



Formally Verifying the Safety of Pipelined Moonshot Consensus Protocol

M. Praveen  

Chennai Mathematical Institute, India
ReLaX, Chennai, India

Raghavendra Ramesh  

Supra Research, Brisbane, Australia

Isaac Doidge  

Supra Research, Brisbane, Australia

Abstract

Decentralized Finance (DeFi) has emerged as a contemporary competitive as well as complementary to traditional centralized finance systems. As of 23rd January 2024, per Defillama [6] approximately USD 55 billion is the total value locked on the DeFi applications on all blockchains put together.

A Byzantine Fault Tolerant (BFT) State Machine Replication (SMR) protocol, popularly known as the consensus protocol, is the central component of a blockchain. If forks are possible in a consensus protocol, they can be misused to carry out double spending attacks and can be catastrophic given high volumes of finance that are transacted on blockchains. Formal verification of the safety of consensus protocols is the golden standard for guaranteeing that forks are not possible. However, it is considered complex and challenging to do. This is reflected by the fact that not many complex consensus protocols are formally verified except for Tendermint [4] and QBFT [5].

We focus on Supra’s Pipelined Moonshot consensus protocol. Similar to Tendermint’s formal verification, we too model Pipelined Moonshot using IVy and formally prove that for all network sizes, as long as the number of Byzantine validators is less than $1/3$, the protocol does not allow forks, thus proving that Pipelined Moonshot is safe and double spending cannot be done using forks. The IVy model and proof of safety is available on [1].

2012 ACM Subject Classification Networks → Protocol testing and verification; Theory of computation → Logic and verification; Theory of computation → Automated reasoning

Keywords and phrases Blockchain consensus, Safety, Formal verification

Digital Object Identifier 10.4230/OASICS.FMBC.2024.3

Supplementary Material *Model (Source-code)*: <https://github.com/Entropy-Foundation/suprabft-fv/tree/master/suprabft>

Funding *M. Praveen*: Funded by Supra

Acknowledgements We acknowledge and thank Chandradeep Dey and Namrata Reddy, who were part of this project during its initial phase. We acknowledge Supra Research for funding M. Praveen, the academic partner of this project and providing other support.

1 Introduction

Public blockchains are revolutionising modern society by rebranding traditional services mainly the traditional finance based services, and offering them on a “decentralized trust” platform. Here no single entity need be trusted as the network is typically open for permissionless participation and tolerates malicious behaviour of the participants up to a certain threshold. Though blockchains are being adopted by multiple domains of applications, finance or Decentralised Finance (DeFi), happens to be the *killer* application that has shot the



© Supra. This work, as part of the collaborative efforts of M. Praveen, Raghavendra Ramesh and Isaac Doidge, falls under the intellectual property rights assigned to Supra in accordance with their agreements.;

licensed under Creative Commons License CC-BY 4.0

5th International Workshop on Formal Methods for Blockchains (FMBC 2024).

Editors: Bruno Bernardo and Diego Marmosoler; Article No. 3; pp. 3:1–3:16



OpenAccess Series in Informatics

OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

blockchain technology to fame as well as towards a popular adoption. As of 23rd January 2024, per DeFillama [6] approximately USD 55 billion is the total value locked on the DeFi applications on all blockchains put together.

Every node in a blockchain network runs a *consensus protocol*, more precisely known as *state machine replication (SMR)* protocol, that enables that node to transition from one blockchain state to the next in a consistent way so that no two nodes in the network end up in different states after processing the same sequence of transitions. The transitions are the clients' submitted ledger transactions that are batched into a block. The sequence of transitions form the chain of blocks, hence the name blockchain. We are interested in consensus protocols in a *partially synchronous network* setting. In this setting it is well known that an SMR protocol tolerates up to one-third of the network nodes being Byzantine – nodes that may crash or deviate arbitrarily from the protocol but are assumed to be unable to break cryptographic primitives like signatures.

Many such protocols have been proposed as well as successfully been adopted in practice such as [10, 11, 27, 15, 20] but only a few protocols have been formally verified to the best of our knowledge: Both the safety and liveness of Tendermint [4] have been formally verified using Microsoft IVy [22]. The safety of QBFT (called IBFT earlier) has been verified [5] using Microsoft's Dafny [18].

Pipelined Moonshot [12] is a novel rotating leader-based Byzantine fault tolerant SMR protocol that leverages *optimistic proposals* to achieve a high block throughput – one block per network hop, and the lowest block finalization latency – 3 network hops, in the scenario of a normal path. It is well known that designing protocols and proving them correct by hand are notoriously prone to errors as many critical errors have been found even in peer-reviewed distributed protocols (see [25] and references therein). In this paper, we focus on formally proving the safety of this protocol.

Safety of a BFT SMR protocol is a critical requirement, ensuring that any two honest processes agree on the set of transactions executed and the order in which they are executed. Formally, if two honest processes have committed chains of blocks, then one of the chains must be (not necessarily strict) prefix of the other one. When a protocol loses safety, forks in the blockchain are possible, essentially yielding to the possibility of *double spending* which is catastrophic to the finances built on top of this blockchain. Pipelined Moonshot protocol is proved to be safe and live [12] with a handwritten proof. The goal of this project is to provide a proof of safety in a formal verification tool.

Our Contributions.

- We provide formal specification of the Pipelined Moonshot protocol in IVy [22], serving as a reference for any implementation.
- We formally verify safety of Pipelined Moonshot successfully. This makes the formal specification a safety-error-free basis for any implementations to be developed.
- We identify several invariants of the protocol to prove it safe. Invariants are useful in generating test cases to test implementations of the protocol [30, 28].
- We record our experience in the form of challenges faced and the corresponding mitigations used. Learning from this experience we extend our wisdom as recommendations for applying formal verification to large projects.

2 Related Work

In this section we detail various formal verification approaches for consensus protocol verification and motivate our choice of IVy for verifying the safety of Pipelined Moonshot protocol.

Model checking approach typically models the protocol as some finite representation of a state transition system and expresses the correctness properties in some logic and enumerates exhaustively the state space validating against the given logical specifications. Various model checking tools like SPIN [2], TLC [3], Apalache [16] etc are popular. Typically for the consensus protocols of interest as long as the number of nodes in the system is not fixed the state space is unbounded and generally does not yield a decidable algorithm to model check. There are bounded model checking approaches where the number of nodes is fixed typically and that yields a finite state machine against which the correctness properties are checked. For instance, [9] model checks the block synchronization protocol of Tendermint after fixing the number of nodes in the network using TLC and Apalache model checker. We are focused on the general problem of safety of the Pipelined Moonshot with no bounds on the network, hence model checking is not applicable and bounded model checking is not satisfactory.

There are other approaches that identify a *small model property* in the given protocol verification problem and apply model checking against the small model. The small model property of a protocol P essentially is a bound on the number of nodes, say k , such that the satisfaction of a property θ by P when run with k nodes imply that P satisfies θ for all $n \geq k$. The threshold automata approach of [17] leverages this and builds a counter abstraction by counting the number of processes in each state. This has been applied to the verification of DBFT [25] asynchronous consensus protocol. However this approach is known to be hard and has not been applied so far for any of the partially synchronous BFT consensus protocols, and we too could not find any direct ways of applying this approach to the safety verification of Pipelined Moonshot.

We now turn to the deductive verification tools. In this approach, the protocol is modeled in some logic (typically first-order logic, its fragments or extensions) and the properties to be verified are also written in the same logic. From these, formulas called Verification Conditions (VCs) are generated, whose unsatisfiability implies that the protocol has the desired property. With interactive theorem provers, proof of unsatisfiability is developed in a proof system (such as natural deduction system or its variants). With automated deductive verification tools, proof of unsatisfiability is given by Satisfiability Modulo Theory solver (SMT solver).

TLA+ [3] supports a very expressive logic called TLA – Temporal Logic of Actions, for specifying state machines and properties. We found that expressing the Pipelined Moonshot protocol in TLA+ to be very complex and huge, and so also the verification in TLAPS – TLA Proof System, to be effortful as each and every lemma has to be proven more or less interactively. We were on the look out for solvers that push more automation and lessen the interaction with the solver. Another requirement of us was that the formal specification should be close to the real world programming languages so that the developer community may be comfortable using the formal specification as the basis for their implementation. We found the TLA+ specifications to be far from the interest of the developer community, unless they are trained specifically towards verification.

Dafny [18] is a verification-aware programming language that facilitates specifying pre and post conditions for procedures and verify at compile time. Correctness by construction is the philosophy here. The code may also be compiled to regular programming languages like C#, Java, JavaScript, Go and Python. The safety of QBFT (previously known as IBFT [20]) has been verified using Dafny.

It is also well known that the formal verification of distributed protocols is an arduous effort. Hence we were on the look out for a tool that maximally uses automation in proof building and we found **IVy** [22] to fit the bill. IVy is a language and a tool for the formal

specification and verification of distributed systems. IVy supports deductive verification using automated provers such as Z3 [8], model checking, automated testing, manual theorem proving and generation of executable code. In order to achieve greater verification productivity, a key design goal for IVy is to allow the engineer to apply automated provers in the realm in which their performance is relatively predictable, stable and transparent. In particular IVy focuses on the use of decidable fragments of first-order logic. IVy supports modularisation of the specifications and proofs, aiding their readability and also ensuring that formulas passed to provers are in decidable fragments. This helps to some extent in getting the provers to return with answers quickly.

As IVy embodies an imperative language, the protocol specification in IVy serves as a sound reference for any implementation. Note that the safety of Tendermint has been verified using IVy [4]. For all these reasons we favoured IVy as the formal verification tool for verifying the safety of Pipelined Moonshot.

To make proofs easier, Pretend Synchrony [26] takes another route of reducing the problem of verifying asynchronous distributed protocols to the problem of verifying synchronous distributed protocols. However it has been applied only in the setting of crash faults setting but not in Byzantine faults setting, which is the focus of this paper.

3 Safety of Pipelined Moonshot Consensus

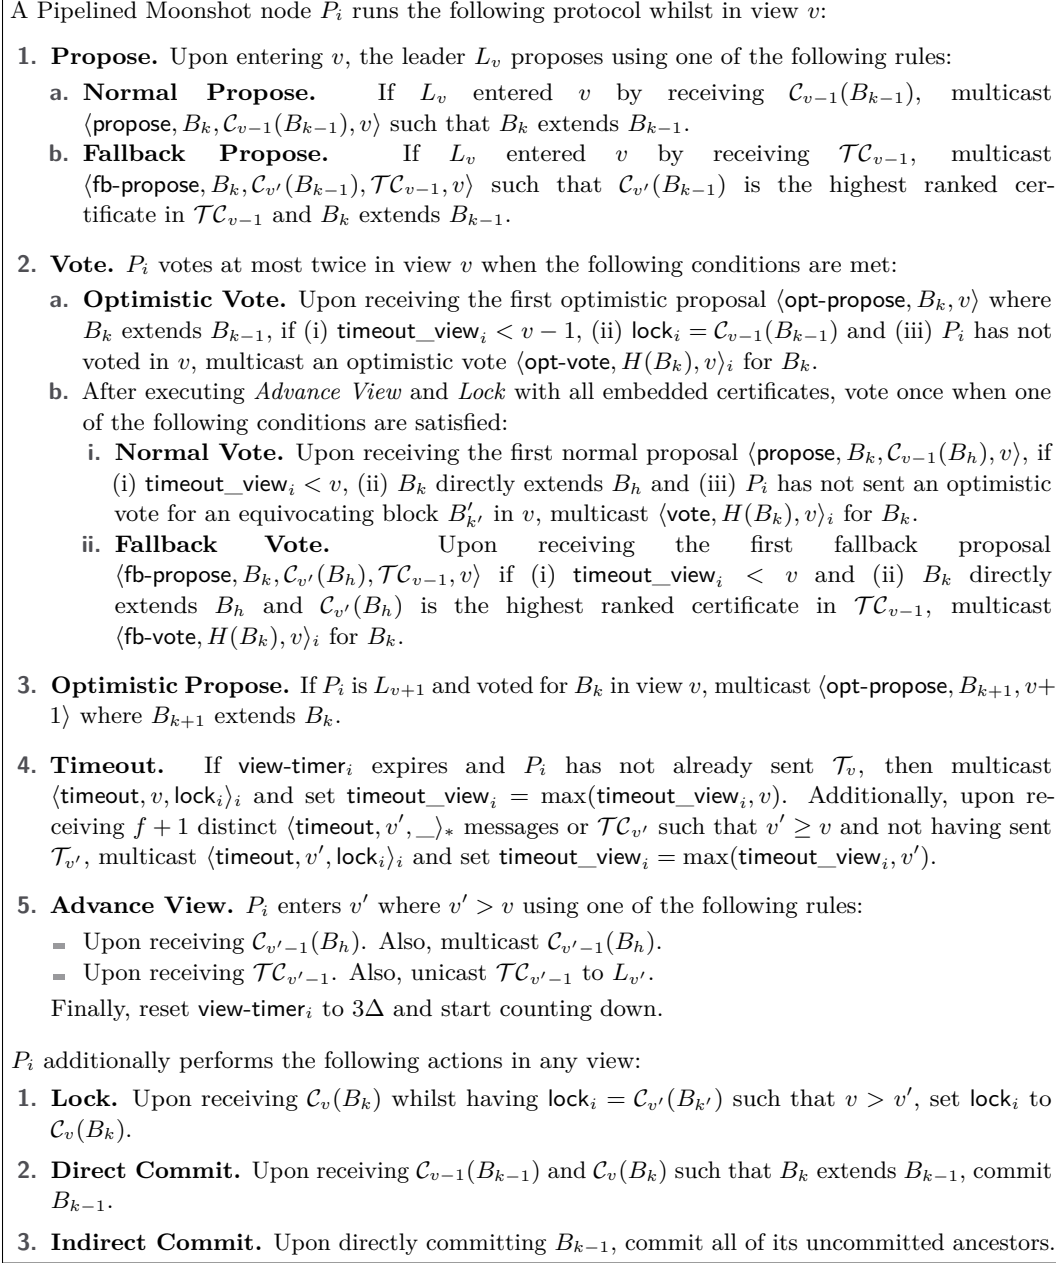
In this section we summarise Pipelined Moonshot and elucidate the scope of our formal verification endeavour.

Pipelined Moonshot

Pipelined Moonshot [12] is a chain-based, rotating leader Byzantine Fault Tolerant (BFT) State Machine Replication (SMR) protocol optimized for wide-area networks. It satisfies the safety and liveness properties of SMR under the partially synchronous network model [13] given at most f of the n total participants in the protocol (which we will call *validators*) are Byzantine, such that $f < \frac{n}{3}$. Without loss of generality, we assume that $n = 3f + 1$ for the rest of the paper and use the term *quorum* to refer to a set of $2f + 1$ validators. The full details of Pipelined Moonshot's setting are provided in [12].

The protocol, presented in Figure 1 as given in [12], constructs a chain of blocks of client transactions (or some abstraction thereof) over a sequence of numbered views advanced by quorum decisions in the form of certificates. In Pipelined Moonshot, a view, say v , may produce two types of certificates; a Quorum Certificate $C_v(B_k)$ comprised of $2f + 1$ votes for some block B_k (where k is the position or *height* of B in the blockchain) proposed for v , or a Timeout Certificate \mathcal{TC}_v comprised of $2f + 1$ timeout messages for v . A Pipelined Moonshot validator in view v votes for B_k when it receives B_k in a valid proposal (described momentarily) and sends a timeout message for v , denoted \mathcal{T}_v , that contains its locked QC – i.e., the QC with the highest view that it has observed so far – when it fails to exit v before its view timer expires or when it observes evidence that at least one honest validator has already sent \mathcal{T}_v . These rules together ensure that the protocol continually generates new certificates, preventing it from halting.

A validator that receives a certificate for view v advances its local view to $v + 1$ and resets its view timer. If it enters $v + 1$ via a QC then it also locks the QC and multicasts it to ensure that its peers enter the view promptly. Otherwise, if it enters $v + 1$ via \mathcal{TC}_v then it unicasts this certificate to the designated leader for $v + 1$, denoted L_{v+1} , enabling it to enter the view and propose promptly.



■ **Figure 1** The Pipelined Moonshot Protocol [12].

Upon entering $v+1$, L_{v+1} creates a new block, say B_l , and multicasts it in a proposal that depends on the type of certificate it used to enter the view. If the view change was triggered by $\mathcal{C}_v(B_k)$, then B_l *directly extends* B_k (i.e. $l = k + 1$ and B_l contains the hash digest of B_k) and L_{v+1} multicasts a Normal Proposal containing both B_l and $\mathcal{C}_v(B)$. Otherwise, B_l extends the block certified by the QC with the highest view included in \mathcal{TC}_v and L_{v+1} multicasts a Fallback Proposal containing both B_l and \mathcal{TC}_v . A validator in $v+1$ that receives a proposal of either type from L_{v+1} that is constructed as previously described and has yet to either send a vote for an equivocating block (as described in Figure 1) or a timeout

message for $v + 1$, multicasts a vote of the corresponding type for B_l for $v + 1$. Importantly, Pipelined Moonshot ensures that votes cannot be aggregated into a QC unless they have the same type.

Upon voting for B_l , if the validator is L_{v+2} then it also creates an Optimistic Proposal for $v + 2$ containing a new block that extends B_l , presuming that B_l will be certified. A validator that receives this proposal votes for it once in $v + 2$ if it has not yet voted in $v + 2$ and it entered the view via $\mathcal{C}_{v+1}(B_l)$ without having sent a timeout message for $v + 1$. Optimistic Proposals, a distinguishing feature of Moonshot protocols, allow votes for the current view to be disseminated in parallel with a proposal for the next view when both leaders are honest. Comparatively, prior protocols require a leader to receive a certificate for the previous view before proposing, inherently sequentializing these actions.

A validator that observes the certification of a block and its immediate successor in the chain for adjacent views commits the block by permanently appending it to its local copy of the blockchain.

Safety

An SMR protocol is *safe* if it ensures that no two validators commit divergent blockchains. Let the local blockchain of validator P_i be denoted by \mathbf{B}_i . More formally, the safety property states that for every run of the protocol, and for each pair of honest validators $(P_i, P_j) \in \mathcal{V} \times \mathcal{V}$, either \mathbf{B}_i is a (not necessarily strict) prefix of \mathbf{B}_j or vice-versa.

The Pipelined Moonshot [12] paper contains the handwritten proof of safety and liveness. As well established in the literature some errors may potentially be present in the handwritten proofs that could go overlooked. A recent example is the Chord [24] protocol for distributed hash tables which, despite having more than 6000 citations, was shown incorrect by Zave [29] almost a decade after its publication. Since only formal verification can conclusively guarantee the absence of errors in a protocol, we aimed at developing mechanically verifiable proofs of the safety of Pipelined Moonshot.

4 Formal Specification and Verification using IVy

In this section, we first present some of the preliminaries of IVy modeling, secondly we present an high level overview of the formal IVy specification of the Pipelined Moonshot consensus protocol, and then finally present the structure of our safety proof.

4.1 IVy modeling setup

IVy is a language and a tool for the formal specification and verification of distributed systems. Systems are represented as state transition machines. States are multi-sorted first-order structures, with relations and functions. Transitions specify how the state is mutated. Any update definable in first-order logic is supported. Update instructions can be given in sequence one after another, giving the syntax the flavor of a developer-friendly imperative programming language. Multiple update instructions can be grouped together into an *action*, a keyword in IVy that is used to denote state transition specifications.

The system under consideration is typically split into multiple modules, with internal states of a module not allowed to be modified directly by other modules. One module can call actions of another module, passing parameters. Modules can reason about one another using *assume-guarantee* specifications, which are formulas specifying properties of the modules' states. Properties of the overall system has to be proved by writing *inductive invariants*,

which are properties satisfying two conditions – initiation and inductiveness. Initiation means that the initial state of the system satisfies the invariant. Inductiveness means that if any of the actions are executed in any state that satisfies the invariant, the resulting state also satisfies the invariant.

Modules in IVy, apart from modularising the protocol specification and proofs, serves another deeper purpose. Multiple formulas used in the proof may together necessitate the use of logics that are undecidable. Modules in IVy allow proving different properties in isolation, ensuring that formulas supporting one invariant are invisible to other modules. This will allow users to control which formulas are passed to the underlying SMT solvers together, so that all calls to the SMT solver are within decidable fragments of first-order logic.

Only a high level abstract specification of the protocol is modeled and verified. Some implementation details are hence modeled with Boolean abstractions. For example, timers used in the protocol are replaced by Boolean propositions that indicate whether or not a timer has expired. In the IVy model, the Boolean proposition can switch value anytime non-deterministically to simulate a timer getting expired, instead of tracking the actual time elapsed since the last reset. This is a sound abstraction for proving safety.

Another abstraction we have adapted from the literature is handling quorums [19]. The protocol specification mandates that a validator needs to receive messages from a super majority of all validators (two-thirds of the entire set) in order to achieve a quorum. Verifying this detail would require having arithmetic in the formulas passed on to SMT solvers, potentially affecting the solver’s performance. Instead, what is modeled is the *quorum intersection property* [19] – any two quorums have at least one common honest validator. It is this property of quorums that are mainly used in correctness proofs and is modeled in IVy as an axiom, avoiding the usage of arithmetic.

Validators receive messages from the network and verify their authenticity by checking digital signatures. It is assumed that Byzantine validators cannot break cryptographic primitives and hence they cannot forge signatures of honest validators. Checking digital signatures is not modeled in IVy – the model assumes messages sent by honest validators are authentic. The model also disallows byzantine validators to send messages on behalf of other honest validator, though they can send any kind of message on behalf of themselves or other byzantine validators, even if such a message is not mandated to be sent by the protocol specification.

4.2 Pipelined Moonshot Specification

We have published our IVy specification and the formal proof of safety of Pipelined Moonshot online on GitHub [7]. Following are the main modules in our IVy specification of Pipelined Moonshot:

Types. This module contains the declarations of the data types used in the IVy specification. The types *round_t*, *height_t* are declared to be instances of *ubd_seq*, a small modification of *unbounded_sequence*, which are finite but unbounded total linear orders. Round is the technical term used in our IVy model for view as used in the protocol specification [12]. The type *process_index_t* is declared to be an instance of *iterable*, which allows a collection of validators to be iterated in a loop in IVy models. The above types are used conventionally while modeling protocols in IVy. Other types declared correspond to message types specified in the protocol specification: *block_t*, *quorum_t*, *qc_t*, *tc_t*. Common properties of these types are also written in this module, including the quorum intersection axiom.

3:8 Formally Verifying the Safety of Pipelined Moonshot Consensus Protocol

Network. This module models the network through which validators interact. It is almost same as the network model in Tendermint’s IVy model [4], except for the kind of messages that can be sent. Here the kind of messages that can be sent are *normal proposal*, *fallback proposal*, *optimistic proposal*, *normal prepare*, *fallback prepare*, *optimistic prepare*, *quorum certificate*, *timeout certificate*, *timeout* and *weak timeout certificate*. A *timeout certificate* is a collection of *timeout* messages from a two thirds majority of validators, whereas a *weak timeout certificate* is a collection *timeout* messages from a number of validators at least one more than the number of Byzantine validators. The network model is that of any asynchronous one, where messages can be dropped or delivered multiple times and/or out of order.

Moonshot. The IVy specification of the Pipelined Moonshot is provided in this module. State variables of individual validators are declared and updates to the state variables are performed in response to specific events as specified in Figure 1. The details of this module are provided in Appendix A.

Quorum verification. In implementation, the integrity of a quorum of messages received by a validator is verified by checking digital signatures accompanying the messages. Here, the integrity is checked by verifying that all honest members of a quorum have actually sent the corresponding messages. It is done in this module using the concept of *monitors* provided by IVy – they are additional updates to state variables that are performed whenever an action is performed by the protocol. This module contains monitors that record prepare and timeout messages sent by the validators. When a validator receives a quorum or timeout certificate, its integrity is checked by verifying from the records that all honest members of the quorum have actually sent the corresponding prepare or timeout message. This can be thought of as some kind of a central authority with a global view of all validators, who records all messages sent by the validators. Of course there is no such central authority in real implementation; it is only modeled here for the sake of proving safety.

Safety. The safety module specifies the desired safety property in the form of an inductive invariant. Numerous supporting invariants are also included here, as detailed in the next sub-section. Following is a code snippet, that states the main safety property.

```
isolate full_safety = {
  relation blockchain_prefix(N1:process_index_t, N2:process_index_t)

  # the latest block committed to b_v by N1 is equal to or ancestor of the
  latest block committed by N2. All blocks committed by N1 are also
  committed by N2. Any block committed by N2 but not by N1 is a descendant
  of the latest block committed by N1
  definition blockchain_prefix(N1,N2) = ...

  # this is the full safety property of the pipelined moonshot protocol: for
  any two honest processors N1,N2 the chain committed by N1 is a prefix of
  N2 or vice versa
  invariant forall N1,N2:process_index_t. is_good(N1) & is_good(N2) ->
  (blockchain_prefix(N1,N2) | blockchain_prefix(N2,N1))
} with block_t, verify_quorum, certified_block_ancestor_m1,
all_ancestors_committed, committed_blocks_ancestors,
latest_committed_ancestors, commit_to_chain, commit_to_chain_m1
```


The definition of the relation `blockchain_prefix` above is not shown fully due to lack of space, but its intention is captured in the comment above. The *with* clause above lists the names of other isolates, containing invariants supporting this one.

Table 1 provides some statistics of these modules. Note that a typical line in `safety.ivy` is much longer than those in other files, since one whole invariant is written in one line of `safety.ivy`. The files also have extensive comments serving the purpose of readability and documentation. There are a total of 190 invariants and 23 monitors. A rough estimate of the ratio of sizes of program code vs. proof is 1:3. Verifying the safety of Pipelined Moonshot took about 140 man hours after the protocol specification itself was stabilized. About 10% of this was needed to model the protocol and the rest to complete the proofs.

■ **Table 1** Modules and their Lines of code.

Module	Contents	Lines
Types	Extended data types	338
Network	Network model	110
Moonshot	Pipelined Moonshot SMR protocol	642
Quorum verification	Validating messages sent by quorum members	165
Safety	Inductive invariants proving safety	1309
	Total	2564

4.3 Structure of the Safety Proof


Mechanically-checked proofs are developed interactively in a dialogue between a Verification engineer and the proof assistant – IVy. The engineer gives the desired specification of the model and of the property, IVy attempts to prove that the model satisfies the property. It may prove, then all is well, it may fail showing a counterexample, or it may not come back for a reasonable amount of time. When satisfiability fails it shows logical errors in the protocol. When it takes an unreasonable amount of time, the engineer has to creatively craft some lemmas that aids the machine in its proof search. This is the standard iterative approach of building mechanised proof.

IVy could not prove the safety specification directly (as is typical). We had to write all the intermediate lemmas given in [12] and many more. We first outline the main steps in the handwritten safety proof.

1. If an honest validator executes direct commit of a block B as given in point 2 at the bottom of Figure 1, then any subsequent block that achieves a quorum is a descendant of B . This is proved in [12, Lemma 2, Lemma 3].
2. If an honest validator commits a block, it also commits all of its ancestors. This is implicit in [12].
3. For any two blocks committed by an honest validator, one is an ancestor of the other. This does not directly correspond to any result stated in [12], but essential in our IVy proof.
4. If B_i (resp. B_j) is the latest block committed by an honest validator v_i (resp. v_j), then B_i is an ancestor of B_j or vice-versa. This is a corollary of item 1 above.
5. If there were two blocks that were divergent, one would be an ancestor of the other (by item 3 above) and both would be ancestors of the latest committed block (by item 4 above). Hence, both would be committed by all honest validators (by item 2 above), contradicting the hypothesis that they are divergent. This argument is essentially the proof of [12, Theorem 3].

3:10 Formally Verifying the Safety of Pipelined Moonshot Consensus Protocol

IVy verifies that properties are inductive invariants by generating formulas in Finite Almost Uninterpreted (FAU) fragment of first-order logic and passing them on to Z3 [19]. Trying to prove too many properties in one step often degrades the performance of the SMT solver. To overcome this, IVy allows to group together a small number of properties in an *isolate*, specifying other isolates as supporting invariants. When verifying one isolate, other isolates that it depends on are assumed to be true. The dependencies can be checked later. The safety invariants in our model are structured into several isolates. The top level of this structure follows the structure of the handwritten proof that is summarized above, and is illustrated below. Here, (5, `full_safety`) means that the point 5 above is proved in the isolate `full_safety`, likewise for other nodes. The arrow from (5, `full_safety`) to (3, `committed_blocks_ancestors`) means that the isolate `full_safety` depends on other isolates: `committed_blocks_ancestors` being one of them.



The isolate `quorum_after_ldc_descendant` is technically the most involved result in both the handwritten proof and IVy proof. This result is proved in IVy by induction on rounds. The principle of induction is taken to be an axiom and applied to the main invariant in the isolate `quorum_after_ldc`. It states that if a block B is committed directly by an honest validator, then any block proposed in later rounds that achieves a quorum has a parent proposed in the same round as B or later rounds. Proving the invariants in the isolate `quorum_after_ldc` itself is lengthy, indirectly involving around 30 other isolates.

5 Challenges

The development and handwritten proof of safety and liveness of the pipelined Moonshot protocol underwent many cycles (some modifications to ensure liveness and some for simplifying the specification and proofs). This naturally resulted in iterating and refining the IVy specification too. The process of formally verifying safety (including analyzing counter examples given by IVy) uncovered some points in the specifications and proofs that were ambiguous and helped better understand many details that were implicit in the handwritten proofs.

Here we document some of the challenges faced in such a verification effort.

Transitive closure

Ancestor relation is the binary transitive closure of the parent relation. For a validator to commit a block to its canonical chain, the block and its ancestors must have got quorums. The statement and proof of the property in the isolate `quorum_after_ldc_descendant` uses the ancestor relation. Thus, many crucial parts of the model and safety proof depend on the ancestor relation. However, transitive closure of binary relations are not definable in first-order logic. To overcome this, we adapted a known technique [19]. If a binary relation is known to be the transitive closure of another base relation, then under some conditions the base relation can be defined from its transitive closure in first-order logic. To use this

technique, a monitor in the isolate `certified_block_ancestor_m1` tracks when blocks get quorums and records the ancestor relation among them. Whenever a block gets quorum, the monitor updates the record, making the newly certified block a descendant of its parent block and all the parent block's ancestors. In the isolate `certified_block_ancestor_m5`, we verify that the base relation obtained from the relation recorded by `certified_block_ancestor_m1` is indeed the parent relation. This challenge is not there for verifying Tendermint [4], where cross dependencies between isolates is lesser.

Nested subroutine calls

As in most programming languages, actions in IVy can invoke other actions, which can themselves invoke more actions and so on. We observed empirically that with higher depth of nesting of action invocations, the performance of the IVy verifier slows down considerably. The protocol specification use subroutines that are called from multiple sites and the natural way to model it would be to have similar actions in IVy called from multiple actions. However, the slowdown in performance was significant enough that we resorted to inlining the subroutine calls. Further work is needed to understand the causes and more elegant workarounds.

IVy verifier getting stuck without giving an answer

This challenge took up most of the time for executing this project. While verifying properties that were expected to be true, the IVy verifier would call the Z3 SMT solver, which would run for a long time without giving any answer. Such behaviour from SMT solvers cannot be entirely avoided, since they try to solve problems that have quite bad complexity theoretical lower bounds. There is no fixed template for handling this. Experience with using the tool and familiarity with the protocol being verified help a little bit. This is a challenge faced by most formal verification efforts; we felt it more since we had many invariants to prove, due to the complexity of the underlying protocol. Here are a few rules of thumb we resorted to, devised from trial and error.

Isolating the cause in the protocol. If the property being verified involved multiple steps in the protocol, we tried commenting out parts of the protocol and trying to verify the property. If the property was proved to be true/false after a particular section was commented out, then we could concentrate on that part to see what can be causing the SMT solver to diverge. This strategy helped us identify some subtle points that were implicitly assumed in the handwritten proof.

Explicitly writing intermediate results. We illustrate this with an example. In the isolate `quorum_after_ldc_descendant_m7`, the third invariant states that under some conditions, the block Bp is an ancestor of B . IVy could not prove this in a reasonable amount of time, so it was not clear whether it is due to lack of supporting invariants or because the SMT solver is diverging. We then added the first two invariants. The first one says that under the same condition, Bp is an ancestor of $Bp1$ and the second one says that additionally, $Bp1$ is an ancestor of B . With the first two invariants added, IVy successfully verifies all the three in short time. There are many more examples like this. The main invariant of `quorum_after_ldc_descendant_m7` is inferred from a similar series of intermediate invariants, starting from `quorum_after_ldc_descendant_m1`, ending at `quorum_after_ldc_descendant_m8` and then finally proving the invariant in `quorum_after_ldc_descendant`.

6 Recommendations

In this section we consolidate our experience with the safety verification of Pipelined Moonshot and attempt to distill some recommendations for applying formal verification techniques in proving the correctness of distributed systems.

- Compared to semi-interactive theorem provers like Coq, the manual effort required with IVy is lesser. The proofs had to be flattened out into small steps manageable by SMT solvers, as explained in the last challenge. More research efforts like [14] are needed to reduce the burden of manually working out minute details, letting users concentrate on understanding protocols and correctness proofs intuitively.
- With complex protocols involving correctness proofs using hundreds of invariants, the success of deductive verifications tools that call SMT solvers in the background depends crucially on modularization of the proof, so that every SMT call is restricted to a small number of closely related formulas. For this, it is important at the outset to have a good idea of how the modules are going to be structured and which modules are meant for what. If this is lacking during the initial phase, chasing minute details during the verification process can quickly lead to huge, monolithic, incomprehensible and unmanageable pile of candidate invariants. In earlier attempts at verifying Pipelined Moonshot, we were sometimes in situations where we changed an invariant written earlier to suitably support a newly written invariant, only to realize later that this change affected an earlier dependency. We had lost track of which invariants supported which others and small changes in one invariant affected seemingly unrelated ones elsewhere.
- A related point is to be disciplined while establishing inter-dependencies among modules, specifically isolates in IVy. If invariant 1 in isolate 1 needs invariant 2 in isolate 2 for support, it is tempting to mention the whole of isolate 2 as a dependency for isolate 1, instead of mentioning just invariant 2 of isolate 2. This may seem to be a time saver in the short term, but will result in unnecessary formulas (invariants in isolate 2 different from invariant 2) being passed to the SMT solver. Such unnecessary formulas can drastically degrade the performance of SMT solvers. Due to this, isolate 1 may pass all verification conditions in a short time currently but may not be able to do so in the future if additional invariants are added to isolate 2. If there is a group of invariants that are always together supporting other invariants, they should be recognized as such and grouped into an isolate, which is why this is related to the previous point of starting with well organized modules.
- The state of the art for formal verification of this scale very much requires experts with advanced knowledge of logic and related topics, who also need to understand the protocol being verified. An ideal team for formal verification would consist of experts with experience in using logic based verification tools on the one hand, and designers who understand the workings of the protocol at both abstract level and minute detail level on the other hand.

7 Conclusion

We have successfully verified the safety of a high performance and complex consensus protocol, namely Pipelined Moonshot, in IVy. This conclusively proves the absence of design or logic errors with respect to the protocol safety. Proving liveness is future work, possibly using [21] to reduce liveness to safety.

This effort has yielded a developer friendly formal specification of the Pipelined Moonshot protocol that helps for any implementation of Pipelined Moonshot to safely base on.

We recorded our experience in the form of challenges faced and the mitigations employed during this project. Learning from this experience we enumerate some recommendations for applying formal verification for large distributed protocols.

References

- 1 IVy modeling of Pipelined Moonshot and its proof of safety. <https://github.com/Entropy-Foundation/suprabft-fv/tree/master/suprabft>.
- 2 SPIN. <https://spinroot.com/spin/whatispin.html>.
- 3 TLA+. <https://lamport.azurewebsites.net/tla/tla.html>.
- 4 Defillama. <https://galois.com/blog/2021/07/formally-verifying-the-tendermint-blockchain-protocol/>, 2021.
- 5 Formal Verification of QBFT Safety. <https://github.com/Consensys/qbft-formal-spec-and-verification>, 2021.
- 6 Defillama. <https://defillama.com>, 2024.
- 7 Moonshot Formal Verification in IVy - GitHub Repository. <https://github.com/Entropy-Foundation/suprabft-fv/tree/master/suprabft>, 2024.
- 8 Z3 SMT Solver. <https://www.microsoft.com/en-us/research/project/z3-3/>, 2024.
- 9 Sean Braithwaite, Ethan Buchman, Igor Konnov, Zarko Milosevic, Iliana Stoilkovska, Josef Widder, and Anca Zamfir. Formal Specification and Model Checking of the Tendermint Blockchain Synchronization Protocol. In Bruno Bernardo and Diego Marmosoler, editors, *2nd Workshop on Formal Methods for Blockchains (FMBC 2020)*, volume 84 of *Open Access Series in Informatics (OASISs)*, pages 10:1–10:8, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/OASISs.FMBC.2020.10.
- 10 Ethan Buchman. *Tendermint: Byzantine fault tolerance in the age of blockchains*. PhD thesis, University of Guelph, 2016.
- 11 Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999. URL: <https://dl.acm.org/citation.cfm?id=296824>.
- 12 Isaac Doidge, Raghavendra Ramesh, Nibesh Shrestha, and Joshua Tobkin. Moonshot: Optimizing chain-based rotating leader bft via optimistic proposals, 2024. doi:10.48550/arXiv.2401.01791.
- 13 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *J. ACM*, 35(2):288–323, apr 1988. doi:10.1145/42282.42283.
- 14 Yotam MY Feldman, James R Wilcox, Sharon Shoham, and Mooly Sagiv. Inferring inductive invariants from phase structures. In *Computer Aided Verification: 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part II 31*, pages 405–425. Springer, 2019. doi:10.1007/978-3-030-25543-5_23.
- 15 Rati Gelashvili, Lefteris Kokoris-Kogias, Alberto Sonnino, Alexander Spiegelman, and Zhuolun Xiang. Jolteon and ditto: Network-adaptive efficient consensus with asynchronous fallback. In *FC*, pages 296–315, 2022. doi:10.1007/978-3-031-18283-9_14.
- 16 Igor Konnov, Jure Kukovec, and Thanh-Hai Tran. Tla+ model checking made symbolic. *Proc. ACM Program. Lang.*, 3(OOPSLA), oct 2019. doi:10.1145/3360549.
- 17 Igor Konnov, Marijana Lazić, Helmut Veith, and Josef Widder. A short counterexample property for safety and liveness verification of fault-tolerant distributed algorithms. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL '17*, pages 719–734, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3009837.3009860.
- 18 K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In Edmund M. Clarke and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning*, pages 348–370, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. doi:10.1007/978-3-642-17511-4_20.

- 19 Kenneth L. McMillan and Oded Padon. Deductive verification in decidable fragments with ivy. In Andreas Podelski, editor, *Static Analysis*, pages 43–55, Cham, 2018. Springer International Publishing. doi:10.1007/978-3-319-99725-4_4.
- 20 Henrique Moniz. The istanbul bft consensus algorithm, 2020. doi:10.48550/arXiv.2002.03613.
- 21 Oded Padon, Jochen Hoenicke, Giuliano Losa, Andreas Podelski, Mooly Sagiv, and Sharon Shoham. Reducing liveness to safety in first-order logic. *Proc. ACM Program. Lang.*, 2(POPL), 2017. doi:10.1145/3158114.
- 22 Oded Padon, Kenneth L. McMillan, Aurojit Panda, Mooly Sagiv, and Sharon Shoham. Ivy: safety verification by interactive generalization. *SIGPLAN Not.*, 51(6):614–630, jun 2016. doi:10.1145/2980983.2908118.
- 23 Supra Research. Moonshot: Optimistic proposal for blockchain-based state machine replication. URL: <https://supraoracles.com/news/moonshot-consensus/>.
- 24 Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '01, pages 149–160, New York, NY, USA, 2001. Association for Computing Machinery. doi:10.1145/383059.383071.
- 25 Pierre Tholoniati and Vincent Gramoli. *Formal Verification of Blockchain Byzantine Fault Tolerance*, pages 389–412. Springer International Publishing, Cham, 2022. doi:10.1007/978-3-031-07535-3_12.
- 26 Klaus v. Gleissenthall, Rami Gökhan Kıcı, Alexander Bakst, Deian Stefan, and Ranjit Jhala. Pretend synchrony: synchronous verification of asynchronous distributed programs. *Proc. ACM Program. Lang.*, 3(POPL), jan 2019. doi:10.1145/3290372.
- 27 Maofan Yin, Dahlia Malkhi, Michael K Reiter, Guy Golan Gueta, and Ittai Abraham. Hotstuff: Bft consensus with linearity and responsiveness. In *PODC*, pages 347–356, 2019. doi:10.1145/3293611.3331591.
- 28 Yuan Yuan, Zeng Fanping, Zhu Guanmiao, Deng Chaoqiang, and Xiong Neng. Test case generation based on program invariant and adaptive random algorithm. In *Advances in Information Technology and Education: International Conference, CSE 2011, Qingdao, China, July 9-10, 2011, Proceedings, Part I*, pages 274–282. Springer, 2011.
- 29 Pamela Zave. Using lightweight modeling to understand chord. *SIGCOMM Comput. Commun. Rev.*, 42(2):49–57, mar 2012. doi:10.1145/2185376.2185383.
- 30 Fanping Zeng, Qing Cao, Liangliang Mao, and Zhide Chen. Test case generation based on invariant extraction. In *2009 5th International Conference on Wireless Communications, Networking and Mobile Computing*, pages 1–4. IEEE, 2009.

A Ivy Specification of the Pipelined Moonshot Protocol

The IVy specification of the Pipelined Moonshot is provided in the module **Moonshot**. The white paper [23] by Supra research describes the same protocol as [12] but in a format that is better suited to serve as a starting point for implementation. The structure of our IVy model closely follows the description in [23], so we use it in the following to explain the organization of the IVy model.

The Moonshot module consists of declarations of state variables to be maintained by honest validators as given in [23, Table II]. These are then followed by *actions*, the keyword in IVy used to denote updates to the state variables performed in response to specific events. Below we list the actions and the corresponding events specified in [23]. Here, f is the maximum number of Byzantine validators tolerated.

■ **Table 2** IVy actions and their corresponding subroutine in the specification.

Action	Subroutine in [23]	Triggering event
qc_processing	Algorithm 2 line 27	Receiving a quorum certificate
optimistic_proposal_processing	Algorithm 2 line 37	Receiving an optimistic proposal
normal_proposal_processing	Algorithm 2 line 49	Receiving a normal proposal
timer_expire	Algorithm 3 line 73	Timer expires
timeout_sync	Algorithm 3 line 76	Receiving timeout messages from $f + 1$ validators
tc_processing	Algorithm 3 line 80	Receiving a timeout certificate
fallback_proposal_processing	Algorithm 3 line 89	Receiving a fallback proposal

Following is a code snippet from the action `normal_proposal_processing`, broadcasting a prepare message.

```
# This function encodes the conditions necessary
for processing a normal proposal
function send_prepare_n_condition(B_pr:block_t, QC:qc_t) : bool
definition send_prepare_n_condition(B_pr, QC) = block_t.round(B_pr,r_c)
& a_f < r_c & t_l < r_c & (a_o < r_c | b_o = B_pr) &
(forall B:block_t. forall R:round_t. qc_t.block(QC,B) &
block_t.round(B,R) -> block_t.parent(B_pr,B) & round_t.succ(R, r_c))

# The procedure in Line 49 of Algorithm 2, executed upon receiving
# a normal proposal
action normal_proposal_processing(b_pr:block_t, qc:qc_t) = {

  require received_proposal_n(b_pr, leader(r_c));
  require received_qc(qc);
  require block_t.cstd(b_pr);

  # Require that the timer has not yet expired for this round
  require ~ t_r;
  require ~ possessed_normal_for_round(r_c);

  #require that the parent of the proposed block b_pr is certified by
  # the accompanying QC qc
  require forall B:block_t. qc_t.block(qc,B) ->
  block_t.parent(b_pr,B);

  possessed_normal_for_round(r_c) := true;

  # If the accompanying qc is not yet processed yet, do that first
  if some b:block_t. qc_t.block(qc,b) & ~ processed_qc(b) {
    call qc_processing(qc);
  }
}
```

3:16 Formally Verifying the Safety of Pipelined Moonshot Consensus Protocol

```
# After processing the accompanying the qc, require that the
# conditions in lines 50-56 of Algorithm are met
require send_prepare_n_condition(b_pr,qc);

# This condition verified by IVy ensures that the parent of the
# proposed block b_pr is for a strictly lesser round
ensure block_t.parent(b_pr,Bp) & block_t.round(Bp,Rp) ->
Rp < r_c;

#Line 58: propose optimistic
##### proposeOptimistic #####
# Line 7,8 of Algorithm 1
var rs := round_t.next(r_c);
if leader(rs) = id & b_o ~= b_pr{

    #Lines 10-11 of Algorithm 1
    var b := block_t.block(rs,b_pr);
    var m : msg;
    m.kind := msg_kind.proposal_o;
    m.block := b;
    m.src := id;

    call shim.broadcast(id, m);
}

if send_prepare_n_condition(b_pr,qc) {

    # Line 59 of Algorithm 2: broadcast prepare normal message
    var m : msg;
    m.kind := msg_kind.prepare_n;
    m.block := b_pr;
    m.src := id;

    call shim.broadcast(id, m);

    # Line 60 of Algorithm 2: a_n is updated to the current
    # round
    a_n := r_c;
}
}
```