

A Survey of Smart Contract Formal Specification and Verification

PALINA TOLMACH, Nanyang Technological University, Singapore and Institute of High Performance Computing (A*STAR), Singapore

YI LI, SHANG-WEI LIN, and YANG LIU, Nanyang Technological University, Singapore

ZENGXIANG LI, Institute of High Performance Computing (A*STAR), Singapore

A smart contract is a computer program which allows users to automate their actions on the blockchain platform. Given the significance of smart contracts in supporting important activities across industry sectors including supply chain, finance, legal and medical services, there is a strong demand for verification and validation techniques. Yet, the vast majority of smart contracts lack any kind of formal specification, which is essential for establishing their correctness. In this survey, we investigate formal models and specifications of smart contracts presented in the literature and present a systematic overview in order to understand the common trends. We also discuss the current approaches used in verifying such property specifications and identify gaps with the hope to recognize promising directions for future work.

Additional Key Words and Phrases: Smart contract, formal verification, formal specification, properties

1 INTRODUCTION

A proposal to store timestamped tamper-resistant information in cryptographically secure chained blocks has gradually evolved from 1991 [100] into what we now know as *blockchain*: a distributed ledger shared between nodes of a peer-to-peer network following a certain consensus protocol, and one of the fastest-growing technologies of the modern time. The blockchain technology owes its initial popularity to Bitcoin [155], which appeared in 2008. Later, Ethereum [220] started the era of Blockchain 2.0 and expanded the capabilities and applications of blockchain by introducing Turing-complete *smart contracts*. The term *smart contract* was originally proposed by Nick Szabo [198], who suggested encoding the contractual terms and protocols of the involved parties in the digital form. In the blockchain context, smart contracts currently refer to computer programs executed on top of blockchain, and we concentrate on this definition of a smart contract in this work. In Ethereum, each contract is associated with an address, a balance, and a piece of executable code. The functions in a contract can be invoked via transaction messages sent from platform users. Most of the smart contracts are written in Turing-complete languages, such as Solidity [6].

The market cap of the Ethereum cryptocurrency was worth \$20 billion as of February 1, 2020, and several hundred of thousands of transactions are processed each day [5]. Unfortunately, the adoption of blockchain and smart contracts is also accompanied by severe attacks, often due to domain-specific security pitfalls in the smart contract implementations. Examples include theft of Ether through notorious attacks on the DAO [188] and Parity Multisig Wallet [124], as well as the freezing of users' funds in the same Parity Multisig Wallet [200] and the King of the Ether Throne [1] smart contracts. Extensive research has been done over the past several years—by technology giants,

This research is partially supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 1 (Award No. 2018-T1-002-069) and Tier 2 (Award No. MOE2018-T2-1-068), and by the National Research Foundation, Singapore, and the Energy Market Authority, under its Energy Programme (EP Award No. NRF2017EWT-EP003-023). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of National Research Foundation, Singapore and the Energy Market Authority.

Authors' addresses: Palina Tolmach, palina001@ntu.edu.sg, Nanyang Technological University, 50 Nanyang Avenue, Singapore, 639798, Institute of High Performance Computing (A*STAR), 1 Fusionopolis Way, #16-16 Connexis North, Singapore, 138632; Yi Li; Shang-Wei Lin; Yang Liu, {yi_li,shang-wei.lin,yangliu}@ntu.edu.sg, Nanyang Technological University, 50 Nanyang Avenue, Singapore, 639798; Zengxiang Li, zengxiang_li@outlook.com, Institute of High Performance Computing (A*STAR), 1 Fusionopolis Way, #16-16 Connexis North, Singapore, 138632.

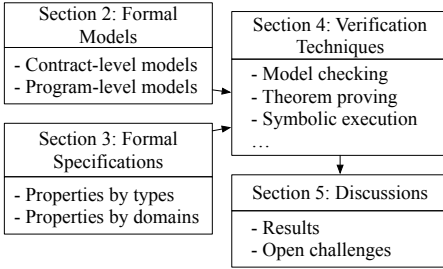


Fig. 1. Organization of the paper.

```

("formal"
 ("verification" OR "modeling" OR "specification")
 "of smart contracts")
OR ("smart contract properties")
OR ("smart contract temporal properties")

("formal"
 ("verification" OR "modeling" OR "specification")
 "of smart contracts")
AND ("properties" OR ("temporal properties"))
  
```

Fig. 2. Search queries used in Google Scholar.

non-profit organizations, and the research community—on the techniques for verifying smart contract correctness. In this survey, we review the recent advances in the applications of formal methods in the analysis of smart contracts, with a focus on the domain-specific properties and various formalisms supporting the specification and verification of such properties.

We organize the paper following a structure outlined in Fig. 1. *Formal verification* proves or disproves the correctness of a system by checking the *formal (mathematical) model* of the system against a certain *formal specification*. A specification is a set of *properties* describing the desired behaviors of a smart contract, usually defined by developers’ intention. Models and specifications can be defined at different levels of abstraction via various types of formalisms, which we overview in Sects. 2 and 3, respectively. Section 3 additionally describes a taxonomy of smart contract properties, which highlights the emerging trends in the literature. Section 4 closes the loop by reviewing formal verification techniques utilized to prove the correctness of smart contract models with respect to the desired properties. Finally, Sect. 5 concludes the survey with our observations on major trends, challenges, and future directions in the formal specification and verification of smart contracts.

1.1 Methodology

Sections 1.1 and 1.2 outline the methodology that we followed to extensively collect the existing work related to smart contract formal specification and verification, and an overview of the body of related works, respectively.

Research Scope. This paper is aimed to overview, analyze, and classify approaches to formal modeling of smart contracts, specification of smart contract properties, and techniques employed in the verification of those. The scope of the survey includes a study of common property specifications of smart contracts, characterizing their functional correctness and security guarantees. The focus of the survey is narrowed down to smart contract functional behavior and does not embody problems related to other blockchain-related execution aspects, such as scalability, consensus, interaction with IoT systems, etc. Therefore, the purpose of the paper is to provide answers to the following *research questions*:

- RQ1:** What are the formal techniques used for modeling, specification, and verification of smart contracts? What are the common formal requirements specified and verified by these techniques?
- RQ2:** What are the challenges introduced by smart contract and blockchain environment in formalizing and verifying smart contracts?
- RQ3:** What are the current limitations in smart contract formal specification and verification and what research directions may be taken to overcome them?

Search Criteria. In order to provide a complete survey covering most of the publications related to smart contract specification and verification since 2014, we created a publication repository that includes 202 papers published from September, 2014 to June, 2020. To collect the relevant literature for analysis, we have performed the search in Google Scholar, ACM, IEEE, and Springer databases using the two queries shown in Fig. 2.

In October 2019, we have explored 10 pages of the results retrieved for each query and obtained 70 papers that fall into the scope of the survey and have been published since 2014. Based on the collected set of papers, we performed snowballing and inverse snowballing to obtain a more comprehensive view by collecting 85 more publications. Additionally, to provide an overview of the state-of-the-art, we included another 47 recently published papers that were retrieved using the two aforementioned queries from October, 2019 to June, 2020 by Google Scholar at the time of writing. Overall, we collected and studied 202 relevant papers. The majority of them were published in the proceedings of 70 conferences and workshops (including ACM CCS, ASE, ESEC/FSE, IEEE SANER, S&P, IEEE/ACM ICSE, NDSS, ISoLA, FC, FM, and the affiliated workshops) or one of 13 journals (including ACM TOSEM, IEEE TSE, and IEEE Access). Using our search methodology, we additionally identified 53 articles that conduct a survey on smart contract types, languages, vulnerabilities, and verification techniques, or propose security and design patterns for their development. We describe the collected articles in an online repository: <https://ntu-srslab.github.io/smart-contract-publications/>.

In the presented work we do not claim to provide a complete set of properties that smart contracts are supposed to meet. However, by means of the presented analysis and classification, we seek to provide a firm formalized foundation and a stable structure for smart contract property specifications, what we hope will facilitate the development of correct, secure, and standardized smart contracts.

1.2 Related Works

By now, a multitude of surveys addressing smart contract analysis has been published. These include a review of security vulnerabilities [26, 97, 149, 181, 224], verification approaches and tools [55, 59, 71, 75, 78, 154, 166, 169, 227], formal specification and modeling techniques [35, 114, 189], and languages for smart contract development [104, 164]. We found several pieces of work discussing common design patterns of smart contracts [218, 219, 221, 223], as well as studies of smart contract platforms and practical applications [27, 34, 212, 216]. Another line of research focuses on revealing challenges and opportunities for safer smart contract development [24, 144, 150, 232, 235]. In addition, some surveys analyze performance and scalability aspects of smart contracts [176].

The surveys most relevant to this work are concerned with the application of formal methods to modeling smart contracts. Singh et al. [189] analyzed 35 research works on formal verification and specification techniques as well as languages for smart contracts, together with the issues and vulnerabilities they address. Bartoletti et al. [35] compared five formal modeling techniques for Bitcoin smart contracts based on their expressiveness, usability, and suitability for verification. Furthermore, Ladleif and Weske [125] proposed a framework for evaluating formal modeling tools with respect to their capability to model a legal smart contract, and assessed eight visual tools for smart contracts modeling using this framework. Different from the discussed works, our survey proposes a coherent conceptual framework that helps to establish the links between the approaches to formal modeling, specification, and verification. Our analysis is also more comprehensive in terms of the types of formal techniques covered and the application platforms considered.

There has been a body of literature dedicated to the collection of both bad and good ways of writing *secure* smart contracts. On one hand, the common pitfalls discovered in past experiences should be documented and avoided in the future development; and on the other hand, the best

practices which are proven to work can be summarized and adopted when facing a similar context. Atzei et al. [26] and some more recent work [97, 149, 181, 224] collected common vulnerability patterns which might compromise the security of smart contracts. Permenev et al. [167] and Bernardi et al. [39] defined several classes of properties for which security analysts should watch out when performing contract audit. Several groups of researchers [34, 218, 219, 221, 223] proposed design patterns for smart contracts, which are general and reusable solutions for specific tasks. In this survey, we look beyond the security aspect and aim to provide a general taxonomy of domain-specific properties, which may also impact the correctness, privacy, efficiency, and fairness of smart contracts. We are also interested in establishing the links between the various types of properties and the formalism used in supporting them. We believe that this effort can help identify the gaps between the analysis needs and the existing technology capability.

We observe a consensus from the previous surveys that specifications for smart contracts are of paramount importance, but how to properly derive and represent them still remains an open challenge. This impacts not only the development of safe smart contracts, but also their verification and utilization [96, 112, 144, 212]. According to a recent survey on smart contract developers by Zou et al. [235], best practices, standards, and provably safe re-usable libraries are listed among the most desired improvements they would like to see to happen in the smart contract ecosystem. While there are ongoing efforts towards this direction, such as the creation of curated datasets of smart contracts [75, 90], we aim to facilitate the adoption of formal specification and verification by contributing the results of our analysis as well.

We adopt the same holistic view of verification techniques and tools for smart contracts. Still, our categorization of formal verification approaches is inspired by the previous studies on smart contract verification techniques performed by Harz et al. [104], Di Angelo et al. [71], and Miller et al. [150]. In this survey, we make the following novel contributions.

- A comprehensive survey of the state of the art with broader scope;
- A four-layer framework for classifying smart contract analysis approaches;
- A taxonomy of smart contracts specifications in various domains;
- Discussions on observed trends and challenges in smart contract development and verification.

2 A TAXONOMY OF SMART CONTRACT MODELING FORMALISM

This section presents the formalisms used for formal modeling of smart contracts in the literature. The formalization approaches for smart contracts can be classified into two broad categories: *contract-level* and *program-level*, based on the level of abstraction at which modeling and analysis is performed. The *contract-level* approaches are concerned with the high-level behavior of a smart contract under analysis and usually do not consider technical details of its implementation and execution. These approaches usually describe the interactions between smart contracts and external agents mainly represented by users and the blockchain. The details of a smart contract are typically abstracted into a set of public functions that are invoked by users and observable results of their executions, such as emitted events or a change in a blockchain state. The *program-level* approaches mostly perform analysis on the contract implementation (i.e., source code or compiled bytecode) itself and thus are platform-dependent. By using program-level representations, such approaches enable precise reasoning about low-level details of the smart contract execution process.

2.1 Contract-Level Models

In the contract-level analyses, smart contracts are viewed as black boxes which accept transaction messages coming from outside, and possibly perform some computation based on them. This may result in externally observable events and even irreversible alteration to the blockchain state. The

internal execution details and intermediate results are abstracted away in these types of models. Essentially, the contract-level models are concerned with the following concepts.

- *Users* – together with their balance, the transactions they initiate and the associated parameters such as the account addresses and the attached values. They can be any type of accounts which interact with the contract of interest—a user account or another contract. Some papers [28, 54, 128] also consider the choices of transactions made by users under different circumstances (i.e., strategies), as well as knowledge that is available to users at a particular moment of a smart contract execution [20, 28, 109, 207].
- *Contracts* – characterized by a set of publicly accessible functions and the externally visible effects of the function executions, e.g., changes of contract owners/balances, and emitted events/operations. As such, the behaviors of a contract are usually described by a set of *traces* [187], which are finite sequences of function invocations (i.e., the executed operations [147]), and properties are predicates on such traces.
- *Blockchain State* – including global variables referred to by the contracts and environment variables such as the timestamps and block numbers. More generally, blockchain state may also include the mining pool [10, 157] and the memory state of the contract execution environment [147].

The contract-level models are effective in expressing properties regarding the interactions between smart contracts and the external environment, and they are usually defined in terms of process algebras, state-transition systems, and set-based methods. The majority of such models are verified using techniques based on model checking [10, 16, 20, 32, 143, 147, 151, 157, 163, 171].

2.1.1 Process Algebras. Process algebra [31] is a family of mathematical approaches used to describe the behaviors of distributed or parallel systems in the form of interacting concurrent processes. The observable behaviors of a smart contract are similar to that of a shared-memory concurrent program [184] or a system with an interleaving mode [171], which makes process algebra a suitable candidate for high-level behavioral modeling. Therefore, process algebra captures the interactions between concurrently acting users and one or several interconnected smart contracts [202]. For these interactions, correctness can be assured through a translation of a smart contract to one of the process-algebraic formalisms. An alternative approach involves implementation of a smart contract in a process-algebraic domain-specific language (DSL).

Given that users interact with an Ethereum contract by calling its public functions, Qu et al. [171] and Li et al. [133] translate such functions from Solidity to process notations defined in *Communicating Sequential Processes* (CSP) and a variant of an applied π -calculus *SAPIC*, respectively. In case transactions are not atomic, the interplay between the functions invoked by different users may lead to risk conditions, which are manifested by vulnerable sequences appeared in the execution trace [171]. The authors of [133] identify violations of trace-level safety properties that are commonly analyzed for token implementations, such as the immutability of the total number of tokens in a smart contract. As an example, Figs. 3 and 4 demonstrate partial source code and a CSP model of the Simple Auction smart contract, which is based on an example from the Solidity documentation [7] and the Blind Auction smart contract provided in [148]. Within the contract, each user is allowed to place his bid until the deadline is reached. After the auction is finished, the highest bid is transferred to the beneficiary, while the unsuccessful bidders can withdraw their bids. Alternatively, the auction can be canceled at the bidding stage. Since the users participate in the auction by invoking functions of a smart contract, Figure 3 shows the Solidity implementation of two functions, namely, `bid()` and `finish()`. Following the approach adopted by the authors of [133, 171], Fig. 4 represents each of these functions as a process: `Bid(msg, blk)` and `Finish(blk)`, respectively. The processes are specified in CSP#—a CSP-based formal language

```

1  contract SimpleAuction {
2      enum State { Bidding, Finished, Cancelled }
3      State public state;
4      address payable public beneficiary;
5      uint public auctionEndTime;
6      address public highestBidder;
7      uint public highestBid;
8      mapping(address => uint) pendingReturns;
9      ...
10     /// Bid on the auction.
11     function bid() public payable {
12         require(state == State.Bidding);
13         require(msg.value > highestBid);
14         require(now <= auctionEndTime);
15         if (highestBid != 0) {
16             pendingReturns[highestBidder] += highestBid; }
17         highestBidder = msg.sender;
18         highestBid = msg.value;
19         emit BidIncreased(msg.sender, msg.value);
20     }
21
22     /// End the auction, send owner the highest bid.
23     function finish() public {
24         require(state == State.Bidding);
25         require(now >= auctionEndTime);
26         state = State.Finished
27         emit AuctionEnded(highestBidder, highestBid);
28         beneficiary.transfer(highestBid);
29     }
30     ...
31 }

```

Fig. 3. Solidity source code of Simple Auction smart contract adapted from [7, 148].

```

///@@SimpleAuction@@
enum{Bidding, Finished, Cancelled};
var state = 0;
var beneficiary = 0;
var auctionEndTime = 0;
var hBidder = 0; var hBid = 0;
var pReturns[2256 - 1];
...
Bid(msg, blk) = if (state == Bidding && msg.value > Bid
&& blk.timestamp <= auctionEndTime){
    if (hBid! = 0){
        pReturns[hBidder] = pReturns[hBidder] + hBid;
        updateBid{hBid = msg.value;
            hBidder = msg.sender;} → hBidIncreased → Skip};
Finish(blk) = if (state == Bidding &&
blk.timestamp >= auctionEndTime){
    endAuction{state = Finished;} →
    transfer.beneficiary.hBid → auctionEnded → Skip};
...

```

Fig. 4. A CSP# model of Simple Auction smart contract.

of the PAT tool [195], which supports both process and programming constructs. Each process describes the corresponding actions of a smart contract and its users in terms of CSP *events* as well as changes in a smart contract state implemented via shared variables. Being a high-level representation, the model abstracts low-level execution details, including that of the `transfer()` function, with the corresponding events. Block and transaction properties are represented using the `msg` and `blk` process parameters, similar to the globally available variables in Solidity. For brevity, the model of user behavior is omitted.

A process-algebraic model of a smart contract is adopted by several domain-specific programming languages. [36, 173]. *Rholang* [173] is a Turing-complete language for the RChain platform, which supports concurrency and is based on the reflective higher-order process calculus ρ -calculus, a variant of π -calculus. Smart contracts in Rholang are processes that are triggered by messages from agents: users or other contracts. *BitML* [36] is a high-level smart-contract language for Bitcoin, which is based on a process calculus with symbolic semantics. It is showcased in a series of publications by Bartoletti et al. [28, 36] how security properties, as well as arbitrary temporal properties, can be expressed and checked in BitML. To ensure correctness of a smart contract under different scenarios, BitML also supports specification of user strategies. Furthermore, smart contracts can also be soundly translated from BitML into transactions—a Bitcoin representation of a smart contract. The behavior of contract users can also be partially defined in a Turing-complete process calculus *scl* proposed by Laneve et al. [128]. *scl* models behaviors of a smart contract and a user as a parallel composition, where a user acts nondeterministically, while the behavior of a smart contract is determined by its execution logic. The *scl* models are in fact transition systems that are summarized by a first-order logic formula describing the values of the objective function for all possible runs and user strategies. Then, the game-theoretical analysis is performed to identify strategies that help users maximize their profits or minimize losses.

2.1.2 State-Transition Systems. The behaviors of a smart contract can be naturally interpreted as a state-transition system, and in fact, Solidity encourages contracts being modeled as state machines [3]. Some next-generation smart contract languages, such as Obsidian [65] and Bamboo [2], follow a state-oriented programming paradigm which makes transitions between states explicit and therefore provides better security guarantees. When it comes to verification, depending on the properties interested, there are also a wide range of choices in modeling smart contracts as state-transition systems. For instance, timed automata, Markov decision processes, and Petri nets are used to capture the time, probabilistic, and multi-agent interactive properties of the systems, respectively. State-transition models are typically verified by model checkers against contract-specific properties, in addition to generic temporal logic properties such as deadlock-/livelock-freedom (C.f. Sect. 4.1). Section 3.2.1 provides a discussion on how the violation of these properties can lead to serious bugs. A variety of analysis frameworks already exist for the validation and verification of state-transition systems, e.g., BIP and cpntools, etc. Such verification capacities combined with the ability to visually define a state-transition system make this formalism useful for smart contract synthesis [147, 148, 236].

Petri net (PN) is a mathematical language used to describe behaviors of concurrent systems. Unlike the aforementioned formalisms, it provides a graphical notation which portrays a system using a set of alternating places and transitions. PN tools allow a user to visually specify a smart contract model and simulate its possible behaviors. Similar capabilities of the *Colored Petri net* (CPN) extension and the supporting cpntools toolkit are utilized for security analysis of a Solidity contract [74, 138]. These two techniques construct CPN models of a contract and its users—manually [138] or automatically [74]—to identify if a vulnerable state of a net is reachable. The authors of [74] make the simulation more realistic by additionally modeling the execution process of EVM, considering gas consumption as well. The ability to visually define and execute a PN model also helps in synthesizing smart contracts from user inputs. A framework by Zupan et al. [236] generates a Solidity smart contract from a model defined as a PN workflow—a subtype of a PN describing business processes. Similar to the discussed PN approaches [74, 138], Zupan et al. automatically validate the model according to typical properties such as the absence of transitions that can never be executed [74, 138, 236] or reachability of a final state [236]. Through the visualization and validation of the model and its execution, a correct-by-design smart contract can be generated by a user without prior knowledge of smart-contract programming. The authors expect the technique to facilitate the application of smart contracts in use cases including supply chain [236] and legal contracts [125].

For a similar purpose, several frameworks synthesize smart contracts from a specification defined using industry standards for visual modeling, such as *Business Process Model and Notation* (BPMN) and *Unified Modeling Language* (UML). Caterpillar [72, 139] and Lorikeet [72] translate BPMN models, which describe collaborative business processes such as supply chain, to Solidity and deploy them to blockchain. A Solidity implementation of a cyber-physical system is generated from a UML state diagram in [87]. BPMN and UML notations provide the capabilities to specify proper access control [72, 87], control flow [139], and asset management [72] in the generated contracts. Still, despite being widely used in the industry, BPMN and UML lack well-established formal semantics, which complicates the analysis of such models. A possible way to validate a BPMN model involves translation to a formal language, such as PN, as performed by Garcia et al. [88].

Behavior-Interaction-Priority (BIP) is a layered framework used to model the interaction between smart contracts and users. BIP models of smart contracts usually consist of two layers: the components of the *Behavior* layer describe the logic of each smart contract as an FSM, while the *Interaction* layer defines the principles of their communication [16, 158]. Translation of an FSM model, visually defined in the FSolidM framework [148], to BIP and, then, an input language of a NuSMV model




Fig. 5. A state-transition model of Simple Auction smart contract (Fig. 3) adapted from [148].

checker, allows VeriSolid [147] to generate provably safe smart contracts (in contrast to UML- and BPMN-based frameworks discussed earlier [72, 87, 139]). Figure 5 presents an FSM model of the Simple Auction smart contract (Fig. 3), inspired by a model of the Blind Auction contract defined using FSolidM [148] and VeriSolid [147]. The model in Fig. 5 depicts possible transitions between three smart contract states: Bidding (B), Finished (F), and Canceled (C).

Timed automata (TA) capture the temporal behavior of a system. TA models allow measuring the impact of time on the execution of a smart contract [10] or reasoning about time constraints, which are essential in legal and commitment smart contracts [64]. Different from other approaches based on BIP [16, 147, 158], Abdellatif et al. [10] build TA models of not only a contract, but also its users, and a simplified blockchain mechanism, including mining of transactions in blocks. The interactions between these components are analyzed by statistical model checking to evaluate the chances of a successful attack on the contract. Andrychowicz et al. [20] use TA to describe behaviors of parties in a Bitcoin timed commitment. Commitment protocols require participants to exchange cryptographic hashes of the secret values they commit to before a stipulated time. To analyze correctness of this protocol, the authors additionally equip each TA with a structure describing the party’s knowledge. The blockchain structure is maintained by a special agent, who is also modeled as an automaton.

A way to reason about the nondeterministic execution of a smart contract involves modeling it through probabilistic and multi-agent extensions of transition systems, such as *Markov decision processes* or *strategic and concurrent games* [43, 54, 128, 207]. To account for all possible user behaviors, Meyden [207] regards an atomic swap smart contract as a concurrent game. Both a model and its properties refer to the strategies of the contract participant. The authors apply the MCK model checker to establish correctness and fairness of a contract under the specified user behaviors. Chatterjee et al. [54] pursue a similar goal, however, their analysis of a concurrent game model is based on its objective function, similar to the *scl* process-algebraic technique [128]. To certify fairness and correctness of a smart contract, the authors of [54] examine the expected payoff of its users and potential incentives for dishonest behavior. Different from the game-theoretical techniques, Madl et al. [143] capture the interaction between the users of a marketplace contract by describing their behaviors as *interface automata*. The authors verify a composition of the automata in a model checker to ensure that participants can successfully cooperate.

2.1.3 Set-Based Methods. Set-based modeling frameworks *Event-B* and *TLA+* utilize set theory and logic to formally specify systems, and have been applied on smart contracts as well. Analogically to other formal languages, such as PN, set-based frameworks support the tools for model analysis, making them useful in the implementation of correct-by-design smart contracts [33, 222] and verification of the existing ones [234]. The authors of the identified publications [33, 222, 234] utilize set-based formalisms to reason about safety properties of smart contracts. Although *TLA+* supports temporal logic for specifying liveness, this feature has not yet been used to formalize smart contract requirements.

Models written in Event-B are discrete transition systems which were demonstrated to be a suitable formal representation of a smart contract. Zhu et al. [234] automatically translate a contract written in a subset of Solidity to an Event-B model for refinement and verification via the Rodin platform. TLA+ is a formal specification language for concurrent systems. That allows Xu et al. [222] to model behaviors of interacting participants of a legal smart contract. Although the authors of [222] do not translate their TLA+ specification to executable code, they propose a set of design patterns to protect them from common security vulnerabilities of smart contracts. In general, the set-based correctness requirements for smart contracts specify values allowed to be taken by contract variables [33, 234], possible states of contract participants [222], or the users that are authorized to invoke contract functions [234]. Xu et al. [222] also verify whether events representing the actions of contract parties are permitted according to the legal contract.

2.2 Program-Level Models

The contract-level models are useful in reasoning about high-level interactions between the contracts and external parties, but do not help much in understanding the internal details about the contract execution. The latter is essential in establishing properties concerning the correctness and security of smart contracts. The program-level models aim to provide a white-box view of the target contracts based on lower-level representations, such as the source code, compiled bytecode, as well as derived analysis artifacts including AST, data- and control-flow graphs.

The following are concepts that are relevant to this level of abstraction.

- *Abstract Syntax Tree (AST)* – representing smart contract source code (e.g., in Solidity) as hierarchical tree structures. It is often used to perform lightweight syntactic analyses on the contract implementations.
- *Bytecode* [167] (or *opcode*) – written in low-level machine instructions, which are closely tied to the execution environments (e.g., Ethereum Virtual Machine (EVM)). Since bytecode is obtained after compilation, it better reflects the machine-level specifics (e.g., gas consumption and exceptions), but at the same time may lose important source-level information.
- *Control Flow Graph (CFG)* – a graph representation of all program paths which might be traversed during execution. The CFG can be obtained from bytecode and is often used in static analyses of smart contracts, such as symbolic execution and automated verification.
- *Program Traces* – sequences of instructions (usually in bytecode) and events collected from executions at runtime. These traces reflect the exact behaviors of the contracts under specific inputs, which can be used as a source for performing dynamic analyses and runtime verification.

The program-level models preserve low-level execution details, and are, therefore, widely used in finding vulnerabilities and checking other security-related properties.

2.2.1 AST-Level Analyses. Smart contract analyses performed directly on ASTs or similar parse-tree structures are often aimed at the checking of predefined code patterns [140, 201, 224]. For example, Yamashita et al. [224] scan ASTs of Hyperledger Fabric [19] contracts written in Go to find dangerous code patterns and library imports. Another AST-based analysis is successfully applied in checking the compliance of a smart contract with respect to the ERC20 [209] standard by Chen et al. [56]. Moreover, AST can be used to facilitate lightweight pre-analyses: e.g., locating arithmetic operations that may cause overflows or assertions [208], decoding the memory layout information to allow deeper inspection [48], and performing systematic instrumentation, such as vulnerability checks, to the contract code [12, 178]. Such pre-analyses can help improve the effectiveness and scalability of the more heavyweight downstream analyses. For instance, AST is the natural abstraction of smart contract code when it comes to statistical [135] and deep learning-based analyses [86, 113].

The obvious drawback of these purely syntactic techniques is that they do not necessarily respect the operation semantics or the execution environment of smart contracts, thereby compromising the soundness/completeness of the analyses. For example, the gas consumption plays an important role in the functional correctness and security of Ethereum smart contracts, but is often neglected [140, 201]. To allow for more sophisticated analyses, ASTs are often translated to derive other intermediate representations, such as CFG [12, 77], inheritance graph [77], and inter-contract call graphs [56].

2.2.2 Control-Flow Automata. *Control-flow graph* (CFG) is often used to describe the operational semantics of a program. As the name suggests, CFG is a labeled directed graph where the graph nodes correspond to program locations and the edges correspond to possible transitions between the program locations. Thus, CFG can also be viewed as a type of state-transition system, except that it more closely reflects contract program executions than the contract-level state transitions discussed in Sect. 2.1. CFGs of smart contracts are mostly constructed from the compiled code—either EVM bytecode of Ethereum contracts [141] or WebAssembly (WASM) bytecode of EOSIO [4] contracts [106, 172]. Still, recovering CFGs from low-level bytecode is non-trivial, at least in the case of Ethereum. There are several reasons for this. First, EVM is a stack-based machine, therefore, recovering the target address of a jumping instruction (destination of an edge) calls for context-sensitive static analysis capable of tracking the state of stack [15, 52]. Furthermore, smart contract execution comprises a sequence of function invocations and, often, calls to other accounts, which demands inter-procedural and inter-contract analyses, respectively.

The types of analyses run on CFGs include control-/data-flow analyses and symbolic execution, which have matured tool supports for traditional software programs. Some of these tools and techniques are applied directly on the CFGs extracted from smart contracts, while others are custom built to accommodate particular language features of contract programs. As one example, Ethainter [49] performs information flow analysis to reveal composite vulnerabilities, i.e., reached only through a sequence of transactions, from traces of Solidity contracts. In order to reuse the existing program analysis tool-chain, namely, Datalog and Soufflé, the EVM bytecode is decompiled using GigaHorse [94] to recover the CFG with data- and control-flow dependencies. The same tool-chain is also used by Securify [205] and Vandal [50] in vulnerability detection after CFGs are constructed. In contrast, many custom symbolic execution engines have been built to traverse CFGs and check for interesting properties along the way. Some of the most notable ones include Oyente [141], Mythril [153], and Manticore [152]. These approaches are designed to identify vulnerable code patterns [141, 152, 153], integer arithmetic bugs [81], honeypots [204], gas-inefficient bytecode patterns [58], etc.

To allow specification and verification of custom properties, several tools support extensions to CFGs with annotations. Annotary [217] extends Mythril [153] to also capture inter-contract and inter-transaction control-flows. Solar [131] allows annotations to a smart contract with constraints and invariants. SmartScopy [79] supports vulnerability patterns to be defined in their query language based on Racket.

2.2.3 Program Logics. To rigorously reason about the correctness of programs, various forms of program logic have been developed over the years, and some of them are successfully applied on smart contracts. A *program logic* is a proof system with a set of formal rules that allow to statically reason about behavior of a program [23]. For example, Hoare logic [18], rewriting logic, and reachability logic [165] were used to prove correctness conditions (C.f. Sect. 3.1.2) for smart contracts. An epistemic logic [109] describes the interactions between smart contract participants in terms of their knowledge, which facilitates verification of commitment and swap protocols (C.f. Sect. 3.2.2). A combination of authority algebra and FOL enables reasoning about trust and accountability in a smart contract [69]. Correctness of smart contracts is also verified through

a dynamic logic developed for Java programs [11, 37, 38]. Finally, formalisms that are based on deontic logic enable specification and verification of legal smart contracts, which are difficult to define formally otherwise [125].

Hoare and reachability logics assist reasoning about smart contract correctness by different formal techniques. Based on a formalization of EVM in *Lem* [108], Amani et al. [18] and Duo et al. [74] defined a sound Hoare-style program logic of Ethereum contracts. In verification of properties, however, the former approach relies on theorem proving, while the latter translates a formal model of a contract to CPN for model checking. In other approaches based on theorem proving, Hoare-style properties are verified for smart contracts written in Michelson [40]—a smart contract language of the Tezos [91] platform, or Lolisa [226]—a formal syntax and semantics of a Solidity language subset. Apart from the tools based on Hoare logic, \mathbb{K} is a framework for design and development of language semantics based on rewriting logic [165]. \mathbb{K} has been used to define a complete executable formal semantics of EVM [165] and Solidity [117], and a formal specification of an intermediate LLVM-like language for smart contracts IELE [119]. Properties that can be automatically proved by \mathbb{K} include pre- and postconditions as well as reachability claims that are verified through matching logic. The expressiveness of \mathbb{K} can be illustrated with a complete formalization of an ERC20 token contract [115]. Despite different choices of a verification technique, the aforementioned approaches are able to capture the dynamic aspects of smart contract execution, such as the memory state and gas consumption.

Such verification systems are usually built upon formal semantics of smart contract languages, e.g., EVM bytecode [18, 74, 108], Solidity [226], Michelson [40]. To make contracts amenable to formal reasoning, developers equip smart contract languages with a logical foundation, such as λ -calculus, and formal semantics [53, 186].

At the same time, smart contracts that are written or translated to Java can be verified using a dynamic logic called JavaDL—a program logic for Java. A Java program annotated with JML specifications is automatically translated to JavaDL by KeY—a framework for deductive verification. While Hyperledger Fabric contracts written in Java can be verified directly [37, 38], Solidity smart contracts are preliminarily translated to Java [11]. Dynamic logic can be seen as an extension of Hoare logic and is also used to define an executable specification for smart contracts [180]. The concepts related to dynamic logic are discussed in more detail in Sect. 3.1.1 and Sect. 3.1.2.

Some researchers consider logic-based languages to be particularly useful for the implementation and verification of *legal smart contracts* [93]. A legal smart contract is a digitized representation of a legal agreement, which typically includes deontic modality statements describing obligation, permission, and prohibition of the parties. To capture contractual aspects, the authors of [93] formalize a smart contract in a logic-based language called *Formal Contract Logic* (FCL). Azzopardi et al. [30] propose a formal deontic logic with small-step operational semantics for Solidity smart contracts [30]. Specifications defined in deontic and defeasible logics are discussed in Sect. 3.1.1.

2.3 Summary

The analyzed work illustrates different aspects of smart contracts that researchers focus on, together with the appropriate formalisms. As such, the concurrent nature of smart contract execution is captured by contract-level models based on process algebra and Petri net, while game-theoretical structures are suitable for modeling a multi-agent environment. Contract-level models, and state-transition systems, in particular, provide an intuitive way to outline high-level behavior of a smart contract, which includes possible interactions with other contracts, users, or blockchain. Still, the abstraction over the execution logic of a smart contract introduces a gap between an original contract and its representation in a chosen formalism. On the other hand, program-level representations give intuition about the lower-level details of the execution process. At the same

time, while bytecode-level modeling provides wide capabilities in terms of low-level reasoning, e.g., about resource consumption, it also complicates specification of semantic properties related to smart contract behavior [22].

3 SURVEY ON SMART CONTRACT SPECIFICATIONS

In this section, we review the literature on smart contract specifications. Requirements on smart contract applications are usually specified as *properties*—statements written in some formal languages about program behaviors. In Sect. 3.1, we survey the formalisms adopted in formal specification, which can be categorized into the contract- and program-level specifications. In Sect. 3.2, we review various types of properties (e.g., security and functional correctness) from different application domains, observed in the published work under study.

3.1 Formal Specifications for Smart Contracts

Similar to the modeling formalisms discussed in Sect. 2, we make a distinction between *contract-level* and *program-level* formal specifications of smart contracts. This is in correspondence with the traditional classification of *model-oriented* and *property-oriented* specifications, rooted from *temporal logic* and *Hoare logic*, respectively. We discuss contract-level specifications about high-level contract properties in Sect. 3.1.1 and program-level specifications governing lower-level program implementations in Sect. 3.1.2.

3.1.1 Contract-Level Specification. Contract-level specifications of smart contracts are expressed in various types of logic. The most prevalent category is a family of temporal logics, which includes linear and branching temporal logics and their dialects. Other logics include deontic and defeasible logics, which help to define the rights and obligations of smart contracts’ parties in accordance with a legal contract. As specifications in temporal [147, 197], dynamic [180], and deontic logics [47, 84] capture high-level behavioral characteristics of a system, they are used to synthesize correct-by-design smart contracts. Propositional and predicate logics usually convey factual statements about the contract’s state over the traces.

Temporal Logics. Temporal logics express properties of a smart contract over time. The execution logic of smart contracts often encloses time constraints, be it legal contractual clauses [64] or ending time of an auction or voting contract [84]. As noted in [157], in relation to smart contracts, the word “temporal” may refer to logical time considering only the order of events. Defined over a sequence of states, temporal formulae constitute a suitable language to describe the behavior of a state-transition system—a common formal representation of a smart contract described in Sect. 2.1.2, which is then verified by a model checker.

The two key temporal properties are *safety* and *liveness*. Safety properties correspond to a statement: “Nothing bad ever happens”, and are often used to express invariance. In addition to general software properties such as deadlock-freedom, common safety properties of smart contracts define permissible values of state variables, establish invariants over a smart contract balance (C.f. Sect. 3.2.3), specify access control for functions, or define a possible order of events.

Liveness properties state that “something good eventually happens”. They characterize the possibility of a program to progress, which, for a smart contract, is often reflected in its ability to give away cryptocurrency. The primary liveness property is *liquidity* [36]—a security property discussed in detail in Sect. 3.2.1. Other liveness properties specify the ability of a smart contract or its users to eventually reach a particular desired state [143, 157], even considering different user strategies [207]. Unlike a safety property, however, there exists no finite counterexample for a liveness property. To make them amenable to symbolic analysis, some authors [160, 187] define under-approximations of liveness properties that are checked on traces of a fixed length.

| | |
|--|---|
| (i) bid cannot happen after close: $\mathbf{G} (\text{close} \rightarrow \neg \mathbf{F} \text{bid})$ | (i) bid cannot happen after close: $\mathbf{AG} (\text{close} \rightarrow \mathbf{AG} \neg \text{bid})$ |
| (ii) withdraw cannot happen before finish: $\text{finish} \mathbf{R} \neg \text{withdraw}$ | (ii) withdraw can happen only after finish: $\mathbf{A} [\neg \text{withdraw} \mathbf{W} \text{finish}]$ |

Fig. 6. Properties of an auction smart contract in LTL from [197] (left) and CTL from [147] (right).

Linear Temporal Logic (LTL) is a fragment of first-order logic (FOL) which is used to define properties over a linear sequence of successive states of a program. LTL properties describe safety and liveness of transition-system models of smart contracts that are verified by a model checker [16, 28, 32, 36, 151]. Temporal safety properties of the Blind Auction smart contract in LTL proposed by Suvorov et al. [197] are depicted in Fig. 6. The properties are also relevant for the Simple Auction smart contract, which is formally modeled in Section 2 (Figs. 3 and 5). *Past-time LTL* (ptLTL) has a set of temporal operators that reference the past states of a trace and allow succinct formulae for safety properties [46, 167]. For example, two ptLTL-based specification languages for Ethereum [46, 167] support an operator *previously*, which refers to the value of a contract variable before the transaction execution. The properties defined in these languages are then translated to source-code assertions. To account for the nondeterministic execution environment of a smart contract, Abdellatif et al. [10] specify properties in *Probabilistic Bounded LTL* (PB-LTL). Probabilistic properties in [10] estimate the probability of a successful attack on a name registration contract on each step of a transaction lifecycle.

Computation Tree Logic (CTL) is a branching-time logic used to describe properties of a computation tree. In contrast to LTL formulae, whose semantics model considers one path (i.e., a sequence of states) without branching, CTL operators quantify over a tree of all possible future (branching) paths starting from the given state. Still, CTL is also used to formalize the specification of an FSM representation of a smart contract for model checking [143, 147, 157]. Figure 6 also illustrates a CTL formalization of temporal properties for an auction smart contract proposed by Mavridou et al. [147]. As in LTL, various dialects of CTL are used to formulate specific requirements of smart contract models. An extension of CTL for Colored Petri Nets (CPN), *ASK-CTL*, is used to describe correctness of CPN models of smart contracts [74, 138]. To reason about the quantitative aspect of time in a timed commitment protocol, Andrychowicz et al. [20] use a *Timed CTL* (TCTL) language of the UPPAAL model checker to define properties of the timed automata models. *Probabilistic CTL* (PCTL) properties of a Markov Decision Process model of a selling contract measure the probability of one of the players to lose a certain percentage of his wealth after a given amount of time [43]. *Alternating Temporal* (ATL) and *Strategic Logics* (SL) are multi-agent extensions of CTL that allow quantification over agents' strategies. They describe the behavior of a concurrent game model of an atomic swap [207].

Dynamic Logic. *Dynamic logic* allows formulating properties in terms of *actions* taking the termination of a program into consideration. The authors of [180] propose a DSL for smart contract implementation based on properties written in *Linear Dynamic Logic on finite traces* (LDL_f). Owing to its better expressiveness compared with LTL, LDL_f is used directly to model a smart contract, instead of encoding the state-transition model in a separate modeling language, such as Promela. Other applications of dynamic logic are related to deductive verification of smart contracts in the Key framework. However, the properties are formulated in a smart contract as JML annotations [11, 37, 38] and are then translated to Java Dynamic Logic (JavaDL) (C.f. Sect. 3.1.2).

Deontic and Defeasible Logics. *Deontic logic* is a division of modal logic operating over the concepts of obligation, permission, and prohibition. As legal prose relies on deontic modality [125], deontic-logic specification of a legal smart contract can bridge the gap between its legal semantics and implementation [64]. However, deontic modality requires special treatment from the specification

perspective, as permission has no direct correspondence in many specification languages [30, 197]. To enable compliance checking between a legal contract and a Solidity smart contract, Azzopardi et al. [30] proposed an operational semantics for contracts written in deontic logic. Deontic statements also constitute one of the components of *ADICO institutional grammar* [67], which captures institutional rules, conventions, and norms. Given a set of ADICO rules, a framework proposed in [84], can generate a skeleton of a Solidity smart contract.

Defeasible logic is a non-monotonic logic which allows for so-called defeasible propositions, in other words, exceptions to some of the rules. Such specifications contain a statement that regulates superiority between the clauses of a contract, which helps to handle possible disputes and breaches. The authors of [93] encode a smart legal contract that regulates a product license evaluation into a set of logical rules in *Formal Contract Logic* (FCL)—a formal contractual language based on defeasible logic and deontic logic of violations. The authors note that such logic-based implementation can be seen as an executable specification and, hence, facilitates monitoring.

Other Logics. We found several approaches that use other types of logic to describe contract-level requirements for smart contracts. For example, Li et al. [133] analyze traces of a process-algebraic contract model based on properties written in a fragment of (temporal) FOL. Temporal properties over traces of a *Scilla* smart contract are formalized using Coq higher-order logic predicates and a proposed temporal connective [185].

Unal et al. [206] use a *formal role-based access control (FPM-RBAC)* security model to specify the behavior of a crypto-wallet smart contract in a 5G network environment. The specification consists of role-based access control statements, generic constraints expressed in predicate logic, and *ambient logic* constraints for location- and time-based access policy.

3.1.2 Program-Level Specification. A dual type of smart contract specification is defined at the program level. In this section, we concentrate on Hoare-style specifications, which are often used together with program logics (C.f. Sect. 2.2.3) or are instrumented in the actual smart contract source code. Then, we review specifications that rely on predefined patterns of instructions and events observed during smart contract execution.

Hoare-Style Properties. A *Hoare-style* property of a program c is captured by a triple $\models \{P\}c\{Q\}$, where P and Q are predicates on the state of a program and its execution environment [18]. P and Q describe the state before and after the execution of c (if it terminates), respectively, and are referred to as a *precondition* and a *postcondition*. An *invariant* here is a predicate preserved by the execution of program c . Generally, Hoare triples express *partial correctness*, while *total correctness* requires proof of program termination. When applied to smart contracts, partial and total correctness are often considered equivalent due to the gas mechanism of Ethereum: a contract is guaranteed to terminate either successfully or due to an out-of-gas exception [18, 89].

A Hoare-style specification of a smart contract consists of preconditions, postconditions, and invariants defined for its functions, loops, and the contract in general. Predicates in these statements include storage and memory variables, environmental variables capturing block or transaction information, events and operations emitted during the execution. Hoare-style properties, therefore, can capture both vulnerable conditions and semantic correctness of smart contract functions. For example, access control in smart contracts is often realized via a precondition check on the identity of a transaction sender. As another example, postconditions are used to check if the constructor has initialized the contract state properly [46, 215]. Hoare-style properties can be specified using program logics (C.f. Sect. 2.2.3) and are often supported by verification languages, such as Why3 [156, 231]. Alternatively, Hoare-style specifications can be defined as smart contract

assertions, that are analyzed by verification tools statically [12, 14, 17, 101, 118, 167, 190, 215] or in runtime [46, 130, 193].

Many program logics listed in Sect. 2.2.3 enable Hoare-style reasoning about smart contracts in theorem provers [18, 40, 132, 226]. The work in [18] demonstrates how this approach is used to describe correctness of an escrow smart contract. To consider all possible invocations, the authors also verify that any input that is not covered by a specification leads to a transaction being reverted.

Nehai et al. [156] define a Hoare-style specification of smart contract functions in the polymorphic FOL of Why3—an instrument for deductive verification. The specification contains requirements for both semantic correctness and lower-level operational details, such as gas consumption and the ordering of array elements. Other techniques based on deductive verification [11, 37, 38] use the KeY framework. In this case, the input consists of a Java smart contract annotated with *Java Modeling Language* (JML) properties expressing pre-/postconditions, invariants, and modifiable variables of contract functions. The authors of [38], however, note the inability of their approach to handle complex properties, including temporal, even if they are expressible in JML. JML specification is also supported by RaziEl [179]—a framework for the implementation of *proof-carrying* smart contracts. RaziEl supports JML for functional correctness annotations and *Obliv-Java*—for privacy variable annotations. A proof-carrying code proposal for safe smart contract upgrade by Dickerson et al. [73] also relies on a Hoare-triple specification.

Assertions that correspond to Hoare-style properties can be added to a smart contract by means of Solidity. The language includes `require` and `assert` statements that can express a precondition or an invariant and postcondition, respectively. `require` statements are widely used to validate the input received from a user—the absence of such checks is known to be a source of serious issues [51, 97]. To advance further program-based specification, several tools provide custom specification languages [46, 101, 130, 131, 167, 192, 193, 215]. They support quantifiers [130], the implication operator [46, 167], and auxiliary functions, such as `sum`, which refers to a sum of values stored in a mapping [101, 130, 167]. The latter is useful in specification of financial smart contract invariants, as shown in Sect. 3.2.3. To check semantic correctness of a smart contract, VeriSol [215] translates a state-machine policy of a Solidity smart contract into Hoare-style source code assertions. It can, therefore, identify an incorrect state transition and an incorrect initial state in a smart contract. Apart from user-defined specifications, several tools detect smart contract vulnerabilities by adding assertions to a smart contract automatically [12, 190]. Overall, instrumentation of a smart contract code with the corresponding checks introduces a transparent and consistent approach to smart contract specification. However, while intended to be side-effect free expressions, source-code assertions generate executable code and, therefore, gas consumption.

Program Path-Level Patterns. Symbolic execution represents each program trace as a propositional formula over symbolic inputs to a program required to follow this execution path. Based on a contract CFG, this approach allows recognizing traces that contain problematic sequences of program states and instructions. In smart contract verification, the existing applications indicate if the contract is vulnerable [52, 81, 106, 141, 153, 160, 172, 233], gas-inefficient [58], or a honeypot [204]. Some authors encode patterns over instructions and flow- and data-dependency in verification languages, such as *Datalog*, to certify contract security [50, 205] including the absence of gas-related vulnerabilities [95]. Figure 7 presents patterns that correspond to *Unchecked Send* and *Accessible Selfdestruct* [50] vulnerabilities. The topmost pattern detects a vulnerability if the result of an external call is not used to update a blockchain state or throw an exception, while the bottom one detects a `selfdestruct` operation that can be called by an arbitrary account (C.f. Sect. 3.2.1). The code snippets illustrating both vulnerabilities are shown in Fig. 9.

```

.decl uncheckedCall(u:Statement)

uncheckedCall(u) :-
  callResult(_, u),
  !checkedCallThrows(u),
  !checkedCallStateUpdate(u).

(a) Unchecked Send

.decl destroyable(stmt:Statement)

destroyable(stmt) :-
  !inaccessible(stmt),
  op(stmt, "SELFDSTRUCT").

(b) Accessible Selfdestruct

```

Fig. 7. Patterns for “Unchecked Send” and “Accessible Selfdestruct” vulnerabilities detected by Vandal in [50].




Fig. 8. An automaton for identification of reentrancy during execution. Reproduced from [134] with permission.

While the pattern-based approaches discussed above operate on an EVM compilation of a smart contract, Perez and Livshits [166] use *Datalog* to formulate vulnerable patterns over execution traces retrieved from the blockchain. During fuzzing, ReGuard [134] compares an execution trace of a smart contract to an automaton that captures reentrancy. The automaton is shown in Fig. 8, where `entry` and `exit` correspond to an entry and exit of a function, and `CALL` is the corresponding EVM instruction. Another common dynamic vulnerability detection technique identifies predefined patterns of instructions in runtime [85, 142, 175].

Another type of path-level specification is based on events. Since the ERC20 standard specifies the events emitted by contract functions, several techniques perform compliance checking considering events too [60, 101]. Events emitted by a smart contract during its execution are analyzed for compliance with business rules by Fournier et al. [82] and Molina-Jimenez et al. [151]. The former specification ensures acceptable transportation conditions in a supply chain use case, while the latter encodes clauses of a selling contract in terms of parties’ responsibilities. Shishkin [187] formalizes functional requirements of an Ethereum contract as predicates over traces of events.

In addition, some publications adopt a different notion of events. For example, EthRacer [120] identifies event-ordering bugs in smart contracts by comparing results of different permutations of function invocations. Therefore, in this work, an event refers to a function call. The authors of [98] verify serializability of smart contract execution by analyzing its traces of events, or transitions between contract states. The corresponding property of being *effectively callback free* (ECF) is satisfied if for all executions containing a callback there exists an equivalent execution without callback that can start in the same state and end in the same final state. A vulnerable pattern of events corresponding to actions performed by a CSP model (C.f. Sect. 2.1.1) of a smart contract helps to detect an attack in its execution [171].

3.2 Properties Classification by Domains

Using the listed formalisms, specifications describe smart contracts’ functional correctness from different perspectives. The considered aspects include the absence of vulnerabilities, respect for privacy requirements, reasonable resource consumption, conformance to business-level rules, fairness to users. This section begins with an overview of existing classifications of domain-specific smart contract properties. Then, based on the property domains, we propose our categorization of smart contract properties that we have frequently seen in the research literature.

Existing Smart Contract Property Classification. Several taxonomies were proposed for smart contract vulnerabilities [26, 55, 97, 181, 224], with the first and, arguably, the most influential published in 2017 [26]. Section 3.2.1 overviews security properties that serve to protect a smart contract from types of vulnerabilities included in these classifications.

Two other publications describe common classes of domain-specific properties for smart contracts [39, 167]. These articles adopt dual forms of property specification: Hoare triples and temporal properties, respectively. Nevertheless, their categorizations share several types, such as *user-based access control* over contract functions and *invariants over aggregated values of collections*. In addition, Permenev et al. [167] define transition-system types of properties for smart contracts and extend invariants to a multi-contract setting. Some of these properties also resemble software design patterns [218], which capture the recurrent problems in smart contract development, e.g., *State Machine*, *Access Restriction*, *Ownership* [218]. Bernardi et al. [39], on the other hand, consider two types of properties related to a cryptocurrency exchange. We review some of these properties in Sect. 3.2.3.

Our domain-based analysis of properties is also shaped by a catalog of smart contract application domains proposed by Bartoletti et al. [34]. They divide smart contracts into *Financial*, *Notary*, *Game*, *Wallet*, and *Library* types. In Sect. 3.2, we describe our findings on the properties that are common for some of these types of smart contracts.

3.2.1 Security. Detection of security flaws in smart contracts received a lot of attention. Important security properties studied by the research community include *liquidity* (“a non-zero contract balance is always eventually transferred to some participants [36, 160]”), *atomicity* (“if one part of the transaction fails, then the entire transaction fails and the state is left unchanged” [141]), *single-entrancy* (“the contract cannot perform any more calls once it has been reentered” [182]), *independence of the mutable state and miner-controlled parameters* [96, 213]. Security guarantees of a smart contract also include proper *access control* [160, 205] for safety-critical operations, *arithmetic correctness* [79, 118, 190], and *reasonable resource consumption* [42, 58, 95]. The similarity of a smart contract to a concurrent object [184] suggests that some security guarantees can be derived from checking for *linearizability* [120] and *serializability*[98] of its executions.

The violation of these properties leads to many well-known smart contract vulnerabilities. For example, contracts violating liquidity are called “*greedy*”, implying that their funds cannot be released under any conditions [160]. A smart contract violates atomicity if it does not check the return value of a (possibly failed) external call operation, which corresponds to the *unchecked call* vulnerability [96, 166]. The execution logic of a smart contract that is not independent of environmental variables, e.g., `block.timestamp`, is prone to *dependence manipulation* [213] by a miner, including the *timestamp dependency* vulnerability [141]. If the business logic of a smart contract depends on its mutable state parameters, such as balance and storage, then it has the *transaction-ordering dependence* (TOD) problem. TOD belongs to a class of *event-ordering* (EO) bugs, which arise from the nondeterministic execution of blockchain transactions [120] and are manifested by different outputs produced by different orderings of the same set of function invocations [120, 213]. The authors of [120] use the notion of linearizability to identify EO bugs, including mishandled asynchronous callbacks from an off-chain oracle. The violation of single-entrancy implies that a smart contract is *reentrant*: provided with enough gas, an external callee can repeatedly call back into a smart contract within a single transaction. In execution traces, reentrancy is also identified as a violation of a serializability-like property [98], i.e., if the result of an execution with external callbacks is not equivalent to that of an execution without. Missing permission checks for the execution of a transfer or a self-destruct operation make a smart contract *prodigal* and *suicidal* (Fig. 9), respectively [160]. The identification of a transaction sender through the

```

1  contract KotET {
2  ...
3  function withdraw(address to,uint amount) public {
4  to.send(amount); //Unchecked Send
5  }
6  function close() public {
7  selfdestruct(msg.sender); //Public Selfdestruct
8  }
9  }

```

Fig. 9. “Unchecked Send” and “Accessible Selfdestruct” vulnerabilities [211].

```

1  contract UnderflowAttack {
2  ...
3  function withdraw (uint amount) public {
4  require(balances[msg.sender] - amount > 0);
5  msg.sender.transfer(amount);
6  balances[msg.sender] -= amount; //Underflow
7  }
8  ...
9  }

```

Fig. 10. “Integer Underflow” vulnerability [99].

`tx.origin` variable is also considered vulnerable [50]. Integer arithmetics of EVM and absence of automatic checks for arithmetical correctness make Ethereum contracts prone to *integer overflow and underflow* [79, 190]. Figure 10 presents smart contract code containing the *integer underflow* vulnerability. Furthermore, the progress of a smart contract can be compromised by gas-exhaustive code patterns [58, 95].

Many symbolic execution tools detect vulnerabilities by scanning for vulnerable patterns in path conditions and traces derived from a contract CFG. For example, to identify *timestamp dependency*, Oyente [141] checks if a path condition of a trace includes the timestamp variable. To detect *suicidal* smart contracts, sCompile [52] checks whether, for the paths that contain `selfdestruct`, the path conditions constrain on the timestamp, block number, or the address of the contract owner. Maian [160] flags traces which transfer funds or control to an address not appearing in a contract state, as *prodigal*. Scanning for vulnerable instruction patterns can also be done in runtime, for example, to identify that a permission check relies on `tx.origin` [57].

He et al. [105] consider prodigal and suicidal contracts to be the most challenging to detect, as it requires an ordered set of transactions. Indeed, if one sets the owner variable of a smart contract equal to his address, he can bypass the aforementioned checks in the following transactions. The authors of [49, 51] call such vulnerabilities *composite* and detect them using information-flow patterns. Other static analysis tools also detect patterns over the data- and control-flow of a smart contract, capturing the necessary context-sensitive information. Figure 7(a) demonstrates a pattern for an *unchecked call* vulnerability detection used by Vandal [50]. The pattern states that the return value of the call operation should be used either to control the execution of the throw operation or to perform storage writes. To prevent *overconsumption of gas*, static analyses infer the upper bound on gas consumed by a function call [13, 42, 146, 156] or detect vulnerable patterns, such as unreachable code [58], trivial conditions [17] or an unbounded number of elements [95] in loops. Race conditions, such as TOD, are identified by static analyses as read-write hazards between the executions of two [96, 118, 205] or several [213] functions. An EO bug caused by an asynchronous callback could as well be prevented by a temporal property proposed in [185]: it demands the state of a contract to remain unchanged until the response from an oracle is received.

As a liveness property, *liquidity* can be intuitively formulated in temporal logics [36]. Bartoletti et al. [36] additionally extend its definitions to user strategies: the default, *strategyless multiparty liquidity*, assumes the existence of a cooperative strategy to unfreeze funds. Liquidity can also be compromised by DoS attacks, for example, Parity Multisig Wallet [200] lost its ability to pay users as it delegated transfers to a smart contract that was destroyed. Nelaturu et al. [158] identify the Parity issue as a violation of deadlock-freedom.

Since arithmetic issues are often caused by non-validated function arguments, their identification is entrusted to pre- and postconditions [108, 119, 156, 165] and assertions [79, 190]. The (formally certified [182, 231]) SafeMath [162] Solidity library implements similar checks via the `require` statements. Source-code assertions are also a popular way to implement role-based *access control* in smart contracts. For example, it is done via equality checks between the values of `msg.sender`

(the usage of `tx.origin` is considered vulnerable [50, 57]) and the corresponding contract variable, e.g., `owner` [101, 167, 215, 217]. Functions with restricted access usually change the value of a safety-critical variable, such as `owner`, or manipulate funds, e.g., `transfer` or `mint` tokens.

Reentrancy is one of the most studied smart contract vulnerabilities. Both static and dynamic data- and control-flow analyses detect the violation of “*no writes after call*” and similar patterns [79, 80, 131, 140, 175, 205] (yet, this pattern is neither sound nor complete approximation of reentrancy [182] and is considered a standalone property in [205]). A similar check is specified in terms of contract operations in CTL [147]. While *single-entrancy* is verified statically using Horn clauses in [182], a similar property is also defined in terms of instructions in the execution trace [134] (Fig. 8).

Finally, the inconsistency caused by reentrancy as well as other issues, such as flawed arithmetics [130], can be detected through invariants. Such invariants are proposed for financial [109, 133, 167, 187] (C.f. Sect. 3.2.3) and voting contracts [159] (C.f. Sect. 3.2.4).

3.2.2 Privacy. Smart contracts that implement auctions, games, and lotteries select a winner among the users who submitted their bids or moves. Blockchain transactions are publicly accessible, so a malicious user can decide on his move based on the actions of other contract participants. In a name registration smart contract, one can intercept the name chosen by another user from an unconfirmed transaction and register his address under this name first [10].

To prevent exposure of user inputs in a smart contract by design, it is recommended to implement cryptographic protocols, such as *commitment* [70] and *timed commitment* [35] schemes. The protocol requires each participant to commit to a secret that corresponds to his input and should be revealed in time—otherwise, the participant has to pay compensation. In general, the correctness of a commitment smart contract is described in terms of the knowledge of each others’ secrets that participants did or did not obtain as well as the corresponding changes in their balance. One example of such a property is proposed by the authors of [20]: “*if an honest Bob did not learn Alice’s secret then he gained Alice’s deposit as a result of the execution*”. Similar aspects are considered in specifications of *atomic-swap* contracts, which allow two parties to exchange their assets *atomically*, i.e., the swap either succeeds or fails for both parties [109, 207]. Alternatively, to prevent contract participants from cheating, Sergey et al. [185] suggested formulating a property as a knowledge argument over the prefix of observed execution history.

The transparency of the transactions and ledger data also makes smart contracts less applicable to privacy-sensitive domains, such as healthcare and voting. Potential solutions for privacy issues in smart contracts include the application of trusted management [122] or hardware [63, 229]. On the whole, privacy-preserving smart contract is an independent topic of intense study, not covered by this survey. Among these approaches, we distinguish two that enforce privacy guarantees by allowing the corresponding annotations in smart contract code. To prevent data leaks, a smart contract language zkay [192] supports variable annotations, which specify who is eligible to read the value of a variable. Similar privacy annotations are supported by the Razi el [179] framework. As a demonstration, a smart contract is implemented in the Obliv-Java programming language with both privacy and correctness requirement annotations. Both tools utilize some zero-knowledge proof protocols to certify that privacy is maintained in a smart contract before it is executed.

3.2.3 Finance. One of the key advantages of smart contracts is their ability to manipulate digital assets, known as cryptocurrency. Smart contracts automate the execution of financial applications that involve storage and transfer of funds, between users and contracts. Common examples include wallets, banks, escrows, cryptocurrency exchanges, and crowdfunding campaigns.

Initial Coin Offering (ICO) is a method of raising funds by selling *tokens*—programmable assets managed by smart contracts on blockchain platforms [61]. Analysis performed by Chen et al. [61] demonstrates that 80% of ICOs are Ethereum-based smart contracts implemented according to the

ERC20 standard [209]. The standard describes an interface of a token contract, i.e., its functions and events. Arguably, even this semi-formal specification gave rise to the application of formal techniques to token contracts [29, 60, 73, 103, 130–133, 163, 167, 190, 196, 208] and development of a complete formal specification of ERC20 in \mathbb{K} [115]. Several tools focus on consistency checking between a token implementation and a standard [60, 130, 131]. For example, TokenScope [60] identifies inconsistency in the histories of smart contract executions, while Li et al. [130] formulate a relevant set of invariants for several standards. Among other properties, several more tools [57, 103, 208] check whether the contract emits the events determined by a specification, which can be expressed in a specification language proposed by Hajdu et al. [103].

Another common way to ensure validity of a token or other similar financial contract implementations is through correctness conditions for functions related to token transfers, e.g., `transfer()`, `transferFrom()`, and `approve()` in an ERC20 contract. For example, preconditions and postconditions of the `transfer()` function typically assert the following statements: “*the sender’s balance is not less than the requested amount of tokens*” and “*the balances of both sender and receiver are updated according to the provided amount*” [11, 29, 118, 131, 132, 138, 165, 193, 196, 208]. To cover the corner cases, some of the listed specifications explicitly state that *balances of non-participating users remain unchanged* [193, 196]. Sun et al. [196] additionally check that *the value to be transferred is greater than zero*, a crowdfunding specified by Kalra et al. [118] only *accepts investments bigger than a threshold limit*, while Dickerson et al. [73] require *user balances to be non-negative*. The latter property can be found in specifications of banking contracts [11, 119], while wallets usually require the opposite: the authors of [52, 118] request a user to *respect the limit* of Ether that can be transferred out of a contract within a transaction or a contract lifetime, respectively. A wallet specification by Kalra et al. [118] does not permit users to transfer funds to themselves, while ERC20-K [115] allows self-transfers but considers it a special case.

Formal specifications of financial contracts often contain vulnerability-related requirements, primarily, integer overflow and underflow that may occur during a transfer [11, 101, 119, 190, 196, 208]. Other potential issues include event-ordering bugs: the authors of [73, 120] show how an unforeseen ordering between the invocations of `approve()` and `transfer()` allows a malicious user to spend more funds than intended, and how it can be prevented by an invariant. Wang et al. [213] detect an extended set of issues referred to as nondeterministic payment bugs, which include dependence manipulation conditions. Other specifications address access control [206, 208, 217], and sending funds to invalid addresses [52, 57]. To assert that a crowdfunding contract can *pay back all investors if the campaign is not successful*, i.e., is not greedy, some specifications introduce temporal properties [167, 185]. To ensure that a smart contract has enough funds to do so, Permenev et al [167] require that *the escrow’s balance must be at least the sum of investor deposits unless the crowdsale is declared successful*. Other temporal properties ensure that funds are raised only until a cap [101] or a deadline is reached [167] or specify the order of states in a crowdfunding lifecycle and a corresponding payment procedure [167, 197].

An encouraged way to ensure the correctness of financial operations is by specification of invariants over a smart contract balance and an aggregation of user balances [101, 130, 133, 167, 196]. The most common invariant states that *the sum of user balances is equal to tokenSupply* [101, 130–132, 167, 187, 190, 196, 208], where the balances of users are stored in a so-called bookkeeping mapping variable, e.g., `balances`, and `tokenSupply` defines the total number of existing tokens. Other token invariants establish *the equivalence between sums of sender and receiver balances before and after the function execution* [17] or require *the difference between the contract balance and the sum of all participants’ bookkeeping balances to remain constant* [210]. These invariants can also certify correctness of inter-contract communications of a smart contract [51, 167]. For example, an invariant proposed by Wang et al. [211] ensures correctness for an account that receives funds from

a smart contract: “*the amount deducted from a contract’s bookkeeping balances is always deposited into the recipient’s account*”. To facilitate specification and verification of these invariants, several tools identify bookkeeping variables automatically [60, 210] or provide support for a special *sum* function, that can be applied to mappings and/or arrays [101, 167, 190]. Invariants are also used to express that the value of `tokenSupply` is immutable, bounded, or does not decrease [17, 39, 132, 167, 208]. In addition, Bernardi et al. [39] consider invariants that apply to cryptocurrency exchanges, such as proportional distribution of tokens: “*one should not be able to exchange an asset worth nothing with a positively valued asset*”.

Fairness in financial smart contracts can be analyzed by statistical means: the authors of [43] estimate the probability of money losses for a participant of a goods exchange protocol. In a game-theoretical analysis of a crowdfunding smart contract, Chatterjee et al. [54] define the amount of sold tokens to be an objective function. If its value can exceed the number of available tokens, the contract is declared unfair.

The authors of [25, 126] define a set of properties for financial legal smart contracts, such as derivatives. The authors of Marlowe [126] verify that all smart contracts written in this language *conserve funds*, i.e., *the money that comes in plus the money in the contract before the transaction must be equal to the money that comes out plus the contract after the transaction, except in the case of an error*. A similar idea that money are neither created nor destroyed by a smart contract is addressed in Move [45] by using linear resource types for assets. In Findel [25], the requirements for financial derivatives include the following: “*derivatives should be free of calculation mistakes*”, “*errors caused by external sources should be handled properly*”, “*accidental swaps should be avoided (i.e., the generated transactions and contracts have the intended issuer and the intended owner)*”.

3.2.4 Social Games. We associate the term *social game* with a class of smart contracts that regulate the participation and interaction of multiple users, according to some predefined game rules. Examples of such applications include auctions [22, 54, 147, 185, 197], voting schemes [22], lotteries [54], and games, such as gambling [11, 76, 118], rock-paper-scissors [37, 54, 185], and Ponzi-like games [147]. The payout to each participant is automatically determined by the contract source code, making the analysis of *fairness* even more relevant. Fairness of social games can be compromised due to privacy issues which occur if participants see each others’ submissions on blockchains (C.f. Sect. 3.2.2). In general, fairness specifications require user submissions to be valid and properly recorded, after which a correctly determined winner receives his payout. While its precise definition is subjective in the presence of several competing users [118], game-theoretical techniques [54, 128] analyze utility functions of participants as a reflection of contract fairness. The theory of mechanism design has also been applied in the verification of fairness properties of game contracts [136].

Social games typically require each user submission to be accompanied by a participation fee, be it a vote in a ballot [40], a bid in an auction [22, 54, 185], or a move in a game [118] (e.g., “*a minimum deposit is required to play the game*” [118]). In addition to that, users should have the permission to participate (“*people may vote if they are registered voters*” [84]), the choice submitted by a user should belong to a set of possible values (“*the vote goes for a registered candidate*” [40]) and, in some cases, *submitted only once* [118, 185]. Temporal properties are often used to assure that submissions are made before a deadline (“*people must not vote after the deadline*” [84]) or before the results are finalized (“*bid cannot happen after close*” [147]). Valid user inputs are to be correctly recorded by a smart contract [40, 185]. As an example, in an auction contract specified by Sergey et al. [185], the `pendingReturns` variable should store a mapping between each account and the sum of all his transfers made to a contract. A specification of a rock-paper-scissors contract by Sergey et al. [185] summarizes several of the listed requirements: “*each player can only submit*

their non-trivial choice once, and this choice will have to be a key from payoffMatrix in order to be recorded in the corresponding contract field". Correctness of recording in voting contracts can also be guaranteed through the use of invariants over the number of votes [22] and individual user weights [130]. For instance, Antonino et al. [22] check that *the number of votes initially available is equal to the sum of votes still to be cast plus votes already cast throughout the election process*.

After user inputs are collected, it is the responsibility of a smart contract to determine the winner and perform the relevant transfers of cryptocurrency. The developers of smart contracts that choose a winner randomly, such as lotteries [28] and casino [120], should be aware of the risks introduced by dependence on external oracles or randomization techniques, as discussed in Sect. 3.2.1. To enforce a correct winner selection in a rock-paper-scissors contract, Beckert et al. [37] encode the game rules in a postcondition of the corresponding function. The property of a simple Ponzi-scheme contract, King of the Ether Throne [1], certifies that *a user should always be crowned after successfully sending Ether to the contract* [147].

Upon the outcome finalization, most of the social game contracts transfer funds to some user, whether he is a successful gambler [76] or an unsuccessful auction participant [185]. In that sense, social games are similar to financial contracts with their properties and potential flaws. For example, the specification of an auction contract proposed by Sergey et al. [185] ensures that the smart contract does not send the funds arbitrarily and is able to pay out the unsuccessful participants: *"an account, which is not the higher bidder, should be able to retrieve the full amount of their bids from the contract, and do it exactly once"*. A similar liveness property is formulated by Shishkin [187] for a DAO-like smart contract, stating that *if some investor deposited some amount of money and did not vote for any proposal afterward, then he will always be able to get a refund by calling the corresponding function*. There are also specifications describing different ways to guarantee that a smart contract is not greedy and has enough funds to perform the required payments. For example, the owner of a casino in [76] is only allowed to *withdraw from the pot when the bet of a player is resolved*. According to Beckert et al. [37], *as long as the auction remains open, the sum of the funds in the auction remains the same or increases but never decreases*.

Smart contract specifications can also describe the principal rules of its execution. For example, temporal properties that define the order of state transitions in an auction contract are shown in Fig. 6. Another auction specification [38] states that *the items belong to the auctioneer as long as the auction is open and to the highest bidder after the auction is closed*. Invariants for an auction contract in [136] require that *the bidder with the highest bid becomes the winner and the highest bid becomes the clear price*.

3.2.5 Asset Tracking. Supply chain is named to be one of the most suitable use cases for permissioned blockchains [82]. A supply chain typically involves several interacting parties responsible for production, transportation, and sales. Alqahtani et al. [16] describe the responsibilities of parties in a fuel supply chain, e.g., *"vendors must always maintain their production under the daily limit and log it into the ledger"*, *"point-of-sales must always adhere to the price set by the regulators"*, and *"once the point-of-sales receives oil, a payment should eventually be triggered"*. Another common requirement for supply chain contracts defines the acceptable transportation conditions, mainly, restrictions on the temperature [82] and humidity [215, 236]. The authors of [215, 236] transfer the contract into a failed state if a critical reading is received, while Fournier et al. [82] report a violation if the amount of such readings exceeds a threshold within an hour.

The next group of asset-management smart contracts includes marketplaces, often the ones that facilitate trading between suppliers and producers of energy resources [147, 156, 157]. Similar to the properties of contracts from other domains, safety properties describe the correct execution flow of a market by declaring the allowed and required chains of actions [147, 157], e.g., *"if a user's*

payment has been received, then his bill is edited before the market closure". Each marketplace contract implements an algorithm that matches the received buying and selling offers [147, 156, 157]. Nehai et al. [157] warrant a satisfactory solution by verifying that *the algorithm implements a repartition of energy proportional to payment* and that *once opened, the market will eventually be closed*. Properties from their later proposal [156] guarantee that the solution is optimal in the sense that it maximizes the total number of tokens exchanged.

A common example of legal smart contracts is a license agreement [92] for product evaluation. Its formal specifications [93, 111, 197] describe the rights, obligations, and prohibitions of a user who participates in the evaluation. For example, one of the clauses states that *the Licensee must not publish comments on the evaluation of the Product, unless the Licensee is permitted to publish the results of the evaluation*. As discussed in Sect. 2.1, logic-based languages are practical for capturing the legal modality of such contracts [93, 111]. Still, Suvorov et al. [197] encoded some of the clauses in temporal logic—although there is no direct representation of deontic modalities, the properties are still sufficient to identify a violation of certain contractual clauses. This approach also enables description of common principles of a contract operation, e.g., that *removal of a comment cannot happen if nothing has been published*. In general, formal representations of rights and obligations in smart legal contracts were discussed with regard to various formalisms including deontic logic [30], automata-based specifications [29], a domain-specific language based on LDL_f [180], set-based frameworks [33, 222], and some others. The authors of [151] verify the formal contractual model of a data selling smart contract against typical contractual problems, such as clause duplication.

Several formal specifications exist for name registration smart contracts that manage domain names or provide aliases to blockchain account addresses [10, 107, 145]. Most of them focus on security analysis [10, 107], e.g., estimating the chance of a stealing attack [10]. On the contrary, Maksimov et al. [145] formulate a property regarding the functional correctness of a smart contract: *"the user can get an alias or a rejection after a number of attempts not exceeding a specified value"*.

3.3 Summary

Smart contract properties express security requirements and semantic conformance in many forms. Vulnerabilities are often detected through path-level patterns in control- and data-flow of smart contract execution. In program-level representations, both vulnerable conditions and semantic correctness of smart contract functions can be captured by Hoare-style properties. The corresponding checks can as well be instrumented into the smart contract code, which, however, increases the amount of gas consumed by a smart contract. Nevertheless, temporal logic is considered a more suitable formalism to specify the multi-agent execution environment of blockchain [207]. Indeed, contract-level behavioral specifications usually employ one of temporal logics to formulate temporal and probabilistic constraints and account for user strategies. The analysis of smart contract properties outlined in Sect. 3.2 also proves that the execution logic of smart contracts often relies on temporal conditions. However, as has been observed in Sect. 2.3, contract-level modeling and specification techniques seldom consider low-level execution details.

4 SMART CONTRACT FORMAL VERIFICATION TECHNIQUES

This section discusses techniques, tools, and frameworks that are employed for verification of smart contract models described in Sect. 2. We also discuss the capabilities of these verification techniques with respect to the properties outlined in Sect. 3. A partial summary of verification techniques and employed tools is provided in Table 1.

Table 1. A (partial) summary of formal verification tools.

| Verification Techniques | Tools | Selected References | Total |
|--------------------------------|--|---------------------------------------|-------|
| Model Checking | NuSMV, nuXmv | [16, 121, 143, 147, 157, 158, 206] | 18 |
| | SPIN | [32, 151, 163] | |
| | CPN | [74, 138] | |
| | BIP-SMC | [10, 145] | |
| | PRISM | [43] | |
| | UPPAAL | [20] | |
| | MCK | [207] | |
| | Maude | [28] | |
| | FDR | [171] | |
| Ambient Calculus | [206] | | |
| LDLf | [180] | | |
| Theorem Proving | Coq | [21, 25, 40, 159, 161, 185, 196, 226] | 20 |
| | Isabelle/HOL | [18, 89, 107, 108, 126, 132] | |
| | Agda | [53] | |
| Program Verification | Datalog & Soufflé | [49, 50, 95, 166, 205] | 42 |
| | Boogie Verifier & Corral | [22, 101, 215] | |
| | LLVM & SMACK | [118, 213] | |
| | SeaHorn | [14, 118] | |
| | F* | [42, 96] | |
| | KeY | [11, 37, 38] | |
| | Why3 | [68, 156, 231] | |
| | \mathbb{K} framework | [117, 119, 165] | |
| Custom static analyses | [17, 77, 83, 126, 146, 182, 186, 187, 190] | | |
| Symbolic & Concolic Execution | Oyente | [81, 141, 204, 233] | 22 |
| | Mythril | [153, 170, 217] | |
| | teEther | [123] | |
| | Maian | [160] | |
| Manticore | [152] | | |
| Runtime Verification & Testing | ContractLarva | [29, 76] | 26 |
| | EVM* | [142] | |
| | ReGuard | [134] | |
| | SODA | [57] | |
| | Solythesis | [130] | |
| ECFChecker | [98] | | |

4.1 Model Checking

Model checking is a well-established technique for automatically verifying a system model (with finite states) against its specification. When applied to smart contracts, model checkers perform verification of contract-level models, mainly transition systems, against a temporal logic specification [10, 16, 28, 43, 121, 143, 145, 147, 157, 207]. We identified 18 verification techniques that are based on model-checking, with 7 utilizing the capabilities of NuSMV and nuXmv model checkers. The popularity of model checking is arguably caused by the suitability of both modeling and specification formalisms to smart contracts description combined with the existence of established automatic frameworks. Although model checking is successful in verifying systems of several smart contracts or users [16, 143, 158], its limitations are induced by the input language of a model checker and the state explosion problem [157].

Support of diverse transition systems and temporal properties (C.f. Sects. 2.1.2 and 3.1.1) help model checking capture different characteristics of smart contract execution, such as concurrency [74, 138, 163, 171], nondeterminism [10, 43, 145, 207], or time constraints [20]. In addition to verifying the functional correctness of one contract [121, 147, 157], model checking handles systems of interacting smart contracts [16, 158] and users [143]. Furthermore, with specifications expressed in temporal logic, model checking is able to verify liveness properties and properties of progress, e.g., liquidity [36]. Model checkers, such as SPIN [110] or PAT [137, 194], additionally verify conventional requirements of concurrent systems, such as deadlock- and livelock-freedom [32, 151].

However, model-checking suffers from the state explosion problem, which requires the users to apply abstraction techniques for smart contracts [28, 36] or assume a set of simplifications to its execution [121, 157]. Furthermore, Nehai et al. [157] claim that precise modeling of the blockchain environment is infeasible due to the limitations of the NuSMV model-checker input language. Indeed, as it checks a high-level representation, model-checking approaches rarely consider the details of smart contract execution on blockchain, such as the gas mechanism or a memory model. One of the two identified exceptions includes a schematic model of blockchain behavior, such as mining of transactions in blocks [10]. A more realistic model of smart contract execution is achieved by Duo et al. [74] by combining CPN with a program logic for EVM bytecode.

4.2 Theorem Proving

Verification based on theorem proving involves encoding a system and its desired properties into a particular mathematical logic. Then, a theorem prover attempts to derive a formal proof of satisfaction of these properties based on the axioms and inference rules of the formal system. Unlike model checking, which only verifies finite-state systems, theorem proving supports verification of infinite systems [177]. On the other hand, theorem-proving techniques are usually semi-automated and require human involvement and expertise. Nevertheless, this drawback discourages neither academia nor industry from developing tools and languages that are backed by theorem proving. Our study identified 20 applications of theorem proving used to derive correctness guarantees for smart contract languages, individual smart contracts, and even verification frameworks.

Coq, Isabelle/HOL, and Agda theorem provers are used to develop formal semantics of low- [18, 40, 53, 89, 109, 161], intermediate- [41, 132, 185], and high-level [117, 225] programming languages for smart contracts, including DSLs for financial contracts [25, 126]. However, Li et al. [132] suggest that, among them, intermediate-level languages are the most suitable for formal verification. Annenkov et al. [21] note that the authors of formal semantics often focus on meta-theory, i.e., verification of language properties, as in case of *Plutus Core* [53] or *Simplicity* [161]. Established through its formal semantics, general properties of *Marlowe* [126] smart contracts include a bound on transactions that a contract can accept and preservation of money within the contract.

ConCert [21] is a Coq-based framework, which allows both meta-theoretic and functional reasoning about a (functional) language and a smart contract, respectively. Together with other publications [18, 21, 25, 40, 126, 132, 159, 226], it illustrates how theorem proving helps to precisely describe and prove correctness conditions of smart contract execution. These conditions include Hoare-style correctness properties over the state of a smart contract and its environment [18, 40, 132, 226], security requirements [25, 196], and gas consumption reasoning [89].

Despite the potential expressiveness of theorem-proving approaches, they seldom consider inter-contract communication and temporal properties of smart contracts. An attempt to examine smart contract interactions in Coq is performed by Nielsen and Spitters [159], who prove a voting contract invariant which approximates a temporal property. Verification of temporal properties, including liveness, in *Scilla* smart contracts, is enabled by a formalization of its trace semantics [185], while an embedding of the language in Coq is under development [186].

4.3 Symbolic Execution

Symbolic execution executes a program symbolically and thus is able to explore multiple concrete execution paths at a time. Symbolic execution of both Ethereum and EOSIO smart contracts is based on a traversal of a CFG reconstructed from the bytecode of a smart contract [106, 141]. Thereby, this approach inherits the difficulties associated with the issues of precise CFG modeling in smart contracts (C.f. Sect. 2.2.2). Other difficulties encountered in symbolic execution for smart

contracts include heavy usage of hash functions (hard to be solved by SMT solvers) and the need to symbolically model memory and smart contract interactions. Similar to model checking, symbolic execution approaches usually explore paths up to a certain length, which calls for the application of abstract interpretation and partial-order reduction. Nevertheless, symbolic execution is one of the principal approaches to detect smart contract vulnerabilities: overall, we found 22 applications of symbolic and concolic executions. For solving path constraints, nine-tenths of these techniques rely on an off-the-shelf constraint solver Z3, however, Frank et al. [83] show that other solvers, especially Yices2, significantly outperforms Z3. Symbolic execution is also used to guide smart fuzzing approaches [105].

In most cases, symbolic execution is used to detect predefined vulnerable patterns in the bytecode level. For example, that is performed by primary symbolic-execution tools that include Oyente [141], Mythril [153], and their extensions [81, 170, 204, 233]. EOSafe [106] verifies EOSIO smart contracts against a set of predefined vulnerabilities. teEther [123] examines execution paths of interest that contain a vulnerable instruction pattern and generates exploits to confirm the vulnerabilities. To extend verification capabilities beyond predefined patterns, SmartScopy [79] and Solar [131] provide support for user-defined vulnerable patterns and constraints, respectively.

Inter-transactional verification helps to uncover smart contract vulnerabilities that are composite, i.e., the vulnerability can only be exploited through a sequence of transactions [49]. However, symbolic execution suffers from the path explosion problem, which makes it unsound and unable to explore all possible sequences due to complex constraints. Therefore, most of the techniques, e.g., Maian [160], bound the invocation depth—the length of a considered sequence of transactions. VerX [167] combines symbolic execution with abstract interpretation and models gas mechanics of Ethereum, which allows elimination of the paths that are unreachable due to an out-of-gas exception. The tool checks reachability of assertions inserted in the smart contract source code from the temporal properties defined by a user. VerX also supports verification of a bundle of interacting smart contracts, which is rarely addressed by static analyses.

4.4 Program Verification

Smart contract correctness and security are verified by a wide range of program verification techniques, which total to 42 publications. The most numerous category among these techniques employs existing verification languages supported by verification frameworks, e.g., Boogie, LLVM, Datalog, or F*. Many of program verification tools are automated [77, 182, 190, 205] and can check for composite vulnerabilities [49] and semantic correctness of a smart contract [215]. The issues associated with program verification (especially static analysis) include the need to precisely abstract components of smart contract execution and memory model [22, 102, 182], while the abstraction is often unsound [182]. Although static analysis tools often do not consider the gas mechanics of smart contract execution [22, 118, 182], others specifically focus on the estimation of gas consumption [13, 146].

Translation of the smart contract source code into verification languages enables the wide capabilities of their toolchains. For example, data- and control flow analyses by Datalog and its Soufflé engine successfully detect vulnerabilities in bytecode [49, 50, 95, 205] and execution traces of Ethereum [166] smart contracts. Translation of Solidity—a high-level language—to Boogie enables verification of Hoare-style semantic properties [22, 215] and secure compilation of annotated smart contracts [101]. LLVM was shown to be a suitable IR for EVM [213], Solidity [118], and Go [118] implementations to check for both vulnerabilities and semantic conformance. Bhargavan et al. [42] performed embedding of both a Solidity subset and EVM bytecode in F* to allow verification of both high-level correctness and low-level properties, e.g., related to gas consumption. Translation of Solidity [156], EVM [231], and Michelson [68] smart contracts to WhyML enables deductive

verification of Hoare-style properties. The authors of [156] note that this technique facilitates the evaluation of a smart contract in a more realistic setting compared to model checking [157], which suffers from state explosion. However, unlike model checking [157], it does not allow specification of liveness properties. Deductive verification of smart contracts written or translated to Java w.r.t. Hoare-style properties is performed by the KeY framework [11, 37, 38]. A correct-by-construction deductive verifier for EVM bytecode has been generated in the \mathbb{K} framework [165].

Although several of the above-listed approaches [118, 205] make soundness claims, Schneidewind et al. [182] demonstrate sources of unsoundness identified in these analyses. The authors of [182] propose a provably sound Horn-clause abstraction for EVM bytecode based on an EVM formalization [96] for reachability analysis, which enables verification of both security and functional Hoare-style properties. Other static analyzers support reasoning about execution traces [60] and smart contract written in Scilla [186], Solidity [17, 77, 187, 190] or EVM bytecode [182], and Tezla [174].

Some of the publications under study actively apply program verification tools that are specifically designed to operate over particular formalisms, such as set-based frameworks TLA+ [222] and Event-B [33, 234], a defeasible logic engine SPINdle [93] for logic-based programs, Tamarin prover for verification of concurrent protocols, C-language verifiers [14].

4.5 Runtime Verification and Testing

Runtime verification is a lightweight verification technique that checks the properties of a running program. Different from the techniques described in previous sections, runtime verification is concerned with one execution trace at a time. In the domain of smart contracts, the term *trace* often denotes a sequence of instructions executed by a blockchain platform, while it can also refer to a sequence of function invocations or events emitted by a smart contract. The availability of runtime information helps dynamic verification techniques mitigate one of the main obstacles to smart contract analysis—the need to model a complex execution environment of blockchain. Runtime verification usually provides a reactive defense against vulnerabilities or violations of correctness at runtime and can potentially identify vulnerable states that could not be reached by model checking or symbolic execution, due to the state or path explosion. Approaches based on runtime verification and testing, including fuzzing, account for 26 publications in our dataset.

Most of the runtime verification techniques perform verification on-chain, which allows them to reject vulnerable transactions acting as an online defense system. The approaches are based on enhancing the source code of a blockchain client [57, 80, 85, 98, 175] or a smart contract itself [12, 29, 30, 46, 130, 193]. Both Go and JavaScript Ethereum clients enhanced with monitoring strategies and/or taint analysis are able to detect predefined vulnerable patterns in the instruction trace [85, 98, 142, 175]. \mathcal{A} egis [80] and SODA [57] additionally support verification of user-defined patterns, which can, for example, indicate an inconsistency between a contract and a corresponding standard. Still, the discussion on the false positives produced by SODA [57] shows that although the tool can access runtime information, its completeness is still subject to the precision of the detection patterns. Another obstacle to wide adoption of such instrumented clients in Ethereum arises from the fact that it would require a hard fork of all EVM clients [130].

Another group of runtime approaches prevents vulnerabilities by inserting protection code into the source code of a smart contract [12, 29, 46, 130, 193]. For example, Solythesis [130] automatically instruments a Solidity smart contract with custom invariants. Compared to Solidity, its specification language for invariants contains additional features, including quantifiers and sums. Azzopardi et al. [30] instrument a smart contract with monitors that verify the compliance between a smart contract and a legal contract. However, the assertions generate additional gas consumption during the execution of an instrumented smart contract [130]. Another limitation of

this approach is caused by the infeasibility to formulate and verify complex temporal properties, including liveness, using such smart contract instrumentation. While the above-listed tools perform on-chain verification, events emitted during the execution of a Solidity or a Hyperledger Fabric smart contract can also be checked for compliance off-chain [82, 151].

Besides runtime verification, smart contracts can be dynamically verified via fuzzing and testing techniques to identify and exploit vulnerabilities [116, 134, 210]. To navigate a fuzzer towards higher path coverage, some authors combine fuzzing with taint analysis [230] and symbolic execution [105, 120, 203]. Symbolic execution is also used to guide invariant-based testing of a smart contract by a federated society of bots [208].

4.6 Other Techniques

Lightweight smart contract analyses include lexical and syntactical checks on Solidity source code and its tree-like structures [140, 201, 224], which detect vulnerable code patterns.

An emerging line of research is focused on the application of statistical and deep learning-based techniques to smart contract verification. One of the first proposals in this domain [135] identifies vulnerable smart contracts as containing irregular token sequences w.r.t. a statistical language model of Solidity. Tann et al. [199] use recurrent neural networks to measure the similarity of smart contract opcode to that of vulnerable smart contracts flagged by Maian [160]. Other machine- and deep-learning analyses are also used to identify vulnerable smart contracts [214], detect Ponzi schemes [62] and code clones [86], or propose code updates [113]. Additionally, machine-learning techniques help Chen et al. [56] identify whether an access control check in a smart contract function is missing.

4.7 Summary

The choice of a verification technique largely depends on the type of a formal model and specification that one aims to analyze. Temporal properties are typically verified by a model checker w.r.t. a contract-level model of one or several interacting smart contracts. Vulnerable patterns in the CFG of a smart contract are identified by symbolic execution techniques. Program verification approaches perform (often unsound) translation of a smart contract to a verification language, which enables data- and control-flow analyses to uncover composite vulnerabilities. Despite a high level of complexity associated with theorem proving, it is actively used to precisely specify and verify correctness of smart contracts (usually expressed in Hoare triples), programming languages, and verification frameworks. As an alternative to static analyses, a smart contract can be verified in runtime, which involves modification of a blockchain client or a smart contract itself. This approach enables rejection of non-compliant transactions in runtime but may introduce gas or runtime overhead.

5 DISCUSSIONS

In this section, we discuss our observations on the formal verification and specification literature of smart contracts, to formulate answers to the research questions raised in Sect. 1.1. Some key observations are also summarized in Table 2.

5.1 Results

With **RQ1**, we aimed to identify formalisms which are currently used in the modeling and specification of smart contracts, as well as approaches to their verification. Through our study, we made the following observations. (1) Among contract-level representations, the dominant combination is state-transition models with specifications in temporal logic, which are then verified by model checkers. As seen from Table 2, this combination is mostly used to formulate functional correctness

Table 2. A (partial) overview of the formalization and verification literature.

| Domains | Applications | Model Formalisms | | | | | Specification Formalisms | | | | Verification Techniques | | | | | |
|----------------|-------------------------|--------------------|------------------------|--|---------------------------|---|--|---|---------------------|--|-------------------------|------------------------|--|--|--|--------------------|
| | | Process Algebra | Transition System | Control-Flow Automata | Program Logic | Temporal Logics | Other Logics | Hoare Logic | Path-Level Patterns | Model Checking | Theorem Proving | Symbolic Execution | Program Verification | Runtime Verification | | |
| Finance | ICO / Token | [133] [⊖] | [147] [⊖] | [167] [⊖] | [165] [⊖] | [147, 167] [⊖] | [165] [⊖] [133] [⊖] | [130, 132] [⊖] | [60] [⊖] | [147] [⊖] | [132] [⊖] | [167] [⊖] | [133, 165] [⊖] | [60, 130] [⊖] | | |
| | Bank | | [74, 121] [⊖] | | [74] [⊖] | [74, 121] [⊖] | | [11] [⊖] | | [42] [⊖] | | | [11, 42] [⊖] | | | |
| | Wallet | | [158] [⊖] | | [40] [✧] | [158] [⊖] | | [40, 68] [✧] | | | | [158] [⊖] | [40] [✧] | [68] [✧] | | |
| | Escrow / Purchase | [171] [⊖] | [74] [⊖] | | [18, 30, 74] [⊖] | [74] [⊖] | | [18] [⊖] | | | [30, 171] [⊖] | [74, 171] [⊖] | [18] [⊖] | [30] [⊖] | | |
| Social Games | Auction | | [147] [⊖] | | | [147] [⊖] | | [231] [⊖] [38] [✧] | | | | [147] [⊖] | [136] [⊖] | [136, 231] [⊖] [38] [✧] | | |
| | Voting | | [74] [⊖] | | [74] [⊖] | [74] [⊖] | | [40] [✧] [22, 130] [⊖] | | | | [74] [⊖] | [40] [✧] | [136] [⊖] | [22, 136] [⊖] [130] [⊖] | |
| | Games / Gambling | | [197] [⊖] | | | [197] [⊖] | | [37] [✧] | | [76] [⊖] | | | [37] [✧] | [76] [⊖] | | |
| Asset Tracking | Supply Chain | | [16] [✧] | | | [16] [✧] | | [215] [⊖] | | [82] [✧] | | [16] [✧] | | [215] [⊖] [82] [✧] | | |
| | Marketplace | | [157] [⊖] | | | [157] [⊖] | | [156] [⊖] | | | | [157] [⊖] | | [156] [⊖] | | |
| | License Agreement | | [197] [⊖] | | [93] [✧] | [197] [⊖] | | [93] [✧] | | | | | | [93] [✧] | | |
| | Name Registration | | [10] [⊖] | | | [10] [⊖] | | [107] [⊖] | | | | [10] [⊖] | | [107] [⊖] | | |
| Protocols | Timed Commitment | [28] [♭] | [20] [♭] | | | [20, 28] [♭] | | | | | | [20, 28] [♭] | | | | |
| | Atomic Swap | | [207] [✧] | | [109] [✧] | [207] [✧] | | [109] [✧] | | | | [207] [✧] | | [109] [✧] | | |
| Security | Reentrancy | | [147] [⊖] | [141] [⊖] | | [147] [⊖] | [108] [⊖] [159] [✧] | | | [134, 141, 213] [⊖] | | [147] [⊖] | [108] [⊖] | [141] [⊖] | [213] [⊖] | [134] [⊖] |
| | Concurrency | | [171] [⊖] | | [120] [⊖] | | | | | [120, 171, 213] [⊖] | | [171] [⊖] | [120] [⊖] | [213] [⊖] | | |
| | Dependence Manipulation | | | [141] [⊖] [106] [†] | | | | | | [141, 142, 213] [⊖] [106] [†] | | | [141] [⊖] [106] [†] | [213] [⊖] | [142] [⊖] | |
| | Unchecked Call | | | [141] [⊖] | | | | | | [57, 141, 213] [⊖] | | | [141] [⊖] | [213] [⊖] | [57] [⊖] | |
| | Access Control | [133] [⊖] | | [167] [⊖] [106] [†] | | [167] [⊖] | [133, 196] [⊖] | [215] [⊖] | | [49, 57] [⊖] | | | [196] [⊖] | [167] [⊖] [106] [†] | [49, 133, 167, 215] [⊖] | [57] [⊖] |
| | Liquidity | [36] [♭] | [147] [⊖] | [160, 205] [⊖] | | [36] [♭] [147] [⊖] | [185] [✧] | | | [160, 205] [⊖] | | | [36] [♭] [147] [⊖] | [185] [✧] | [160] [⊖] | [205] [⊖] |
| | Resource Consumption | | | [58, 95] [⊖] | | | [89] [⊖] | [156] [⊖] | | [58, 95] [⊖] | | | [89] [⊖] | [58] [⊖] | [95, 156] [⊖] | |
| | Arithmetic | | | [81] [⊖] | | | | [196] [⊖] | [190] [⊖] | | [81, 142] [⊖] | | [196] [⊖] | [81] [⊖] | [190] [⊖] | [142] [⊖] |

⊖: Ethereum, ♭: Bitcoin, ✧: Hyperledger Fabric, ✧: Tezos, †: EOS, ✧: Other

properties, and some security properties involving the concept of progress, such as liquidity [36], are modeled as liveness. (2) Another popular approach is automated program-level verification against Hoare-style specifications. In terms of functional correctness, program logics defined at the bytecode-level allow the reasoning about smart contract behaviors using theorem provers, while source-level annotations are usually discharged by program verification techniques. Verification of security properties is dominated by existing well-established symbolic execution and program verification tool-chains, e.g., Boogie. With a few exceptions, these approaches identify security issues along program paths that match predefined vulnerable patterns. (3) Runtime verification checks for similar issues but on runtime execution traces. These techniques often require instrumentation of smart contract code with proper monitors, e.g., assertions. Runtime monitoring can also be implemented off-chain with tracking of events for compliance checking for escrow and supply chain applications [82, 151].

5.2 Open Challenges

With **RQ2** and **RQ3**, we aimed to identify current challenges and limitations in formal verification of smart contracts, respectively.

Pattern-Based Verification. As demonstrated in Table 2, smart contract security analysis is mostly performed at the program level. These techniques typically rely on *vulnerable patterns* that were defined by experts, which limits these tools to the identification of a known set of vulnerabilities [213]. The generalization to larger classes of vulnerabilities can be achieved by focusing on the root causes

of bugs, such as the influence of nondeterministic factors on smart contract execution [213]. To expand the list of detectable vulnerabilities, some works propose specification languages which allow users to define vulnerabilities [79], temporal properties [167], or annotate Solidity contracts with requirements, in forms of invariants [101, 211].

Limited Verification of Temporal Properties. Temporal logic is used as a universal formalism for high-level correctness specification for smart contracts. In nearly all cases, temporal properties are specified over a contract-level model of a smart contract, which typically abstract technical details of its operation. At the same time, program-level techniques that provide better precision in terms of the execution environment modeling rarely support specification and verification of temporal properties. Nevertheless, the wide use of temporal properties in smart contract specification suggests that they should be analyzed to the precision of program-level techniques.

Complications in Execution Environment Modeling. Arguably, the major difficulty of all techniques that perform static analyses lies in the modeling of intricate aspects of smart contract execution, which also differ across platforms. For example, a verifier for Ethereum smart contracts should precisely model the gas mechanism [167] and the memory model [22, 83, 102] with its hash-based object allocation [167]. Gas consumption in Ethereum is determined by the executed instructions at runtime, so gas-aware analyses have to be carried out at the bytecode-level, which makes specification of high-level properties less straightforward [22]. Precise reasoning about temporal properties concerned with physical time is also complicated by the absence of a global clock [54] and the delay in transaction processing [150]. Execution of Ethereum contracts also introduces many sources of nondeterminism, including transaction ordering and the fallback mechanism [213], only determined at runtime. Still, comprehensive verification of functional correctness and security of smart contracts should allow for inter-contractual reasoning [123], which accounts for a nondeterministic behavior of an external callee. In some platforms [19], nondeterminism can also arise from the development of smart contracts in general-purpose languages, which can demonstrate nondeterministic behavior. However, to ensure consensus checking between blockchain nodes, the usage of nondeterministic methods is strongly discouraged and can be detected statically [224] or in combination with dynamic analysis [191]. Other aspects of blockchain that require modeling include the possibility of forks [27, 109] and reverted transactions.

Limitations of Runtime Verification. Since smart contracts are immutable once deployed, the discussed techniques are mostly to be applied in design-time. Even in the domain of runtime verification, many techniques [29, 30, 46, 130, 193] rely on the instrumentation of contracts with protecting code, which has to be performed pre-deployment as well. Moreover, instrumentation introduces additional gas consumption [130] and increases the size of bytecode, which is limited [57]. Runtime verification of smart contracts is also performed through the application of instrumented Ethereum clients [57, 85, 98, 142]. They, however, can only operate locally—otherwise, a hard fork is required for its integration into a platform [130]. In case the information about the execution is retrieved from the blockchain, information about private contract variables may not be observable [208]. Furthermore, if transactions are analyzed offline, suspicious transactions cannot be reverted in runtime [57].

Fragmented Standards and Ecosystems. Despite the extensive literature, there is still not a clear path towards safe smart contract development [235]. This is particularly reflected by the lack of standard development tool-chains and the difficulty in choosing the right techniques. Arguably, the difference between the chosen formal specification and verification techniques as well as different threat models are the key reasons for a low agreement between smart contract verification tools demonstrated by recent studies [166, 213]. In addition, our observations coincide with the opinion

of He et al. [106] that smart contracts operating on different platforms require substantially different approaches to specification and vulnerability detection. Although Table 2 demonstrates that a vast majority of approaches concentrate on verification of Ethereum smart contracts, our study includes the work on Bitcoin [20, 28], Tezos [40, 41, 174], Hyperledger Fabric [16, 38, 82, 118, 143, 180, 224], EOS [106, 129, 172].

5.3 Future Directions

We envision a few promising research directions in the area of formal specification and verification of smart contracts.

Safe Languages. One of the emerging directions is to develop *safe smart contract languages* which eliminate many security risks by design or facilitate program verification. For example, Tezos supports three formally certified programming languages varying in platform abstraction level [40, 41, 174], shipped with associated verification frameworks. Besides formal semantics, other guarantees are derived from linear resource types in Move [45] and Flint [183], explicit state changes of Bamboo [2] and Obsidian [2, 65], functional programming principles and a restricted instruction set in Scilla [186], user-centered design as in Obsidian [65], etc. To allow automated formal verification, some languages (e.g., DAML [8], Pact [168], and TEAL [9]) are also intentionally designed as non-Turing complete. We also noticed an increasing number of *domain-specific languages*, for example, Marlowe [126, 127] and Findel [25, 44], which should simplify the development and verification of financial smart contracts.

Self-Healing Contracts. Another prominent direction is concerned with recovering from runtime failures/violations, a.k.a. *automated repair* of contracts. In 2017, Magazzeni et al. [144] suggested that AI planning techniques can be used to automatically patch the misaligned traces of smart contracts. One of the recent examples is an automated gas-aware repair technique for Ethereum smart contracts proposed by Yu et al. [228], which is based on a genetic mutation technique guided towards lower gas consumption. To ensure validity of a generated patch w.r.t. the legitimate behavior of a vulnerable smart contract, the authors automatically construct a test suite based on its previous blockchain transactions. The generated patches should, therefore, fix the identified vulnerability without breaking any of these previously passing tests. The gas constraints of the Ethereum blockchain introduce an additional criterion for the validity of a patch, which should minimize the gas consumption. The ability of a smart contract to recover from a violation could also be relevant in the context of smart legal contracts, e.g., to resolve disputes and conflicts between participants of a contract [66].

Hybrid Analyses. In the verification of smart contracts, better results can be achieved through a combination of online and offline techniques [57, 130]. A combination of contract-level and program-level approaches to smart contract modeling can additionally help mitigate the trade-off between the high-level behavioral specification and the precise modeling of execution details. Filling in the gap between high- and low-level analyses can enable accurate verification of interesting temporal properties. The study of smart contract fairness is complicated by the multi-user environment of its execution model and unknown intentions of contract designers. Still, to analyze such high-level properties, including fairness, researchers currently apply techniques from the domains of game theory and mechanism design [54, 128, 136].

Collaborative Development of Standards. Some work has been done to promote standards and best practices in the development of smart contracts, for example, collections of design and security [218, 219, 223] patterns. Two recent publications provided taxonomies of common smart

contract properties and invariants [39, 167], while the authors of [97] summarized their observations on the performance of static and dynamic analysis tools w.r.t. different security issues. The adoption of formally verified libraries is another step towards safer smart contract development. Two recent works [182, 231] used formal verification to assure the safety of the popular SafeMath library for Solidity. Other contributions include creation of a benchmark and curated datasets of vulnerable smart contracts [75] as well as evaluation frameworks for static analyses based on vulnerable code injection [12, 90]. The development of standard benchmark is facilitating advances in the ML/DL approaches to smart contract analysis—another trend in smart contract analysis that we observed in recent publications. We expect the work in this direction to continue.

6 CONCLUSION

In view of the increasing adoption of blockchain and smart contract technologies, formal specification and verification of smart contracts has been an active topic of research in recent years. Within this survey, we have presented a comprehensive analysis and classification of existing approaches to formal modeling, specification, and verification of smart contracts. Furthermore, we outlined common properties from different smart contract domains and correlated them with the capabilities of existing verification techniques. Our findings suggest that a combination of contract-level models and specifications with model checking is widely used to reason about the functional correctness of smart contracts from all the considered domains. On the other hand, program-level representations are often analyzed from the perspective of security properties, which are checked by means of symbolic execution, program verification toolchains, or theorem proving. Security issues can also be identified in execution traces of smart contracts by runtime verification techniques. A holistic perspective on smart contract analysis, which was adopted in this survey, allowed us to uncover the existing limitations to effective formal verification of smart contracts and suggest future research directions. We envision further advancement in the areas including safe smart contract language design, automated repair of smart contracts, hybrid contract- and program-level approaches to formal modeling and specification of smart contracts, and collaborative development of standards for safe smart contract implementation.

REFERENCES

- [1] 2016. King of the Ether Throne — Post-mortem investigation. Retrieved July 14, 2020 from <https://www.kingoftheether.com/postmortem.html>.
- [2] 2018. Bamboo: A morphing smart contract language. Retrieved July 14, 2020 from <https://github.com/cornellblockchain/bamboo>.
- [3] 2020. Common Patterns — Solidity 0.6.11 documentation. Retrieved July 14, 2020 from <https://solidity.readthedocs.io/en/v0.6.11/common-patterns.html>.
- [4] 2020. EOS.IO Technical white paper v2. Retrieved July 14, 2020 from <https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md>.
- [5] 2020. Ethereum Charts and Statistics | Etherscan. Retrieved February 13, 2020 from <https://etherscan.io/charts>.
- [6] 2020. Solidity — Solidity 0.6.11 documentation. Retrieved July 14, 2020 from <https://solidity.readthedocs.io/en/v0.6.11/>.
- [7] 2020. Solidity by Example — Solidity 0.6.11 documentation. Retrieved July 14, 2020 from <https://solidity.readthedocs.io/en/v0.6.11/solidity-by-example.html>.
- [8] 2021. Daml Programming Language. Retrieved April 1, 2021 from <https://daml.com>.
- [9] 2021. Transaction execution approval language (TEAL) specification. Retrieved April 1, 2021 from <https://developer.algorand.org/docs/reference/teal/specification>.
- [10] T. Abdellatif and K. L. Brousmiche. 2018. Formal verification of smart contracts based on users and blockchain behaviors models. In *Proc. of IFIP NTMS*, Vol. 2018-Janua. IEEE, 1–5.
- [11] W. Ahrendt, R. Bubel, J. Ellul, G. J. Pace, R. Pardo, V. Rebiscoul, and G. Schneider. 2019. Verification of smart contract business logic. In *Proc. of FSEN*. Springer International Publishing, Cham, 228–243.
- [12] S. Akca, A. Rajan, and C. Peng. 2019. SolAnalyser: A framework for analysing and testing smart contracts. In *Proc. of APSEC*. 482–489.

- [13] E. Albert, J. Correas, P. Gordillo, G. Román-Díez, and A. Rubio. 2019. GASOL: Gas analysis and optimization for Ethereum smart contracts. arXiv:1912.11929
- [14] E. Albert, J. Correas, P. Gordillo, G. Román-Díez, and A. Rubio. 2019. SAFEVM: A safety verifier for Ethereum smart contracts. abs/1906.04984 (2019). arXiv:1906.04984
- [15] E. Albert, J. Correas, P. Gordillo, G. Román-Díez, and A. Rubio. 2020. Analyzing smart contracts: From EVM to a sound control-flow graph. arXiv:2004.14437
- [16] S. Alqahtani, X. He, R. Gamble, and P. Mauricio. 2020. Formal verification of functional requirements for smart contract compositions in supply chain management systems. In *Proc. of HICSS*.
- [17] L. Alt and C. Reitwiessner. 2018. SMT-based verification of Solidity smart contracts. In *Proc. of ISoLA*, Vol. 11247 LNCS. Springer Verlag, 376–388.
- [18] S. Amani, M. Bégel, M. Bortin, and M. Staples. 2018. Towards verifying Ethereum smart contract bytecode in Isabelle/HOL. In *Proc. of ACM CPP*. ACM Press, 66–77.
- [19] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, N. Binh, M. Sethi, G. Singh, K. Smithm, A. Sorniotti, C. Stathakopoulou, M. Vukolic, S. W. Cocco, and J. Yellick. 2018. Hyperledger Fabric: A distributed operating system for permissioned blockchains. In *Proc. of EuroSys*. ACM, Article 30.
- [20] M. Andrychowicz, S. Dziembowski, D. Malinowski, and L. Mazurek. 2014. Modeling Bitcoin contracts by Timed Automata. In *Proc. of FORMATS*. Springer International Publishing, 7–22.
- [21] D. Annenkov, J. B. Nielsen, and B. Spitters. 2020. ConCert: A smart contract certification framework in Coq. In *Proc. of ACM CCP*. ACM, 215–228.
- [22] P. Antonino and A. W. Roscoe. 2020. Formalising and verifying smart contracts with Solidifier: a bounded model checker for Solidity. arXiv:2002.02710
- [23] A. W. Appel, R. Dockins, A. Hobor, L. Beringer, J. Dodds, G. Stewart, S. Blazy, and X. Leroy. 2014. *Program logics for certified compilers*. Cambridge University Press.
- [24] E. J. G. Arias. 2019. Towards principled compilation of Ethereum smart contracts (SoK). In *Proc. of IFIP NTMS*. 1–5.
- [25] A. Arusoai. 2020. Certifying Findel derivatives for blockchain. arXiv:2005.13602
- [26] N. Atzei, M. Bartoletti, and T. Cimoli. 2017. A survey of attacks on Ethereum smart contracts. In *Proc. of POST*. Springer, 164–186.
- [27] N. Atzei, M. Bartoletti, T. Cimoli, S. Lande, and R. Zunino. 2018. SoK: Unraveling Bitcoin smart contracts. In *Proc. of POST*. Springer International Publishing, 217–242.
- [28] N. Atzei, M. Bartoletti, S. Lande, N. Yoshida, and R. Zunino. 2019. Developing secure Bitcoin contracts with BitML. In *Proc. of ACM ESEC/FSE*. ACM, 1124–1128.
- [29] S. Azzopardi, J. Ellul, and G. J. Pace. 2019. Monitoring smart contracts: Contractlarva and open challenges beyond. In *Proc. of RV*, Vol. 11237. Springer Verlag, 113–137.
- [30] S. Azzopardi, G. J. Pace, and F. Schapachnik. 2018. On observing contracts: Deontic contracts meet smart contracts. In *Proc. of JURIX*, Vol. 313. IOS Press, 21–30.
- [31] J.C.M. Baeten. 2005. A brief history of Process algebra. *Theor. Comput. Sci.* 335, 2 (2005), 131–146. Process Algebra.
- [32] X. Bai, Z. Cheng, Z. Duan, and K. Hu. 2018. Formal modeling and verification of smart contracts. In *Proc. of ACM ICSCA*. ACM Press, 322–326.
- [33] R. Banach. 2020. Verification-led smart contracts. In *Proc. of FC*. Springer International Publishing, 106–121.
- [34] M. Bartoletti and L. Pompianu. 2017. An empirical analysis of smart contracts: Platforms, applications, and design patterns. In *Proc. of FC*, Vol. 10323 LNCS. Springer Verlag, 494–509. arXiv:1703.06322
- [35] M. Bartoletti and R. Zunino. 2019. Formal models of Bitcoin contracts: A survey. *Frontiers in Blockchain* 2 (2019), 8.
- [36] M. Bartoletti and R. Zunino. 2019. Verifying liquidity of Bitcoin contracts. In *Proc. of POST*. Springer International Publishing, 222–247.
- [37] B. Beckert, M. Herda, M. Kirsten, and J. Schiffl. 2018. Formal specification and verification of Hyperledger Fabric chaincode. In *Proc. of SDLT*. Institute for Integrated and Intelligent Systems, Griffith University, Australia, 44–48.
- [38] B. Beckert, J. Schiffl, and M. Ulbrich. 2019. Smart contracts: Application scenarios for program verification. Retrieved July 14, 2020 from <https://www.key-project.org/wp-content/uploads/2019/11/sc-verification.pdf>.
- [39] T. Bernardi, N. Dor, A. Fedotov, S. Grossman, N. Immerman, D. Jackson, A. Nutz, L. Oppenheim, O. Pistiner, N. Rinetzky, M. Sagiv, M. Taube, J. A. Toman, and J. R. Wilcox. 2020. WIP: Finding bugs automatically in smart contracts with parameterized invariants. Retrieved July 14, 2020 from <https://www.certora.com/pubs/sbc2020.pdf>. (2020).
- [40] B. Bernardo, R. Cauderlier, Z. Hu, B. Pesin, and J. Tesson. 2019. Mi-Cho-Coq. A framework for certifying Tezos smart contracts. (sep 2019). arXiv:1909.08671
- [41] B. Bernardo, R. Cauderlier, B. Pesin, and J. Tesson. 2020. Albert, an intermediate smart-contract language for the Tezos blockchain. arXiv:2001.02630

- [42] K. Bhargavan, N. Swamy, S. Zanella-Béguelin, A. Delignat-Lavaud, C. Fournet, A. Gollamudi, G. Gonthier, N. Kobeissi, N. Kulatova, A. Rastogi, and T. Sibut-Pinote. 2016. Formal verification of smart contracts. In *Proc. of ACM PLAS*. ACM Press, 91–96.
- [43] G. Bigi, A. Bracciali, G. Meacci, and E. Tuosto. 2015. *Validation of decentralised smart contracts through game theory and formal methods*. Springer International Publishing, 142–161.
- [44] A. Biryukov, D. Khovratovich, and S. Tikhomirov. 2017. Findel: Secure derivative contracts for Ethereum. In *Proc. of FC*, Vol. 10323 LNCS. Springer Verlag, 453–467.
- [45] S. Blackshear, D. L. Dill, S. Qadeer, C. W. Barrett, J. C. Mitchell, O. Padon, and Y. Zohar. 2020. Resources: A safe language abstraction for money. arXiv:2004.05106
- [46] Vera Bogdanich Espina. 2019. *Ethereum smart contracts verification: a survey and a prototype tool*. Master’s thesis.
- [47] Kees Boogaard. 2018. *A model-driven approach to smart contract development*. Master’s thesis.
- [48] S. Bragagnolo, H. Rocha, M. Denker, and S. Ducasse. 2018. SmartInspect: Solidity smart contract inspector. In *Proc. of IEEE IWBOSE*. 9–18.
- [49] L. Brent, N. Grech, S. Lagouvardos, B. Scholz, and Y. Smaragdakis. 2020. Ethainter: A smart contract security analyzer for composite vulnerabilities. In *Proc. of ACM PLDI*. ACM, 454–469.
- [50] L. Brent, A. Jurisevic, M. Kong, E. Liu, F. Gauthier, V. Gramoli, R. Holz, and B. Scholz. 2018. Vandal: A scalable security analysis framework for smart contracts. arXiv:1809.03981
- [51] E. Cecchetti, S. Yao, H. Ni, and A. C. Myers. 2020. Securing smart contracts with information flow. Retrieved July 14, 2020 from <http://www.cs.cornell.edu/andru/papers/fab20/fab20.pdf>.
- [52] J. Chang, B. Gao, H. Xiao, J. Sun, Y. Cai, and Z. Yang. 2019. sCompile: Critical path identification and analysis for smart contracts. In *Proc. of ICFEM*. Springer International Publishing, 286–304.
- [53] J. Chapman, R. Kireev, C. Nester, and P. Wadler. 2019. System F in Agda, for fun and profit. In *Proc. of MPC*. Springer International Publishing, 255–297.
- [54] K. Chatterjee, A. K. Goharshady, and Y. Velner. 2018. Quantitative analysis of smart contracts. In *Proc. of ESOP*, Vol. 10801 LNCS. Springer Verlag, 739–767. arXiv:1801.03367
- [55] H. Chen, M. Pendleton, L. Njilla, and S. Xu. 2020. A survey on Ethereum systems security: Vulnerabilities, attacks, and defenses. *ACM Comput. Surv.* 53, 3, Article 67 (June 2020).
- [56] J. Chen, X. Xia, D. Lo, and J. Grundy. 2020. Why do smart contracts self-destruct? Investigating the selfdestruct function on Ethereum. arXiv:2005.07908
- [57] T. Chen, R. Cao, T. Li, X. Luo, G. Gu, Y. Zhang, Z. Liao, H. Zhu, G. Chen, Z. He, Y. Tang, X. Lin, and X. Zhang. 2020. SODA: A generic online detection framework for smart contracts. In *Proc. of NDSS*.
- [58] T. Chen, X. Li, X. Luo, and X. Zhang. 2017. Under-optimized smart contracts devour your money. In *Proc. of IEEE SANER*. 442–446.
- [59] T. Chen, Z. Li, Y. Zhang, X. Luo, T. Wang, T. Hu, X. Xiao, D. Wang, J. Huang, and X. Zhang. 2019. A large-scale empirical study on control flow Identification of smart contracts. In *Proc. of ACM/IEEE ESEM*. 1–11.
- [60] T. Chen, Y. Zhang, Z. Li, X. Luo, T. Wang, R. Cao, X. Xiao, and X. Zhang. 2019. TokenScope: Automatically detecting inconsistent behaviors of cryptocurrency tokens in Ethereum. In *Proc. of ACM CCS*. ACM, 1503–1520.
- [61] W. Chen, T. Zhang, Z. Chen, Z. Zheng, and Y. Lu. 2020. Traveling the token world: A graph analysis of Ethereum ERC20 token ecosystem. In *Proc. of WWW*. ACM, 1411–1421.
- [62] W. Chen, Z. Zheng, J. Cui, E. Ngai, P. Zheng, and Y. Zhou. 2018. Detecting Ponzi schemes on Ethereum: Towards healthier blockchain technology. In *Proc. of WWW*. 1409–1418.
- [63] R. Cheng, F. Zhang, J. Kos, W. He, N. Hynes, N. Johnson, A. Juels, A. Miller, and D. Song. 2019. Ekiden: A platform for confidentiality-preserving, trustworthy, and performant smart contracts. In *Proc. of IEEE EuroS&P*. 185–200.
- [64] C. D. Clack and G. Vanca. 2018. Temporal aspects of smart contracts for financial derivatives. In *Proc. of IsoLA*, Vol. 11247 LNCS. Springer Verlag, 339–355. arXiv:1805.11677
- [65] M. Coblenz, R. Oei, T. Etzel, P. Koronkevich, M. Baker, Y. Bloem, B. A. Myers, J. Sunshine, and J. Aldrich. 2019. Obsidian: Typestate and assets for safer blockchain programming. arXiv:1909.03523
- [66] C. Colombo, J. Ellul, and G. J. Pace. 2018. Contracts over smart contracts: Recovering from violations dynamically. In *Proc. of IsoLA*, Vol. 11247 LNCS. Springer Verlag, 300–315.
- [67] S. E. S. Crawford and E. Ostrom. 1995. A grammar of institutions. *Am. Political Sci. Rev.* 89, 3 (1995), 582–600.
- [68] L. P. A. da Horta, J. S. Reis, M. Pereira, and S. M. de Sousa. 2020. WhySon: Proving your Michelson smart contracts in Why3. arXiv:2005.14650
- [69] Z. Dargaye, A. Del Pozzo, and S. Tucci Piergiovanni. 2018. Pluralize: a trustworthy framework for high-level smart contract-draft. abs/1812.05444 (2018). arXiv:1812.05444
- [70] K. Delmolino, M. Arnett, A. Kosba, A. Miller, and E. Shi. 2016. Step by step towards creating a safe smart contract: Lessons and insights from a cryptocurrency lab. In *Proc. of FC*. 79–94.

- [71] M. di Angelo and G. Salzer. 2019. A survey of tools for analyzing Ethereum smart contracts. In *Proc. of IEEE DAPPCON*. IEEE, 69–78.
- [72] C. Di Ciccio, A. Cecconi, M. Dumas, L. García-Bañuelos, O. López-Pintado, Q. Lu, J. Mendling, A. Ponomarev, A. Binh Tran, and I. Weber. 2019. Blockchain support for collaborative business processes. *Informatik Spektrum* 42, 3 (01 Jun 2019), 182–190.
- [73] T. Dickerson, P. Gazzillo, M. Herlihy, V. Saraph, and E. Koskinen. 2019. Proof-carrying smart contracts. In *Proc. of FC*. Springer Berlin Heidelberg, 325–338.
- [74] W. Duo, H. Xin, and M. Xiaofeng. 2020. Formal analysis of smart contract Based on Colored Petri Nets. *IEEE Intelligent Systems* 35, 3 (2020), 19–30.
- [75] T. Durieux, J. F. Ferreira, R. Abreu, and P. Cruz. 2020. Empirical review of automated analysis tools on 47,587 Ethereum smart contracts. *Proc. of IEEE/ACM ICSE* (Jun 2020). arXiv:1910.10601
- [76] J. Ellul and G. J. Pace. 2018. Runtime verification of Ethereum smart contracts. In *Proc. of EDCC*. IEEE, 158–163.
- [77] J. Feist, G. Greico, and A. Groce. 2019. Slither: A static analysis framework for smart contracts. In *Proc. of IEEE/ACM WETSEB*. IEEE Press, 8–15.
- [78] X. Feng, Q. Wang, X. Zhu, and S. Wen. 2019. Bug searching in smart contract. (may 2019). arXiv:1905.00799
- [79] Y. Feng, E. Torlak, and R. Bodik. 2019. Precise attack synthesis for smart contracts. (2019). arXiv:1902.06067
- [80] C. Ferreira Torres, M. Baden, R. Norvill, and H. Jonker. 2019. ÆGIS: Smart shielding of smart contracts. In *Proc. of ACM CCS*. ACM, 2589–2591.
- [81] C. Ferreira Torres, J. Schütte, and R. State. 2018. Osiris: Hunting for integer bugs in Ethereum smart contracts. In *Proc. of ACSAC*. ACM, 664–676.
- [82] F. Fournier and I. Skarbovsky. 2019. Enriching smart contracts with temporal aspects. In *Proc. of ICBC*, Vol. 11521 LNCS. Springer Verlag, 126–141.
- [83] J. Frank, C. Aschermann, and T. Holz. 2020. ETHBMC: A bounded model checker for smart contracts. In *Proc. of USENIX Security*. USENIX Association.
- [84] C. K. Frantz and M. Nowostawski. 2016. From institutions to code: Towards automated generation of smart contracts. In *Proc. of IEEE FAS-W*. IEEE, 210–215.
- [85] J. Gao, H. Liu, C. Liu, Q. Li, Z. Guan, and Z. Chen. 2019. EASYFLOW: Keep Ethereum away from overflow. In *Proc. of IEEE/ACM ICSE-Companion*. 23–26.
- [86] Z. Gao, V. Jayasundara, L. Jiang, X. Xia, D. Lo, and J. Grundy. 2019. SmartEmbed: A tool for clone and bug detection in smart contracts through structural code embedding. *Proc. of IEEE ICSME* (2019).
- [87] P. Garamvölgyi, I. Kocsis, B. Gehl, and A. Klenik. 2018. Towards model-driven engineering of smart contracts for cyber-physical systems. In *Proc. of IEEE/IFIP DSN-W*. 134–139.
- [88] L. García-Bañuelos, A. Ponomarev, M. Dumas, and I. Weber. 2017. Optimized execution of business processes on blockchain. In *BPM*. Springer International Publishing, 130–146.
- [89] T. Genet, T. Jensen, and J. Sauvage. 2020. *Termination of Ethereum’s smart contracts*. Research Report. Univ Rennes, Inria, CNRS, IRISA. <https://hal.inria.fr/hal-02555738>
- [90] A. Ghaleb and K. Pattabiraman. 2020. How effective are smart contract analysis tools? Evaluating smart contract static analysis tools using bug injection. arXiv:2005.11613
- [91] L. M. Goodman. 2014. Tezos — A self-amending crypto-ledger. White paper. Retrieved July 14, 2020 from https://tezos.com/static/white_paper-2dc8c02267a8fb86bd67a108199441bf.pdf.
- [92] G. Governatori. 2014. Thou Shalt is not You Will. abs/1404.1685 (2014). arXiv:1404.1685
- [93] G. Governatori, F. Idelberger, Z. Milosevic, R. Riveret, G. Sartor, and X. Xu. 2018. On legal contracts, imperative and declarative smart contracts, and blockchain systems. *Artificial Intelligence and Law* 26, 4 (01 Dec 2018), 377–409.
- [94] N. Grech, L. Brent, B. Scholz, and Y. Smaragdakis. 2019. Gigahorse: Thorough, declarative decompilation of smart contracts. In *Proc. of IEEE/ACM ICSE*. 1176–1186.
- [95] N. Grech, M. Kong, A. Jurisevic, L. Brent, B. Scholz, and Y. Smaragdakis. 2018. MadMax: Surviving out-of-gas conditions in Ethereum smart contracts. *Proc. of ACM OOPSLA*, Article 116 (Oct. 2018), 116:1–116:27 pages.
- [96] I. Grishchenko, M. Maffei, and C. Schneidewind. 2018. A semantic framework for the security analysis of Ethereum smart contracts. In *Proc. of POST*. Springer, Cham, 243–269.
- [97] A. Groce, J. Feist, G. Grieco, and M. Colburn. 2019. What are the actual flaws in important smart contracts (and how can we find them)? arXiv:1911.07567
- [98] S. Grossman, I. Abraham, G. Golan-Gueta, Y. Michalevsky, N. Rinetzky, M. Sagiv, and Y. Zohar. 2017. Online detection of effectively callback free objects with applications to smart contracts. *Proc. of ACM POPL* (dec 2017), 1–28.
- [99] NCC Group. 2018. TOP 10 — Arithmetic Issues. Retrieved April 05, 2021 from <https://www.dasp.co/>.
- [100] S. Haber and W. S. Stornetta. 1991. How to Time-Stamp a Digital Document. *J. Cryptol.* 3, 2 (Jan. 1991), 99–111.
- [101] Á. Hajdu and D. Jovanović. 2019. solc-verify: A modular verifier for Solidity smart contracts. (2019). arXiv:1907.04262

- [102] Á. Hajdu and D. Jovanović. 2020. SMT-Friendly formalization of the Solidity memory model. *Proc. of ESOP (2020)*, 224–250.
- [103] Á. Hajdu, D. Jovanović, and G. Ciocarlie. 2020. Formal specification and verification of Solidity contracts with events. arXiv:2005.10382
- [104] D. Harz and W. Knottenbelt. 2018. Towards safer smart contracts: A survey of languages and verification methods. (sep 2018). arXiv:1809.09805
- [105] J. He, M. Balunović, N. Ambroladze, P. Tsankov, and M. Vechev. 2019. Learning to fuzz from symbolic execution with application to smart contracts. In *Proc. of ACM CCS*. ACM, 531–548.
- [106] N. He, R. Zhang, L. Wu, H. Wang, X. Luo, Y. Guo, T. Yu, and X. Jiang. 2020. Security analysis of EOSIO smart contracts. arXiv:2003.06568
- [107] Y. Hirai. 2016. Formal verification of Deed contract in Ethereum Name Service. Retrieved July 14, 2020 from <https://yoichihirai.com/deed.pdf>.
- [108] Y. Hirai. 2017. Defining the Ethereum Virtual Machine for interactive theorem provers. In *Proc. of FC*. Springer, Cham, 520–535.
- [109] Y. Hirai. 2018. Blockchains as Kripke models: An analysis of atomic cross-chain swap. In *Proc. of ISoLA*. Springer International Publishing, 389–404.
- [110] G. J. Holzmann. 1997. The model checker SPIN. *IEEE Trans. Softw. Eng.* 23, 5 (1997), 279–295.
- [111] J. Hu and Y. Zhong. 2018. A method of logic-based smart contracts for blockchain system. In *Proc. of ICDPA*. ACM, 58–61.
- [112] Y. Huang, Y. Bian, R. Li, J. L. Zhao, and P. Shi. 2019. Smart contract security: A software lifecycle perspective. *IEEE Access* 7 (2019), 150184–150202.
- [113] Y. Huang, Q. Kong, N. Jia, X. Chen, and Z. Zheng. 2019. Recommending differentiated code to support smart contract update. In *Proc. of IEEE/ACM ICPC*. 260–270.
- [114] A. Imeri, N. Agoulmine, and D. Khadraoui. 2020. Smart contract modeling and verification techniques: A survey. In *Proc. of ADVANCE*. 1–8.
- [115] Runtime Verification Inc. 2018. ERC20-K: Formal executable specification of ERC20. Retrieved July 14, 2020 from <https://github.com/runtimeverification/erc20-semantics>.
- [116] B. Jiang, Y. Liu, and W. K. Chan. 2018. ContractFuzzer: Fuzzing smart contracts for vulnerability detection. In *Proc. of ACM/IEEE ASE*. ACM, 259–269.
- [117] J. Jiao, S. Kan, S. Lin, D. Sanán, Y. Liu, and J. Sun. 2020. Semantic understanding of smart Contracts: executable operational semantics of Solidity. In *Proc. of IEEE S&P*. IEEE Computer Society, 1265–1282.
- [118] S. Kalra, M. Goel, S. and Dhawan, and S. Sharma. 2018. ZEUS: Analyzing safety of smart contracts. In *Proc. of NDSS*.
- [119] T. Kasampalis, D. Guth, B. Moore, T. F. Serbanuta, Y. Zhang, D. Filaretti, V. Serbanuta, R. Johnson, and G. Roşu. 2019. IELE: A rigorously designed language and tool ecosystem for the blockchain. In *Proc. of FM*.
- [120] A. Kolluri, I. Nikolic, I. Sergey, A. Hobor, and P. Saxena. 2019. Exploiting the laws of order in smart contracts. In *Proc. of ACM ISSTA*. ACM, 363–373.
- [121] J. Kongmanee, P. Kijsanayothin, and R. Hewett. 2019. Securing smart contracts in blockchain. In *Proc. of ACM/IEEE ASEW*. 69–76.
- [122] A. Kosba, A. Miller, E. Shi, Z. Wen, and C. Papamanthou. 2016. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *Proc. of IEEE S&P*. 839–858.
- [123] J. Krupp and C. Rossow. 2018. teEther: Gnawing at Ethereum to automatically exploit smart contracts. In *Proc. of USENIX Security*. ACM, 1317–1333.
- [124] A. Juels L. Breidenbach, P. Daian and E. G. Sire. 2017. An in-depth look at the Parity Multisig bug. Retrieved February 13, 2020 from <https://hackingdistributed.com/2017/07/22/deep-dive-parity-bug/>.
- [125] J. Ladleif and M. Weske. 2019. A unifying model of legal smart contracts. In *Proc. of ER*. 323–337.
- [126] P. Lamela Seijas, A. Nemish, D. Smith, and S. Thompson. 2020. Marlowe: Implementing and analysing financial contracts on blockchain. Retrieved July 14, 2020 from <https://iohk.io/en/research/library/papers/marloweimplementing-and-analysing-financial-contracts-on-blockchain/>.
- [127] P. Lamela Seijas and S. Thompson. 2018. Marlowe: Financial contracts on blockchain. In *Proc. of ISoLA*, Vol. 11247 LNCS. Springer Verlag, 356–375.
- [128] C. Laneve, C. S. Coen, and A. Veschetti. 2019. *On the prediction of smart contracts' behaviours*. Springer International Publishing, 397–415.
- [129] S. Lee, D. Kim, D. Kim, S. Son, and Y. Kim. 2019. Who spent my EOS? On the (in)security of resource management of EOS.IO. In *Proc. of USENIX WOOT*. USENIX Association.
- [130] A. Li, J. A. Choi, and F. Long. 2020. Securing smart contract with runtime validation. In *Proc. of ACM PLDI*. ACM, 438–453.
- [131] A. Li and F. Long. 2018. Detecting standard violation errors in smart contracts. (dec 2018). arXiv:1812.07702

- [132] X. Li, Z. Shi, Q. Zhang, G. Wang, Y. Guan, and N. Han. 2019. Towards verifying Ethereum smart contracts at intermediate language level. In *Proc. of ICFEM*. Springer International Publishing, 121–137.
- [133] X. Li, C. Su, Y. Xiong, W. Huang, and W. Wang. 2019. Formal verification of BNB smart contract. In *Proc. of BIGCOM*. 74–78.
- [134] C. Liu, H. Liu, Z. Cao, Z. Chen, B. Chen, and B. Roscoe. 2018. ReGuard: Finding reentrancy bugs in smart contracts. In *Proc. of IEEE/ACM ICSE*. ACM, 65–68.
- [135] H. Liu, C. Liu, W. Zhao, Y. Jiang, and J. Sun. 2018. S-gram: Towards semantic-aware security auditing for Ethereum smart contracts. In *Proc. of ACM/IEEE ASE*. ACM, 814–819.
- [136] Y. Liu, Y. Li, S.-W. Lin, and R. Zhao. 2020. Towards automated verification of smart contract fairness. In *Proc. of ACM ESEC/FSE*.
- [137] Y. Liu, J. Sun, and J. S. Dong. 2011. PAT 3: An Extensible Architecture for Building Multi-domain Model Checkers. In *Proc. of IEEE ISSRE*.
- [138] Z. Liu and J. Liu. 2019. Formal verification of blockchain smart contract based on Colored Petri Net models. In *Proc. of IEEE COMPSAC*. IEEE, 555–560.
- [139] O. López-Pintado, L. García-Bañuelos, M. Dumas, I. Weber, and A. Ponomarev. 2018. CATERPILLAR: A business process execution engine on the Ethereum blockchain. *Softw., Pract. Exper.* 49 (2018), 1162–1193.
- [140] N. Lu, B. Wang, Y. Zhang, W. Shi, and C. Esposito. 2019. NeuCheck: A more practical Ethereum smart contract security analysis tool. *Softw., Pract. Exper.* (2019).
- [141] L. Luu, D.-H. Chu, H. Olickel, P. Saxena, and A. Hobor. 2016. Making smart contracts smarter. In *Proc. of ACM CCS*. ACM, 254–269.
- [142] F. Ma, Y. Fu, M. Ren, M. Wang, Y. Jiang, K. Zhang, H. Li, and X. Shi. 2019. EVM*: From offline detection to online reinforcement for Ethereum Virtual Machine. In *Proc. of IEEE SANER*. 554–558.
- [143] G. Madl, L. Bathen, G. Flores, and D. Jadav. 2019. Formal verification of smart contracts using interface automata. In *Proc. of IEEE Blockchain*. 556–563.
- [144] D. Magazzeni, P. Mcburney, and W. Nash. 2017. Validation and verification of smart contracts: A research agenda. *Computer* 50, 9 (2017), 50–57.
- [145] D. B. Maksimov, I. A. Yakimov, and A. S. Kuznetsov. 2020. Statistical model checking for blockchain-based applications. *IOP Conf. Ser. Mater. Sci. Eng.* 734 (jan 2020), 012152.
- [146] M. Marescotti, M. Blicha, A. E. J. Hyvärinen, S. Asadi, and N. Sharygina. 2018. Computing exact worst-case gas consumption for smart contracts. In *Proc. of ISO/LA*. Springer International Publishing, 450–465.
- [147] A. Mavridou, A. Laszka, E. Stachtari, and A. Dubey. 2019. VeriSolid: Correct-by-design smart contracts for Ethereum. (jan 2019). arXiv:1901.01292
- [148] A. Mavridou, A. and Laszka. 2018. Tool demonstration: FSolidM for designing secure Ethereum smart contracts. In *Proc. of POST*. 270–277.
- [149] A. Mense and M. Flatscher. 2018. Security vulnerabilities in Ethereum smart contracts. In *Proc. of iiWAS2018 (iiWAS2018)*. ACM, 375–380.
- [150] A. Miller, Z. Cai, and S. Jha. 2018. Smart contracts and opportunities for formal methods. In *Proc. of ISO/LA*, Vol. 11247 LNCS. Springer Verlag, 280–299.
- [151] C. Molina-Jimenez, I. Sfyarakis, E. Solaiman, I. Ng, M. Weng Wong, A. Chun, and J. Crowcroft. 2018. Implementation of smart contracts using hybrid architectures with on and off-blockchain components. In *Proc. of IEEE SC2*. 83–90.
- [152] M. Mossberg, F. Manzano, E. Hennenfent, A. Groce, G. Grieco, J. Feist, T. Brunson, and A. Dinaburg. 2019. Manticore: A user-friendly symbolic execution framework for binaries and smart contracts. In *Proc. of ACM/IEEE ASE*. 1186–1189.
- [153] B. Mueller. 2018. Smashing Ethereum smart contracts for fun and real profit. Retrieved July 14, 2020 from <https://github.com/b-mueller/smashing-smart-contracts/blob/master/smashing-smart-contracts-1of1.pdf>.
- [154] Y. Murray and D. A. Anisi. 2019. Survey of formal verification methods for smart contracts on blockchain. In *Proc. of IFIP NTMS*. IEEE, 1–6.
- [155] S. Nakamoto et al. 2008. Bitcoin: A peer-to-peer electronic cash system. Retrieved July 14, 2020 from <https://bitcoin.org/bitcoin.pdf>. (2008).
- [156] Z. Nehai and F. Bobot. 2019. Deductive proof of Ethereum smart contracts using Why3. (2019). arXiv:1904.11281
- [157] Z. Nehai, P. Piriou, and F. Daumas. 2018. Model-checking of smart contracts. In *Proc. of IEEE iThings/GreenCom/CP-SCoM/SmartData*. IEEE, 980–987.
- [158] K. Nelaturu, A. Mavridou, A. Veneris, and A. Laszka. 2020. Verified development and deployment of multiple interacting smart contracts with VeriSolid. Retrieved July 14, 2020 from <http://www.eecg.utoronto.ca/~veneris/20ICBC.pdf>.
- [159] J. B. Nielsen and B. Spitters. 2019. Smart Contract Interactions in Coq. arXiv:1911.04732
- [160] I. Nikolić, A. Kolluri, I. Sergey, P. Saxena, and A. Hobor. 2018. Finding the greedy, prodigal, and suicidal contracts at scale. *Proc. of ACSAC* (2018), 653–663.
- [161] R. O’Connor. 2017. Simplicity: A new language for blockchains. In *Proc. of ACM PLAS*. ACM, 107–120.

- [162] OpenZeppelin. 2020. SafeMath library. Retrieved July 14, 2020 from <https://github.com/OpenZeppelin/openzeppelin-contracts/blob/master/contracts/math/SafeMath.sol>.
- [163] T. Osterland and T. Rose. 2020. Model checking smart contracts for Ethereum. *Pervasive and Mobile Computing* (2020), 101–129.
- [164] R. M. Parizi, A. Singh, and A. Dehghantanha. 2018. Smart contract programming languages on blockchains: An empirical evaluation of usability and security. In *Proc. of ICBC*, Vol. 10974 LNCS. Springer Verlag, 75–91.
- [165] D. Park, Y. Zhang, M. Saxena, P. Daian, and G. Roşu. 2018. A formal verification tool for Ethereum VM bytecode. In *Proc. of ACM ESEC/FSE*. ACM, 912–915.
- [166] D. Perez and B. Livshits. 2019. Smart contract vulnerabilities: Does anyone care? (feb 2019). arXiv:1902.06710
- [167] A. Permenev, D. Dimitrov, P. Tsankov, D. Drachler-Cohen, and M. Vechev. 2020. VerX: Safety verification of smart contracts. In *Proc. of IEEE S&P*. IEEE Computer Society, 414–430.
- [168] S. Popejoy. 2017. The Pact Smart-Contract Language. Retrieved April 1, 2021 from <https://www.kadena.io/whitepapers>.
- [169] P. Pratheeshan, L. Pan, J. Yu, J. Liu, and R. Doss. 2019. Security analysis methods on Ethereum smart contract vulnerabilities: A survey. (aug 2019). arXiv:1908.08605
- [170] D. Prechtel, T. Groß, and T. Müller. 2019. Evaluating spread of ‘Gasless Send’ in Ethereum smart contracts. In *Proc. of IFIP NTMS*. 1–6.
- [171] M. Qu, X. Huang, X. Chen, Y. Wang, X. Ma, and D. Liu. 2018. Formal verification of smart contracts from the perspective of concurrency. In *Proc. of SmartBlock*, Vol. 11373 LNCS. Springer Verlag, 32–43.
- [172] L. Quan, L. Wu, and H. Wang. 2019. EVulHunter: Detecting fake transfer vulnerabilities for EOSIO’s smart contracts at Webassembly-level. (2019). arXiv:1906.10362
- [173] RChain. 2018. Contract Design — RChain Architecture 0.9.0 documentation. Retrieved July 14, 2020 from <https://architecture-docs.readthedocs.io/contracts/contract-design.html>.
- [174] J. S. Reis, P. Crocker, and S. M. de Sousa. 2020. Tezla, an intermediate representation for static analysis of Michelson smart contracts. arXiv:2005.11839
- [175] M. Rodler, W. Li, G. O. Karame, and L. Davi. 2018. Sereum: Protecting existing smart contracts against re-entrancy attacks. abs/1812.05934 (2018). arXiv:1812.05934
- [176] S. Rouhani and R. Deters. 2019. Security, performance, and applications of smart contracts: A systematic survey. *IEEE Access* 7 (2019), 50759–50779.
- [177] J. Rushby. 2001. *Theorem proving for verification*. Springer Berlin Heidelberg, 39–57.
- [178] N. F. Samreen and M. H. Alalfi. 2020. Reentrancy vulnerability identification in Ethereum smart contracts. In *Proc. of IEEE IWBOSE*. 22–29.
- [179] D. C. Sánchez. 2018. Raziell: Private and verifiable smart contracts on blockchains. (2018). arXiv:1807.09484
- [180] N. Sato, T. Tateishi, and S. Amano. 2018. Formal requirement enforcement on smart contracts based on Linear Dynamic Logic. In *Proc. of IEEE iThings/GreenCom/CPSCoM/SmartData*. 945–954.
- [181] S. Sayeed, H. Marco-Gisbert, and T. Caira. 2020. Smart contract: Attacks and protections. *IEEE Access* 8 (2020), 24416–24427.
- [182] C. Schneidewind, I. Grishchenko, M. Scherer, and M. Maffei. 2020. eThor: Practical and provably sound static analysis of Ethereum smart contracts. arXiv:2005.06227
- [183] F. Schrans, S. Eisenbach, and S. Drossopoulou. 2018. Writing safe smart contracts in Flint. In *Proc. of ACM Programming Companion*. ACM, 218–219.
- [184] I. Sergey and A. Hobor. 2017. A concurrent perspective on smart contracts. In *Proc. of FC*, Vol. 10323 LNCS. Springer Verlag, 478–493.
- [185] I. Sergey, A. Kumar, and A. Hobor. 2018. Temporal properties of smart contracts. In *Proc. of ISoLA*. Springer, Cham, 323–338.
- [186] I. Sergey, V. Nagaraj, J. Johannsen, A. Kumar, A. Trunov, and K. C. G. Hao. 2019. Safer smart contract programming with Scilla. *Proc. of ACM OOPSLA*, Article 185 (Oct. 2019), 185:1–185:30 pages.
- [187] E. Shishkin. 2018. Debugging smart contract’s business logic using symbolic model-checking. (dec 2018). arXiv:1812.00619
- [188] D. Siegel. 2016. Understanding the DAO attack — CoinDesk. Retrieved July 14, 2020 from <https://www.coindesk.com/understanding-dao-hack-journalists>.
- [189] A. Singh, R. Parizi, Q. Zhang, K.-K. R. Choo, and A. Dehghantanha. 2019. Blockchain smart contracts formalization: Approaches and challenges to address vulnerabilities. *Comput. Secur.* 88 (10 2019), 101654.
- [190] S. So, M. Lee, J. Park, H. Lee, and H. Oh. 2020. VeriSmart: A highly precise safety verifier for Ethereum smart contracts. In *Proc. of IEEE S&P*. IEEE Computer Society, 825–841.
- [191] F. Spoto. 2020. Enforcing Determinism of Java Smart Contracts. In *Proc. of FC*, Matthew Bernhard, Andrea Bracciali, L. Jean Camp, Shin’ichiro Matsuo, Alana Maurushat, Peter B. Rønne, and Massimiliano Sala (Eds.). Springer International Publishing, Cham, 568–583.

- [192] S. Steffen, B. Bichsel, M. Gersbach, N. Melchior, P. Tsankov, and M. Vechev. 2019. Zkay: Specifying and enforcing data privacy in smart contracts. In *Proc. of ACM CCS*. ACM, 1759–1776.
- [193] L. Stegeman. 2018. *Solitor: Runtime verification of smart contracts on the Ethereum network*. Master’s thesis. <http://essay.utwente.nl/76902/>
- [194] J. Sun, Y. Liu, and J. S. Dong. 2008. Model Checking CSP Revisited: Introducing a Process Analysis Toolkit. In *Proc. of ISoLA*. Springer Berlin Heidelberg.
- [195] J. Sun, Y. Liu, J. S. Dong, and J. Pang. 2009. PAT: Towards flexible verification under fairness. In *Proc. of CAV*, Vol. 5643. Springer, 709–714.
- [196] T. Sun and W. Yu. 2020. A formal verification framework for security issues of blockchain smart contracts. *Electronics* 9, 2 (feb 2020), 255.
- [197] D. Suvorov and V. Ulyantsev. 2019. Smart contract design meets state machine synthesis: Case studies. (jun 2019). arXiv:1906.02906
- [198] Nick Szabo. 1997. Formalizing and Securing Relationships on Public Networks. *First Monday* 2, 9 (Sep. 1997).
- [199] W. J.-W. Tann, X. J. Han, S. S. Gupta, and Y.-S. Ong. 2018. Towards safer smart contracts: A sequence learning approach to detecting security threats. (nov 2018). arXiv:1811.06632
- [200] Parity Technologies. 2017. A postmortem on the Parity Multi-Sig library self-destruct | Parity Technologies. Retrieved February 13, 2020 from <https://www.parity.io/a-postmortem-on-the-parity-multi-sig-library-self-destruct/>.
- [201] S. Tikhomirov, E. Voskresenskaya, I. Ivanitskiy, R. Takhaviev, E. Marchenko, and Y. Alexandrov. 2018. SmartCheck: Static analysis of Ethereum smart contracts. In *Proc. of IEEE/ACM WETSEB*. 9–16.
- [202] P. Tolmach, Y. Li, S.-W. Lin, and Y. Liu. 2021. Formal Analysis of Composable DeFi Protocols. arXiv:2103.00540
- [203] C. Ferreira Torres, A. K. Iannillo, A. Gervais, and R. State. 2020. Towards smart hybrid fuzzing for smart contracts. arXiv:2005.12156
- [204] C. Ferreira Torres, M. Steichen, and R. State. 2019. The art of the scam: Demystifying honeypots in Ethereum smart contracts. In *Proc. of USENIX Security*. USENIX Association, 1591–1607.
- [205] P. Tsankov, A. Dan, D. Drachslér-Cohen, A. Gervais, F. Bünzli, and M. Vechev. 2018. Securify: Practical security analysis of smart contracts. In *Proc. of ACM CCS*. ACM, 67–82.
- [206] D. Unal, M. Hammoudeh, and M. S. Kiraz. 2019. Policy specification and verification for blockchain and smart contracts in 5G networks. *ICT Express* (jul 2019).
- [207] R. v. d. Meyden. 2019. On the specification and verification of atomic swap smart contracts (extended abstract). In *Proc. of IEEE ICBC*. 176–179.
- [208] E. Vighianisi, M. Ceccato, and P. Tonella. 2020. A federated society of bots for smart contract testing. *J. Syst. Softw.* 168 (2020), 110647.
- [209] F. Vogelsteller and V. Buterin. 2015. ERC-20 token standard. Retrieved July 14, 2020 from <https://eips.ethereum.org/EIPS/eip-20>.
- [210] H. Wang, Y. Li, S.-W. Lin, C. Artho, L. Ma, and Y. Liu. 2019. Oracle-supported dynamic exploit generation for smart contracts. arXiv:1909.06605
- [211] H. Wang, Y. Li, S.-W. Lin, L. Ma, and Y. Liu. 2019. VULTRON: Catching vulnerable smart contracts once and for all. In *Proc. of IEEE/ACM ICSE*. IEEE Press, 1–4.
- [212] S. Wang, L. Ouyang, Y. Yuan, X. Ni, X. Han, and F.-Y. Wang. 2019. Blockchain-enabled smart contracts: Architecture, applications, and future trends. *IEEE Trans. Syst., Man, Cybern. Syst.* (feb 2019), 1–12.
- [213] S. Wang, C. Zhang, and Z. Su. 2019. Detecting nondeterministic payment bugs in Ethereum smart contracts. *Proc. of ACM OOPSLA*, Article Article 189 (Oct. 2019).
- [214] W. Wang, J. Song, G. Xu, Y. Li, H. Wang, and C. Su. 2020. ContractWard: Automated vulnerability detection models for Ethereum smart contracts. *IEEE Trans. on Netw. Sci. Eng.* (2020), 1–1.
- [215] Y. Wang, S. Lahiri, S. Chen, R. Pan, I. Dillig, C. Born, and I. Naseer. 2019. Formal specification and verification of smart contracts for Azure blockchain. Retrieved July 14, 2020 from <https://www.microsoft.com/en-us/research/publication/formal-specification-and-verification-of-smart-contracts-for-azure-blockchain/>.
- [216] X. Wei, C. Lu, F. R. Ozcan, T. Chen, B. Wang, D. Wu, and Q. Tang. 2019. A behavior-aware profiling of smart contracts. In *Security and Privacy in Communication Networks*. Springer International Publishing, 245–258.
- [217] K. Weiss and J. Schütte. 2019. Annotary: A concolic execution system for developing secure smart contracts. In *Proc. of ESORICS*. Springer International Publishing, 747–766.
- [218] M. Wohrer and U. Zdun. 2018. Design patterns for smart contracts in the Ethereum ecosystem. In *Proc. of IEEE iThings/GreenCom/CPSCoM/SmartData*. IEEE, 1513–1520.
- [219] M. Wohrer and U. Zdun. 2018. Smart contracts: Security patterns in the ethereum ecosystem and solidity. In *Proc. of IEEE IWBOSE*, Vol. 2018-Janua. IEEE, 2–8.
- [220] G. Wood. 2014. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum Project Yellow Paper* 151 (2014), 1–32.

- [221] C. R. Worley and A. Skjellum. 2019. Opportunities, challenges, and future extensions for smart-contract design patterns. In *Lecture Notes in Business Information Processing*, Vol. 339. Springer Verlag, 264–276.
- [222] W. Xu and G. A. Fink. 2019. Building executable secure design models for smart contracts with formal methods. arXiv:1912.04051
- [223] X. Xu, C. Pautasso, L. Zhu, Q. Lu, and I. Weber. 2018. A pattern collection for blockchain-based applications. In *Proc. of ACM EuroPLoP*. ACM, ACM.
- [224] K. Yamashita, Y. Nomura, E. Zhou, B. Pi, and S. Jun. 2019. Potential risks of Hyperledger Fabric smart contracts. In *Proc. of IEEE IWBOSE*. 1–10.
- [225] Z. Yang and H. Lei. 2018. Lolisa: formal syntax and semantics for a subset of the Solidity programming language. (2018). arXiv:1803.09885
- [226] Z. Yang, H. Lei, and W. Qian. 2019. A hybrid formal verification system in Coq for ensuring the reliability and security of Ethereum-based service smart contracts. (2019). arXiv:1902.08726
- [227] J. Ye, M. Ma, T. Peng, Y. Peng, and Y. Xue. 2019. Towards automated generation of bug benchmark for smart contracts. In *Proc. of IEEE ICSTW*. IEEE, 184–187.
- [228] X. L. Yu, O. Al-Bataineh, D. Lo, and A. Roychoudhury. 2019. Smart contract repair. arXiv:1912.05823
- [229] F. Zhang, E. Cecchetti, K. Croman, A. Juels, and E. Shi. 2016. Town Crier: An authenticated data feed for smart contracts. In *Proc. of ACM CCS*. ACM, 270–282.
- [230] Q. Zhang, Y. Wang, J. Li, and S. Ma. 2020. EthPloit: From fuzzing to efficient exploit generation against smart contracts. In *Proc. of IEEE SANER*. 116–126.
- [231] X. Zhang, Y. Li, and M. Sun. 2020. Towards a formally verified EVM in production environment. In *Proc. of COORDINATION*. Springer International Publishing, 341–349.
- [232] Z. Zheng, S. Xie, H.-N. Dai, W. Chen, X. Chen, J. Weng, and M. Imran. 2020. An overview on smart contracts: Challenges, advances and platforms. *Future Gener. Comp. Sy.* 105 (2020), 475–491.
- [233] E. Zhou, S. Hua, B. Pi, J. Sun, Y. Nomura, K. Yamashita, and H. Kurihara. 2018. Security assurance for smart contract. In *Proc. of IFIP NTMS*. 1–5.
- [234] J. Zhu, K. Hu, M. Filali, J.-P. Bodeveix, and J.-P. Talpin. 2020. Formal verification of Solidity contracts in Event-B. arXiv:2005.01261
- [235] W. Zou, D. Lo, P. S. Kochhar, X.-B. D. Le, X. Xia, Y. Feng, Z. Chen, and B. Xu. 2019. Smart contract development: Challenges and opportunities. *IEEE Trans. Softw. Eng.* (sep 2019), 1–1.
- [236] N. Zupan, P. Kasinathan, J. Cuellar, and M. Sauer. 2020. *Secure smart contract generation based on Petri Nets*. Springer Singapore, 73–98.