, the  $\text{pad}(params)$  vector is fed into the sponge by loading it on  $q_a$  while simultaneously loading the all-zero vector on  $q_b$ . Both constants are stored at known addresses in the BRAM, and are absorbed in a separate 12-round *Bootstrap* state. During bootstrapping, the duplex only receives an input vector in the first round. Hence, for subsequent rounds,  $q_a$  and  $q_b$  output the all-zero vector, and their sum is passed to the duplex via its input MUX.

2) *Setup Phase*: We split the setup phase into three distinct phases for convenience, namely *Setup0*, *Setup1*, and *Setup2*, which correspond to Lines 6–8, Lines 9–11, and Lines 12–20 of Algorithm 2, respectively. Similarly to the bootstrapping phase, the setup phase uses the all-zero vector stored in the BRAM. In the *Setup0* state, the squeezes input an empty message into the duplex and directly write the duplex output to the BRAM. To achieve that, the all-zero vector is output on  $q_a$ ,  $q_b$ , and  $q_c$ . *Setup1* reads the all-zero vector on  $q_b$ , but a specific vector from the BRAM on  $q_a$ . *Setup2* reads two vectors from  $q_a$  and  $q_b$ . Both the duplex output and the rotated duplex output are XORed with two other vectors from the BRAM, requiring the two XOR blocks in parallel illustrated in Fig. 2. On the control path, counters keep track of the various rows ( $row^0, row^1, prev^1$ ) and their corresponding columns to generate read and write addresses for the RAM.

3) *Wandering Phase*: The input to the duplex in the wandering phase is always the word-wise addition of two RAM cells. Both XOR blocks connected to the duplex output are used. As mentioned before, the pseudorandom and deterministic rows used in the wandering phase can collide. In hardware, this special case requires the output of one XOR block to input

to the other, while the write port of the first XOR block needs to be disabled to prevent write collisions on the RAM.

4) *Wrap-Up Phase*: Wrap-up inputs one RAM cell into the sponge and then processes it using a full-round absorb. For the following squeeze, the requested hashed-output length  $k$  is lower than the bitrate  $b$ , meaning that the duplex state at that point directly provides the output hash.

### B. Memory Matrix

In the wandering phase, up to two RAM cells need to be written and three RAM cells need to be read per CC. These operations cannot be spread over multiple CCs without affecting the overall throughput of the design. Therefore, we use standard two-port BRAMs along with multipumping and replication techniques [11] in order implement the required functionality. Replication provides extra read ports by physically replicating the BRAM while connecting the write ports to keep the two copies coherent. Multipumping operates the BRAM at double the clock frequency of the surrounding logic, which, together with replication, effectively provides four read ports and two write ports.

### C. Pipelined Architecture

Pipelining the BLAKE2b round function can greatly reduce the delay of the critical path, which extends from the RAM read ports to the RAM write ports in the basic iterative version described above. Eight pipeline stages in the round were found to optimally increase throughput/area. Each hash that is concurrently being processed by the core needs its own memory. However, extra RAM-based memory is readily available since the current Lyra2REv2 parameters result in a RAM depth much shallower than that of the FPGA BRAMs. With adequate scheduling, concurrent hashes write to the same BRAMs in distinct CCs. While read ports  $q_a$  and  $q_b$  feed the duplex,  $q_c$  and  $q_d$  feed the XORs with duplex outputs. When pipelining the round function,  $q_c$  and  $q_d$  therefore need to be delayed by as many CCs as there are pipeline stages. The extra read port that is unused in the basic architecture allows delaying the control path for  $q_d$  rather than using a delayed version of  $q_b$ , avoiding a long chain of 768-bit registers. Eight pipeline stages in the round increase the latency to 544 CCs per hash. On the other hand, eight hashes are processed concurrently and the achievable clock frequency more than doubles, so the overall hashing throughput is improved significantly.

## V. HARDWARE IMPLEMENTATION RESULTS

To the best of our knowledge, there is no FPGA implementation of Lyra2 in the literature. For this reason, we can unfortunately not provide comparative FPGA implementation results. Moreover, since in this work we only present a core for Lyra2 and not a full miner architecture, we can also not compare our implementation to existing FPGA miners for cryptocurrencies based on other PoW algorithms.

Table I presents post-fitting results for the Xilinx Virtex 7 485T FPGA featured on the Xilinx VC707 Evaluation Kit as well as for the Xilinx Zynq Ultrascale+ 7EV FPGA from the ZCU104 Evaluation Kit. The power-consumption estimation

TABLE I  
POST-FITTING RESULTS FOR XILINX FPGAS.

FPGA	Virtex 7 485T	Zynq Ultrascale+ 7EV
Area (slices or CLBs)	2 163 (2.85%)	1 153 (4.00%)
LUTs	6 047 (1.99%)	6 010 (2.61%)
Registers	8 296 (1.37%)	8 296 (1.80%)
RAM (kbits)	1 548 (4.27%)	1 548 (5.39%)
Frequency (MHz)	175	250
T/P (MHash/s)	2.58	3.69
Dyn. Power (W)	1.12	1.19
Energy Eff. (nJ/Hash)	432	323

was obtained using Xilinx's Vivado Power Estimator tool, where the timing constraints are those required for the operating frequencies of Table I, the switching activity is from the simulation of the Lyra2 core processing random input vectors, and the post-fitted design provided to the tool meets all timing constraints. Table I reports the estimated dynamic power for the Lyra2 core. The functionality of the Lyra2 core was verified against test vectors that were generated using CPUminer [12].

From Table I, looking at the number of slices or CLBs required for the Virtex and Zynq FPGAs, respectively, it can be seen that the proposed Lyra2 core amounts to less than 4% of the resources available. The amount of RAM occupied is the same for both FPGAs, however the usage share is greater for the Zynq as it has less RAM blocks than FPGAs from the Virtex series. Also, from Table I, it can be seen that the throughput is 2.58 MHash/s and 3.69 MHash/s for the Virtex and Zynq FPGAs, respectively. The estimated dynamic power consumption of the Lyra2 core is under 1.2 W for both FPGAs. As a result, the energy efficiency is estimated to be in the vicinity of 325 to 435 nJ/Hash.

## VI. CONCLUSION

In this paper, we presented the first hardware implementation of the Lyra2 hashing algorithm, tailored to Lyra2REv2, an ASIC-resistant chained hashing algorithm employed by a few cryptocurrencies. The key to achieve good throughput and energy efficiency is to efficiently map the memory matrix to FPGA RAM blocks and to pipeline the BLAKE2b round function. Based on post-fitting results for two Xilinx FPGAs, we believe that the proposed Lyra2 implementation is a promising core for the purpose of FPGA-based Lyra2REv2 mining.<sup>1</sup> For example, we showed that, for a Zynq Ultrascale+ FPGA featured on an affordable evaluation kit, the achievable throughput is of 3.7 MHash/s and the energy efficiency of 323 nJ/Hash, for a resource usage of 4%.

## ACKNOWLEDGMENT

The authors thank Jean-François Têtu for useful feedback. The authors also gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan Xp GPU.

<sup>1</sup>Our VHDL code and relevant scripts are publicly available at <https://github.com/Michielvb/lyra2-hw>.

## REFERENCES

- [1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.
- [2] "Vertcoin." [Online]. Available: <http://vertcoin.org>
- [3] "MonaCoin." [Online]. Available: <https://monacoin.org>
- [4] M. A. Simplicio Jr, L. C. Almeida, E. R. Andrade, P. C. dos Santos, and P. S. Barreto, "Lyra2: Password hashing scheme with improved security against time-memory trade-offs," Cryptology ePrint Archive, Report 2015/136, 2015. [Online]. Available: <https://eprint.iacr.org/2015/136>
- [5] E. R. Andrade, M. A. Simplicio, P. S. L. M. Barreto, and P. C. F. d. Santos, "Lyra2: Efficient password hashing with high security against time-memory trade-offs," *IEEE Trans. Comput.*, vol. 65, no. 10, pp. 3096–3108, Oct 2016.
- [6] G. Bertoni, J. Daemen, M. Peters, and G. V. Assche, "Cryptographic sponge functions," Tech. Report v0.1, Jan. 2011.
- [7] NIST, "SHA-3 standard: Permutation-based hash and extendable output functions," FIPS Publication 202, Aug. 2015.
- [8] M. A. Simplicio Jr, L. C. Almeida, E. R. Andrade, P. C. dos Santos, and P. S. Barreto, "The Lyra2 reference guide," Tech. Report v2.3.2, 2014.
- [9] J.-P. Aumasson, S. Neves, Z. Wilcox-O'Hearn, and C. Winnerlein, "BLAKE2: simpler, smaller, fast as MD5," in *Int. Conf. on Applied Crypto. and Netw. Security (ACNS)*. Springer, 2013, pp. 119–135.
- [10] J.-P. Aumasson, L. Henzen, W. Meier, and R. C.-W. Phan, "SHA-3 proposal BLAKE," Tech. Report v1.3, Dec. 2010.
- [11] C. E. LaForest and J. G. Steffan, "Efficient multi-ported memories for FPGAs," in *Ann. ACM/SIGDA Int. Symp. on FPGAs*, 2010, pp. 41–50.
- [12] T. Pruvot, "CPUMiner-Multi," GitHub repository, 2017. [Online]. Available: <https://github.com/tpruvot/cpuminer-multi>