

BitVMX: A CPU for Universal Computation on Bitcoin

Sergio Demian Lerner^{1,2}, Ramon Amela¹, Shreemoy Mishra¹, Martin Jonas², and Javier Álvarez Cid-Fuentes¹

¹ RootstockLabs

{sergio,ramon.amela,shreemoy,javier}@rootstocklabs.com

² FairGate Labs

{sergio,martin.jonas}@fairgate.io

Abstract. BitVMX is a new design for a virtual CPU to optimistically execute arbitrary programs on Bitcoin based on a challenge response game introduced in BitVM. Similar to BitVM1 we create a general-purpose CPU to be verified in Bitcoin script. Our design supports common architectures, such as RISC-V or MIPS. Our main contribution to the state of the art is a design that uses hash chains of program traces, memory mapped registers, and a new challenge-response protocol. We present a new message linking protocol as a means to allow authenticated communication between the participants. This protocol emulates stateful smart contracts by sharing state between transactions. This provides a basis for our verification game which uses a graph of pre-signed transactions to support challenge-response interactions. In case of a dispute, the hash chain of program trace is used with selective pre-signed transactions to locate (via *n*-ary search) and then recover the precise nature of errors in the computation. Unlike BitVM1, our approach does not require the creation of Merkle trees for CPU instructions or memory words. Additionally, it does not rely on signature equivocations. These differences help avoid complexities associated with BitVM1 and make BitVMX a compelling alternative to BitVM2. Our approach is quite flexible, BitVMX can be instantiated to balance transaction cost vs round complexity, prover cost vs verifier cost, and precomputations vs round complexity.

1 Introduction

BitVMX is a new design for a simple and efficient CPU that can optimistically verify the execution of arbitrary programs on Bitcoin. TrueBit [8, 9] introduced the *verification game* paradigm to blockchain systems. In this paradigm, if all parties agree on the result of an off-chain computation, then no computation is performed on-chain. In case of a dispute, an on-chain interaction formally solves the dispute through a challenge-response type verification game.

This paradigm gained mainstream awareness with the introduction of *fraud proofs* and *optimistic rollups*. More recently, we have seen the emergence of rollups using a different paradigm relying on *validation proofs* created using *Zero Knowledge* cryptography.

Such scaling solutions have been harder to build on top of Bitcoin. Some reasons include Bitcoin’s restricted scripting language, the lack of features related to transaction introspection, and a restricted execution environment that does not offer any native

method to persist state across transactions. This is where a recent proposal by Robin Linus, BitVM [6] made an important contribution by presenting the first attempt to bring the Disputable Computation paradigm to Bitcoin. Remarkably, Linus' proposal does not require any consensus changes or soft forks on Bitcoin. This paradigm is now a rapidly growing area of research. Most of the community interest is focused on two primary use cases: (i) building trust-minimized bridges for Bitcoin sidechains and rollups and (ii) optimistic verification of ZK proofs for rollups and other applications.

The original BitVM design was a theoretical proposition with restricted applicability. Later versions became more practical. BitVMX offers an alternative to these newer designs. BitVMX allows funds to be locked in a UTXO with a spending constraint that depends on the result of running a predefined program with a given input. In its simplest form, BitVMX is a two-party protocol where the first participant is the *prover* (also called the *operator*) and the counterparty is called the *verifier*. BitVMX can be extended to the N verifier setting with the assumption that at least one of the $N \geq 1$ verifiers is honest (as in BitVM).

When the operator wants to access the funds locked in the UTXO, he claims that he has executed the program correctly for some input of mutual interest, and shares that input with the verifier (off-chain). The operator first claims that the computation yields an approval for the spending of the UTXO. If the verifier does not contest this claim, then the operator can access the funds after a timelock. The timelock allows the verifier to execute the program locally with the reported input and verify the operator's claim. If the verifier detects cheating, she can challenge the operator and both enter a dispute protocol (verification game) on the Bitcoin blockchain.

BitVMX improves BitVM by simplifying the way computations are represented. The protocol requires both parties to run the program locally. When executing the program, each party simultaneously generates an *execution trace* and a *hash chain*, with one hash for each step. The hash chain is constructed by (i) appending the current step's trace to the previous step's hash, and then (ii) hashing the result securely. Starting from an initial state, each step's hash is computationally unforgeable and part of a *recursive hash chain*. Any difference in the program's execution trace across the parties will lead to diverging hash chains.

The trace representation allows BitVMX to perform *n-ary* searches, rather than just binary ones as in BitVM. This can significantly reduce the number of challenges required to find the conflicting computational step. Once a fault has been localized, a variety of specially crafted challenges can be used to determine the exact nature of the fault. For instance, BitVMX uses a new challenge-response protocol to detect faulty memory reads based on tracking the last time a memory word was written to.

All responses provided by the responder are uniquely linked to a specific challenge. The challenger can use the responder's own response at later stages of the protocol to prove cheating and ultimately force a dishonest operator to concede the funds. Using the responder's own signed message helps BitVMX avoid the use of on-chain transactions to prove equivocations, which results in a simpler protocol. Equivocations can be used, as in BitVM, they are simply not core to BitVMX.

2 Related work

TrueBit was the first blockchain system to propose a verification game to validate complex off-chain computations without running them on-chain. In TrueBit, a *solver* runs an arbitrary computation off-chain and posts the result on a smart contract for validation. The validation is performed through a verification game triggered by a challenger. As part of this game, the challenger forces the solver to reveal certain details of the execution in order to find a conflicting computational step. In case of equivocation, the smart contract executes the conflicting step to penalize the solver. This reduces the amount of computation that has to be executed on-chain to a single step and extends the computational capabilities of the underlying blockchain.

TrueBit's base concept was popularized by optimistic rollups on Ethereum. However, there was no such implementations on Bitcoin - until BitVM [6]. The first version of BitVM proposed a mechanism to create logic gate commitments in Bitcoin that enable the verification of arbitrary computations in a way similar to TrueBit. This is based on the fact that any computation can be expressed as a circuit of NAND gates. In BitVM, a prover and a verifier create a Taproot tree where each leaf represents one NAND gate. The two parties then pre-sign a sequence of challenge-response transactions that allow the verifier to challenge the inputs and output of any NAND gate in the circuit. The prover is forced to respond to the challenges by revealing hash preimages for the gate's inputs and output. When challenged, gates are executed and verified on-chain. After a series of challenges, the verifier can punish the (dishonest) prover by forcing the publication of incorrect or conflicting data on-chain.

Though conceptually interesting, working directly with logic gates is not a practical approach for verifying complex computations that must ultimately be run using Bitcoin script. Linus [2] presents an improved version, BitVM1, where computations are represented as a sequence of complex instructions of a virtual CPU. In BitVM1 each instruction consists of two inputs, an opcode, and an output. The state at each computational step is defined by the root of a Merkle tree that stores the contents of all the memory locations. A trace Merkle tree is built upon leaves that contain the memory states. The verification game consists of the verifier first identifying a conflicting computational step by binary search (or *bisection*) over the trace Merkle tree, and then showing that the prover committed incorrect information on-chain. To detect incorrect memory reads or writes, an additional binary search is performed over the hashes of a Merkle path in the memory Merkle tree related to the step in question. As in TrueBit, at the end of the process, the system executes at most a single instruction on-chain. Unlike BitVM1, BitVMX does not use Merkle trees to store the computational trace, the program, or the memory. Our use of hash chains for the trace allows BitVMX to perform *n-ary* searches. BitVM1 needs to partially or totally recompute the memory Merkle tree to reflect all memory updates performed at each step. BitVMX does not use trees, but it marks each memory word with an integer indicating the last step it was updated. This is more memory-efficient and also reduces computation costs.

Recently, the development of BitVM1 was paused in favor of a new approach in BitVM2 [5]. In BitVM2, the full computation to be verified is divided into a sequence of connected steps representing intermediate computations. The prover commits to the inputs, the final result of the entire computation, and to all of the intermediate results in

a single transaction. After this, *any arbitrary* party can penalize the prover for submitting an incorrect intermediate result by forcing the on-chain execution of the particular step re-generating the result. BitVM2 allows anyone to become the verifier, which is a significant security improvement. BitVM2 also reduces the number of transactions needed in the verification game. However, this approach lacks the generality afforded by BitVMX (and BitVM1). For every program, each intermediate computational step will require a custom implementation in Bitcoin script. To avoid having to write custom implementations for each program, BitVM2 proposes writing a SNARK [3] verifier in Bitcoin script. This will enable the verification of any computation – for which a SNARK proof can be generated. BitVMX, by contrast, is designed from the start as a general-purpose framework that enables verifying any program that can be compiled to common architectures, such as RISC-V, MIPS or any standard CPU.

3 Message linking scheme

Message linking is a simple cryptographic scheme used in BitVMX to link data referenced in sequential Bitcoin transactions. If a computation is disputed, then our verification protocol plays out on-chain over challenge and response transactions, the details of which appear in Section 5. We want these back and forth transactions to transfer some information about the state of the computation we wish to verify. Bitcoin has no native mechanism to transfer state data across transactions. So we present a framework to share state data - with a cryptographic link between a response message and the corresponding challenge message.

We also want to ensure that if an operator is challenged by a verifier, then the operator must respond, and do so in a manner that is consistent with previously agreed upon terms. The verifier’s challenges must also fit a pre-arranged format. Implementing such restrictions can be challenging because Bitcoin script is not Turing complete. The execution environment has no access to transaction data. So a coin’s spending restrictions are typically limited to checking signatures, hash locks and time locks. There is no op-code for concatenation and arithmetic operations are limited to words of at most 32 bits (4 bytes). One consequence of these limitations is the current impossibility of *stateful smart contracts* or *covenants* on Bitcoin.

Some of the above restrictions on introducing state into transactions can be softened by all parties pre-signing a set of related transactions that form a directed acyclic graph. The parties can commit to all inputs and outputs as desired, and then a subset of these pre-signed transactions can be broadcast in some sequence to reproduce covenant-like behavior. To make such configurations robust, every party has to be sure that the other parties have no way to spend the locked funds in a way that can break the intended flow. This is related to the concept of *connectors* introduced in the Ark protocol [1]. BitVMX uses pre-signed transactions to use connectors and force the operator to engage in the message linking protocol.

Our message linking scheme is composed as follows. Both parties commit to the state data they intend to communicate by using *one-time signature* schemes such as Lamport [4] or Winternitz [7] signatures. We use transaction templates to specify the format of challenge and response transactions. Since enforcing a predefined sequence of

transaction templates is not possible using Bitcoin script, the parties must jointly pre-sign the transactions generated from templates during a *setup phase*. When pre-signing, they must use appropriate SIGHASH flags to commit to all inputs and outputs - leaving only the witness data to be filled in later. This way, the transaction *ids* will be fixed from the start and we can create sequential transactions with dependencies, so state information can be communicated across them.

An on-chain verification game can begin only after the operator broadcasts a *claim* or *kickoff* transaction to initiate a withdrawal of funds from the locked deposits. The actual challenge-response state information to be exchanged in the process of the verification game will obviously not be available during the setup phase. This implies that the commitment of state data using one-time signatures must happen later, by adding them to the segregated witness stack of appropriate pre-signed transactions (created from templates). It is only after the valid state data is added to the witness stack that the pre-signed templates become valid transactions that can be broadcast.

More formally, to create a link between two messages, a challenger C and a responder R start by pre-signing three transactions together: a start transaction T_s , and two segregated witness transaction templates T_c and T_r that do not include their witnesses. Additionally, C and R generate one-time signatures with public keys P_c^o and P_r^o respectively.

The starting transaction T_s contains an output that can be spent by C by providing an arbitrary message x along with a one-time signature $\sigma_c^o(x)$ that is consistent with C 's one-time public key P_c^o . Figure 1 shows a diagram of the linking scheme. To start the message linking scheme C adds a message x and a one-time signature $\sigma_c^o(x)$ to the witness stack of transaction T_c . The function `OTVerify` in Figure 1 verifies a one-time signature against a pre-defined (one-time) public key. `OTVerify` is an example of a *meta* op-code. This is something that we must implement using Bitcoin script.

T_c creates an output that can be spent by the responder R . In order to spend this output, R must duplicate C 's message and signature $(x, \sigma_c^o(x))$ and add them to the witness along with a message y and one-time signature $\sigma_r^o(xy)$ of their own. R actually signs y after appending it to x . Thus, R 's signature $\sigma_r^o(xy)$ is over the concatenation of both messages $(x||y)$ and this must be consistent with R 's one-time public key P_r^o . Throughout this paper, $||$ denotes the concatenation operator.

Note that C can use their private key to rebroadcast a different value x' in the future. To avoid this, we require that if C has to refer to their own message x in some future transaction, then they must use the value from R 's response which contains both x and y . Since only C knows their one-time *private* key, it is safe for anyone else (including R) to rebroadcast C 's message x .

Note that C is forced to sign and reveal x in order to start the challenge and that R is forced to repeat the provided x and respond with their own signed message. T_c contains another output that pays to C if R does not publish T_r before a deadline, which forces R to sign and reveal y once T_c is published. The sequence of transactions can be extended to an arbitrary length by replicating the scheme after T_r .

Figure 1 is a highly simplified illustration of how message linking can be used in BitVMX. An actual implementation of the protocol will be more complex - because message linking is only a part of the challenge-response protocol. For instance, the initial deposits from each party will typically be locked in a transaction that is mined much

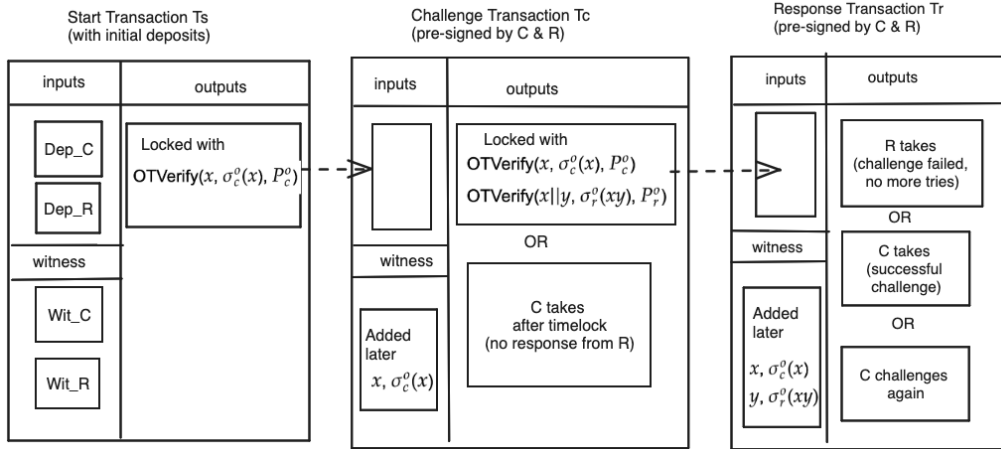


Fig. 1: Simplified illustration of using one-time signatures and pre-signed templates to link messages across transactions

earlier than shown here. Once the deposits are locked, all parties are committed to the protocol. The full protocol, presented later, will require pre-signing a set of transactions that form a directed acyclic graph. It is possible to pre-sign several transactions with the same input but different outputs. By doing so, we can create a decision tree where the party in charge of publishing a certain message at a particular step of the protocol chooses the branch through which the protocol should continue. When one path is chosen, some others get discarded as UTXOs associated with the alternative paths have been spent. The initial transaction generation, pre-signing and sharing process is called a *setup ceremony*.

4 CPU specification

BitVMX defines a virtual CPU to represent the computation that needs to be verified. This virtual CPU is specified by a set of instructions, addressable memory, and a representation of the CPU state. BitVMX provides a flexible design that allows the implementation of common architectures, from 8-bit CPUs to 32-bit RISC-V, without having to build a specific compiler for BitVMX.

A verifiable program in BitVMX is a sequence of executed instructions or steps in this virtual CPU, where each step is structured as follows:

Inputs							Outputs			
read1			read2			readPC	write		writePC	
address	value	lastStep	address	value	lastStep	address	opcode	address	value	pc

Table 1: Layout of an instruction executed by the CPU.

Every executed instruction reads at most two values and an op-code from memory, performs an operation, and writes the result to memory. Additionally, each instruction sets the address of the next instruction by modifying the program counter (`pc`). The field named `lastStep` stores the program step number where each of the read values was last written. Using this instruction structure, we provide the following definitions:

Definition 1. We define the `full_trace` of an execution step i as the concatenation of all the input and output fields in Table 1

$$\text{full_trace}_i = \text{read1.address}_i || \text{read1.value}_i || \text{read1.lastStep}_i || \dots || \text{writePC.pc}_i$$

Definition 2. We define the `trace` of an execution step i as the concatenation of data from just the output fields, that is,

$$\text{trace}_i = \text{write.address}_i || \text{write.value}_i || \text{writePC.pc}_i$$

The length of the memory addresses depends on the architecture. For example, they are 32-bit long for a 32-bit RISC-V. Registers are mapped to pre-defined memory addresses. For example, we can map r registers to the memory addresses from 0 to $r - 1$, and the program counter is the only register that does not have a memory address.

The state of the virtual CPU consists of the values in every memory address at each program step plus the `pc`. We represent *state transitions* using a hash chain where each member, denoted by `stepHashi`, is constructed recursively as follows.

Definition 3. Given a hash function h , we define `stepHash` for an instruction i as

$$\text{stepHash}_i = h(\text{stepHash}_{i-1} || \text{trace}_i).$$

With this recursive construction, each `stepHash` implicitly depends on every memory write operation from the start to the current instruction step. By using a secure hashing function, we can be almost certain that each `stepHash` is a unique representation of the memory state at each program step.

Arithmetic operations in Bitcoin script take 32-bit signed inputs. Thus, implementing 32 or 64-bit architectures requires a mechanism to handle data in smaller chunks using Bitcoin primitives. This means we cannot use the built-in hash functions available in Bitcoin script to compute the hash of the above objects. The chosen hash function, h , has to be implemented using basic arithmetic and logical primitives of Bitcoin script. Therefore, both the choice of a hash function and its implementation are important for efficiency and security of the protocol.

To initialize the CPU, we insert write instructions at *negative full_trace* positions that load the program into memory. That is, if a program to be executed has m instructions, then `full_trace-m` to `full_trace-1` contain only write operations for each instruction.

We set the starting point of the hash chain as

$$\text{stepHash}_{-m-1} = 0$$

and then use Definition 3 to compute `stepHashi` for all steps from $i = -m$ up to $i = -1$. In this way, `stepHash-1` represents the initial state of the CPU, where each program

instruction is in a particular memory address. The range of memory addresses that store the program is read-only. Additionally, we define another range of read-only memory addresses to store the program input. The rest of the memory addresses are initialized with the value 0 and their `lastStep` field is initialized with the special value `INITIAL`. Given some input, we can execute the program. If the program exits successfully at step $i = k$, then we set `stepHashi` = 1 for all $i > k$. In case of an exception, we set these values to a negative number.

5 Challenge-response protocol

In this section, we describe the challenge-response protocol that BitVMX uses to verify computations on-chain. The protocol is based on the message linking scheme described in Section 3 and involves a prover and a verifier.

During setup, the prover and the verifier agree on `stepHash-1`, and before running the protocol both parties initialize the CPU as described in Section 4. After this, the prover can trigger the challenge-response process by providing the input of the program and the `stepHash` value of the very last step. The protocol continues if the verifier disagrees with the final state of the CPU for the given input. In that case, the goal of the verifier is to identify an error in the prover’s computation. Towards this, the verifier starts by identifying the first conflicting computational step and then forcing the prover to reveal the data associated with that step. Then, the verifier identifies the precise nature of the error in this step and validates the error on-chain. In the following sections, we describe various parts of the challenge-response process in more detail.

5.1 Search mechanism

At various parts of the challenge-response protocol, the prover and the verifier use an n -ary search algorithm. Given a sequence of `stepHash` values, we split the search space into




Fig. 2: n -ary search mechanism with 64 hashes ($m = 64$, $n = 4$, $l = 3$). The concatenation of the interval identifiers selected produces $011001_b = 25$.

n intervals defined by $n - 1$ equally spaced hashes, where n is a power of two. We identify each interval with the *binary representation* of a number in the range 0 to $n - 1$. We then chose one of the intervals and split it into a further n sub-intervals. After $\log_n(m)$ such splits, let us assume we reach the i th hash in the sequence. Suppose we choose m such that $m = n^l$ for some fixed number of iterations l . Then the concatenation of the identifier for each interval selected during the search provides the index (in binary) of the i -th hash. Figure 2 illustrates an n -ary search process that reaches the 25th hash in the sequence after $l = 3$ iterations with $m = 64$ and $n = 4$.

Note that every interval of the search process can be defined with $n - 1$ hashes, because both left and right boundaries are provided in the previous iteration. For example, in Figure 2, the interval h_{16}, \dots, h_{31} can be defined by h_{19}, h_{23}, h_{27} .

5.2 Conflicting step search

A challenge is needed only when the two parties disagree on the last `stepHash` reported by the prover. Since the prover and the verifier agree on `stepHash-1` during setup and disagree on the last one, there must be two consecutive hashes `stepHashi` and `stepHashj`, such that `stepHashi` is correct and `stepHashj` is incorrect according to the verifier. The goal of this stage of the protocol is to identify a pair i, j for which this is true. We refer to `stepHashj` as the first conflicting step.

Figure 3 illustrates the first stage of the challenge-response protocol, which begins when the verifier challenges the final state of the CPU published by the prover. To identify the first conflicting step, the verifier runs an n -ary iterative search with the prover for a number of rounds as described in Section 5.1. In this process, the prover reveals $n - 1$ hashes and the verifier responds by choosing the lowest interval defined by i, j , such that `stepHashi` is correct and `stepHashj` is incorrect. In the following round, the prover reveals $n - 1$ new hashes between the previous i, j interval. After l rounds,

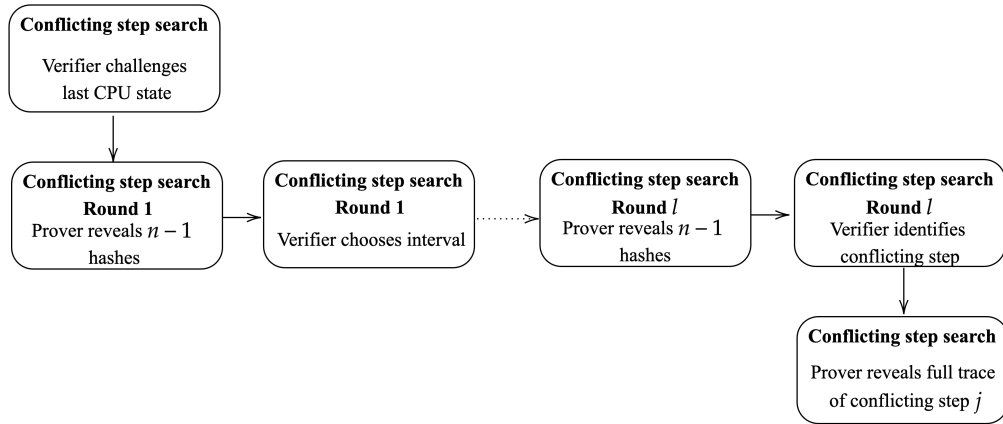


Fig. 3: First stage of the challenge-response protocol used to identify the first conflicting step in the sequence of `stepHash` hashes.

the verifier reaches consecutive values for i, j and the first conflicting step in the chain is `stepHashj`.

Note that this n -ary search process creates a trade-off between its duration and the amount of data published per round. For example, for 2^{32} computational steps, the protocol needs 32 rounds with $n = 2$, 16 rounds with $n = 4$, and 8 rounds with $n = 16$. During the setup stage, the prover and the verifier decide on a maximum number of rounds l for the search process, which limits the maximum number of program steps to $m = n^l$.

Once the verifier identifies the first conflicting step j , the prover is asked to reveal `full_tracej`, which contains all the information of the conflicting step and sets up the next challenge.

5.3 Conflicting step challenge

The second stage of the challenge-response protocol continues after the prover reveals `full_tracej` with the full description of the first conflicting step j . In this stage, the verifier challenges specific data of the conflicting step by choosing one of several possible paths in the protocol. In case the verifier disagrees with multiple parts of the conflicting step, the rational choice is to challenge the part that requires the fewest interactions to verify. In the following, we describe each challenge type in detail.

5.3.1 Trace hash challenge The first information about the conflicting step that the verifier can challenge is the `stepHash` value itself. At this point, the prover has revealed `full_tracej` (including `tracej`), `stepHashj` and `stepHashj-1` for the first conflicting step j . Therefore, the verifier can show that the prover is wrong by validating the following inequality on-chain.

$$\text{stepHash}_j \neq h(\text{stepHash}_{j-1} || \text{trace}_j)$$

5.3.2 Program input challenge During setup, the prover and the verifier agree on a read-only memory range for storing the input of the program. The prover must provide such input in order to validate the computation. The verifier can challenge an input value in either `read1` or `read2` by showing on-chain that `read.addressj` is an input memory address and that `read.valuej` is different from the value provided by the prover for that input.

5.3.3 Read value challenge The verifier can challenge `read.valuej` for either `read1` or `read2` by checking the `read.lastStepj` value provided by the prover. The verifier can easily validate on-chain that $0 \leq \text{read.lastStep}_j \leq j$, so we assume that the `lastStep` value is within boundaries.

The verifier requests different information from the prover depending on the value in `read.lastStepj` to show that `read.valuej` is incorrect. If the prover reveals incorrect data during this challenge, the verifier runs the challenge described in Section 5.3.4. There are three possible scenarios:

1. The `lastStep` value is marked as `INITIAL` but the address has been written to in a previous instruction

2. The `lastStep` value revealed is greater than or equal to the correct value
3. The `lastStep` value revealed is lower than the correct value

In the first scenario, the verifier requests `stepHashk` and `tracek` from the prover, such that

$$\text{write.address}_k = \text{read.address}_j,$$

which proves on-chain that `read.lastStepj ≠ INITIAL`.

In the second scenario, the verifier requests `stepHashk` and `tracek` from the prover, such that $k = \text{read.lastStep}_j$ and shows on-chain that

$$\text{write.address}_k \neq \text{read.address}_j$$

or

$$\text{write.value}_k \neq \text{read.value}_j,$$

which proves that either the instruction k does not write `read.addressj` or that the value written is not the value read.

In the third scenario, the verifier requests `stepHashk` and `tracek` from the prover, such that step k is the actual last step that writes the value. This shows on-chain that

$$\text{write.address}_k = \text{read.address}_j$$

$$k > \text{read.lastStep}_j,$$

which proves that `read.lastStepj` is incorrect.

If the prover reveals an incorrect `stepHashk` in any of the three scenarios, the verifier triggers a new n -ary search to challenge the hash. We describe this process in Section 5.3.4.

5.3.4 Last step hash challenge The verifier triggers a new n -ary search process on `stepHash` if the prover reveals an incorrect `stepHashk` as part of the read value challenge (Section 5.3.3). The goal of the verifier is to find q such that `stepHashq-1` is incorrect and `stepHashq` is correct. This process is equivalent to the search process described in Section 5.2 but with the correct and incorrect hashes reversed. The verifier is guaranteed to find such q because the prover revealed the correct value for `stepHashj-1` during the search for the first conflicting step and now the prover has revealed an incorrect hash in step $k < j - 1$. After identifying q , the verifier forces the prover to reveal `traceq` and shows on-chain that

$$h(\text{stepHash}_{q-1} || \text{trace}_q) \neq \text{stepHash}_q$$

and thus, that `stepHashq-1` is incorrect.

Figure 4 shows a diagram of the read value and last step hash challenges. We omit the rounds of the search process for simplicity.

5.3.5 Program counter challenge The verifier can show an error in the program counter of the first conflicting step, that is `readPC.addressj`, by providing `stepHashj-2` and `tracej-1`, and validating on-chain that

$$h(\text{stepHash}_{j-2} || \text{trace}_{j-1}) = \text{stepHash}_{j-1}$$

$$\text{writePC.address}_{j-1} \neq \text{readPC.address}_j.$$

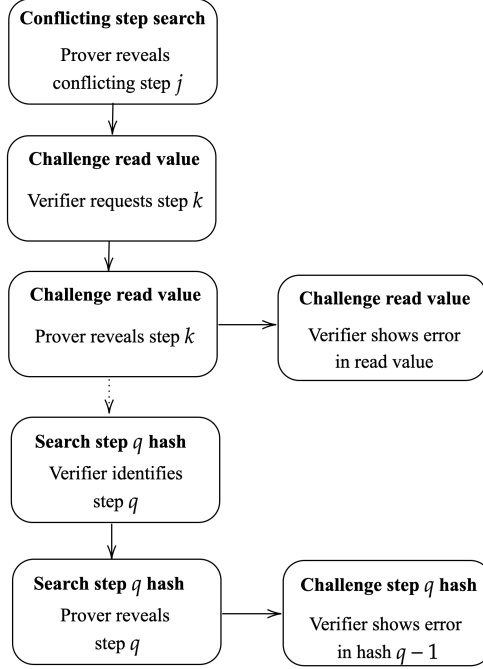


Fig. 4: Read value and last step hash challenges. See Figure 3 for the detailed search process.

The first validation guarantees that the verifier provides correct values for stepHash_{j-2} and trace_{j-1} by checking them against stepHash_{j-1} , which the prover revealed in the first stage of the protocol. The second validation shows that readPC.address_j is incorrect.

5.3.6 Op-code challenge The verifier can challenge the value in readPC.opcode_j by initiating a step search as described in Section 5.3.4 in the negative positions of stepHash . That is, from stepHash_{-p-1} to stepHash_{-1} . Recall that, during setup, the prover and the verifier insert write operations for each instruction in $\text{full_trace}_{(-p-1)\dots-1}$ and agree on stepHash_{-1} . This search process allows the verifier to force the prover to reveal full_trace_{-k} , where $-k$ is the write operation for instruction j . The verifier can then show on-chain that the address where the instruction is written matches readPC.address_j but that the op-code written is different from readPC.opcode_j .

5.3.7 Execution challenge The verifier can challenge the execution of full_trace_j if either write.value_j or writePC.pc_j are incorrect. This is done by executing the instruction on-chain and validating both writes.

6 Conclusion and future work

BitVMX expands the current state of the art in optimistic execution of off-chain programs on Bitcoin – an area of research pioneered by BitVM. In this paper we laid out the primitives of BitVMX, focusing on the single verifier case. Future work can explore ways to extend this to the N verifier setting in a robust manner. Our goal is to create a secure, extensible, open-source, peer-reviewed and sidechain-agnostic framework that can be used to develop blockchain bridges, aggregator oracles, and SNARK/STARK verifiers.

Further research is also needed on the economic incentives for the core protocol and to match the needs of specific use cases. This includes the size of deposits from prover and verifier, bounties for verifiers, the cost of capital when operating a bridge, and the crypto-economic security of the protocol. Applications may have varying priorities with regards to transaction costs or round complexity. Such trade-offs depend on opportunity costs inherent to each application. Future research can explore how BitVMX can be customized to suit the needs of distinct applications.

References

1. Ark, <https://arkdev.info/>
2. BitVM, <https://bitvm.org/>
3. Bitansky, N., Canetti, R., Chiesa, A., Tromer, E.: From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. p. 326–349 (2012)
4. Lamport, L.: Constructing digital signatures from a one way function. Tech. rep. (1979)
5. Linus, R.: BitVM2: Permissionless Verification on Bitcoin, <https://bitvm.org/bitvm2.html>
6. Linus, R.: BitVM : Compute Anything on Bitcoin. Tech. rep., ZeroSync (2023)
7. Merkle, R.C.: A Certified Digital Signature. In: Advances in Cryptology — CRYPTO’ 89 Proceedings. pp. 218–238 (1989)
8. Teutsch, J., Reitwießner, C.: A scalable verification solution for blockchains (2017), <https://arxiv.org/abs/1908.04756v1>
9. Teutsch, J., Reitwießner, C.: A scalable verification solution for blockchains. In: Aspects of Computation and Automata Theory with Applications. pp. 377–424 (2023)