

Information Flow Analysis for Detecting Non-Determinism in Blockchain

Luca Olivieri  

University of Verona, Corvallis Srl

Luca Negrini  

Corvallis Srl

Vincenzo Arceri  

University of Parma

Fabio Tagliaferro  

CYS4 Srl

Pietro Ferrara  

Ca' Foscari University of Venice

Agostino Cortesi  

Ca' Foscari University of Venice

Fausto Spoto  

University of Verona

Abstract

A mandatory feature for blockchain software, such as smart contracts and decentralized applications, is determinism. In fact, non-deterministic behaviors do not allow blockchain nodes to reach one common consensual state or a deterministic response, which causes the blockchain to be forked, stopped, or to deny services. While domain-specific languages are deterministic by design, general-purpose languages widely used for the development of smart contracts such as Go, provide many sources of non-determinism. However, not all non-deterministic behaviours are critical. In fact, only those that affect the state or the response of the blockchain can cause problems, as other uses (for example, logging) are only observable by the node that executes the application and not by others. Therefore, some frameworks for blockchains, such as Hyperledger Fabric or Cosmos SDK, do not prohibit the use of non-deterministic constructs but leave the programmer the burden of ensuring that the blockchain application is deterministic. In this paper, we present a flow-based approach to detect non-deterministic vulnerabilities which could compromise the blockchain. The analysis is implemented in GoLiSA, a semantics-based static analyzer for Go applications. Our experimental results show that GoLiSA is able to detect all vulnerabilities related to non-determinism on a significant set of applications, with better results than other open-source analyzers for blockchain software written in Go.

2012 ACM Subject Classification Security and privacy → Distributed systems security; Theory of computation → Program analysis; Theory of computation → Program verification; Software and its engineering → Software notations and tools

Keywords and phrases Static Analysis, Program Verification, Non-determinism, Blockchain, Smart contracts, DApps, Go language

Digital Object Identifier 10.4230/LIPIcs.ECOOP.2023.16

Funding *Vincenzo Arceri*: Bando di Ateneo per la ricerca 2022, founded by University of Parma, project number: MUR_DM737_2022_FIL_PROGETTI_B_ARCERI_COFIN, *Formal verification of GPLs blockchain smart contracts*

Pietro Ferrara: SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU, iNEST-Interconnected NordEst Innovation Ecosystem funded by PNRR (Mission 4.2, Investment 1.5) NextGeneration EU - Project ID: ECS 00000043, and SPIN-2021 "Static Analysis for Data



© Luca Olivieri, Luca Negrini, Vincenzo Arceri, Fabio Tagliaferro, Pietro Ferrara, Agostino Cortesi, Fausto Spoto;

licensed under Creative Commons License CC-BY 4.0

37th European Conference on Object-Oriented Programming (ECOOP 2023).

Editors: Karim Ali and Guido Salvaneschi; Article No. 16; pp. 16:1–16:25



Leibniz International Proceedings in Informatics

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

46 Scientists” funded by Ca’ Foscari University
 47 *Agostino Cortesi*: SERICS (PE00000014) under the NRRP MUR program funded by the EU - NGEU,
 48 iNEST-Interconnected NordEst Innovation Ecosystem funded by PNRR (Mission 4.2, Investment
 49 1.5) NextGeneration EU - Project ID: ECS 00000043, and SPIN-2021 "Ressa-Rob" funded by Ca’
 50 Foscari University

51 **1** Introduction

52 In the last decade, blockchain software has undergone a notable evolution. In 2008, Bitcoin [33]
 53 introduced a Turing-incomplete low-level language to specify locking conditions that must
 54 hold for a transaction to be accepted by the network [3]. In 2013, Ethereum [8, 4] provided a
 55 Turing-complete bytecode where smart contract rules are enforced by the blockchain consensus.
 56 The execution of the code takes place on the Ethereum Virtual Machine (EVM), resulting in
 57 software identified as *decentralized applications* (DApps). EVM bytecode is supported by high-
 58 level domain-specific languages (DSLs), such as Solidity and Vyper, that have been designed
 59 from scratch for the purpose of being executed in the restricted environment of blockchain.
 60 Subsequently, thanks to frameworks such as Hyperledger Fabric [2], Tendermint [6, 29], and
 61 Cosmos SDK [30], general-purpose programming languages (GPLs) such as Go, Java and
 62 JavaScript can also be used to develop smart contracts and DApps, with Go being the most
 63 popular in industrial blockchains.

64 The popularity of GPLs for writing smart contracts and DApps is steadily increasing.
 65 Their success is mostly due to the maturity of the languages themselves, directly resulting in
 66 wide communities, consolidated tools (such as IDEs and debuggers), and most importantly a
 67 pool of expert and knowledgeable developers that can write highly efficient smart contracts.
 68 Yet, GPLs were not conceived solely for blockchain ecosystems: code that is harmless and
 69 bug-free in other contexts may result in vulnerabilities and errors. Among these, one of
 70 the most insidious is non-determinism. When the result of an operation on a blockchain
 71 is non-deterministic, there is no guarantee that a common state can be reached by the
 72 network’s nodes, possibly preventing it from reaching consensus. This can manifest, among
 73 other possibilities, as transaction failures or denial of service. Nevertheless, not all instances
 74 of non-determinism are intrinsically dangerous: logging the time of a transaction can result
 75 in different timestamps appearing in each node’s logs, but it does not endanger consensus as
 76 it is not *observable* by other nodes. In fact, non-deterministic instructions are problematic
 77 only if they can affect the shared blockchain state.

78 As an example, consider the code in Figure 1, reporting an excerpt of the `ValidateBasic`
 79 method from module `x/authz` (part of the Cosmos SDK versions 0.43.x and 0.44.{0,1}) and
 80 affected by the vulnerability reported in CVE-2021-41135¹. The code is meant to fail the
 81 validation of expired grants. Note that the guard at line 2 involves the local clock of nodes
 82 (`time.Now()`) rather than leveraging the timestamp included in the Block header provided
 83 by the Byzantine Fault Tolerant clock, that is agreed upon by the consensus. As reported in
 84 the official Cosmos forum [12]:

85 *Local clock times are subjective and thus non-deterministic. An attacker could craft*
 86 *many Grants, with different but close expiration times (e.g., separated by a few seconds),*
 87 *and try to exercise the granted functionality for all of them close to their expiration*

¹ <https://nvd.nist.gov/vuln/detail/CVE-2021-41135>.

```

1 func (g Grant) ValidateBasic() error {
2     if g.Expiration.Unix() < time.Now().Unix() {
3         return sdkerrors.Wrap(ErrInvalidExpirationTime, "Time can't be in the past")
4     }
5     // [...]
6 }

```

■ **Figure 1** Cosmos SDK code affected by CVE-2021-41135

88 *time. It is likely in such a scenario that some nodes would consider a grant to have*
 89 *expired while others would not, leading to a consensus halt.*

90 The code was then fixed in version 0.44.2, but is still a clear example of a vulnerability arising
 91 from non-deterministic constructs.

92 The problem of non-determinism in blockchain software is clearly felt by the communities
 93 of the blockchain frameworks treated in this paper. As a representative example, the
 94 Tendermint Core documentation [27], while discussing non-determinism, reports:

95 *While programmers can avoid non-determinism by being careful, it is also possible*
 96 *to create a special linter or static analyzer for each language to check for determinism.*
 97 *In the future we may work with partners to create such tools.*

98 Paper contribution

99 This paper presents a software verification approach based on static analysis for the detection
 100 of non-deterministic vulnerabilities in blockchain ecosystems, covering the most popular
 101 frameworks for developing this kind of software, such as Hyperledger Fabric, Tendermint
 102 Core and Cosmos SDK. We shift the classical focus that has been applied in this context
 103 beyond the mere syntactic absence of non-deterministic constructs. In fact, we aim at
 104 distinguishing *harmful* usages of non-determinism, that is, constructs affecting the blockchain
 105 state and response, from *harmless* ones. As a consequence, the set of alarms issued to the
 106 user sensibly shrinks, as shifting from a syntactic approach towards a semantic one leads to
 107 a sensible reduction in false positives. We propose a semantic flow-based static analysis for
 108 detecting flows from non-deterministic constructs to blockchain state modifiers and response
 109 builders. The choice of a flow-based analysis seems natural when the problem is phrased
 110 as “*is there execution where a non-deterministic value affects the blockchain state or the*
 111 *contract’s response?*”. We thus exploit the well-consolidated literature in this area to adopt
 112 scalable solutions that soundly over-approximate all program executions.

113 We provide a static analyzer implementing our approach: GoLiSA², a sound static analyzer
 114 based on abstract interpretation [10] for Go applications. Intuitively, we use our analyzer’s
 115 fixpoint engine to mark *all* program variables (local variables, objects’ fields, ...) that can
 116 contain values affected, directly or indirectly, by a non-deterministic construct or computation.
 117 Specifically, we can perform a shallower analysis detecting only explicit flows using *Taint*
 118 analysis [43, 14], where non-deterministic constructs and blockchain state modifiers are
 119 modeled as sources and sinks, respectively. Alternatively, we can perform a deeper analysis
 120 able to also detect implicit flows by means of the *Non-interference* analysis [24, 25], where
 121 non-deterministic constructs and blockchain state modifiers are instead modeled as low and

² <https://github.com/lisa-analyzer/go-lisa>

high variables, respectively. Both solutions are implemented in GoLiSA, whose analysis starts by syntactically visiting the input application to annotate all sources and sinks. The annotations are dynamically generated depending on the kind of application of interest (i.e., Hyperledger Fabric, Cosmos SDK, or Tendermint Core). Since there is no predefined set of sources in the target program, both *Taint* analysis and *Non-interference* are parametric: they consider as *harmful* (i.e., tainted or low integrity, depending on the analysis that is to be executed) only variables that are annotated as sources. The fixpoint engine then takes care of propagating values coming from sources on the entirety of the program, exploiting our analyses implementations. After the fixpoint converges, a mapping stating if each program variable is the result of a non-deterministic computation is available at each program point. These are then used by our non-deterministic semantic checkers, that visit the whole application searching for statements annotated with the sink annotation. Whenever one is found, the mappings are used to determine if the values used as parameters of the call are critical or, in the case of *Non-interference*, if the call happens on a critical state.

Our approach, as highlighted by our evaluation, shows a significant decrease of false positives on real-world blockchain applications compared to other analyzers for blockchain non-determinism. The solution has been experimented on a benchmark of more than 600 real-world blockchain programs written in Go. These show that GoLiSA is able to perform the analysis on the totality of smart contracts in this significant benchmark, and to successfully report their non-determinism vulnerabilities.

The analyses are then evaluated in terms of precision of the results (true positive, false positive, and false negative alarms). Based on these criteria, GoLiSA outperforms existing open-source static analyzers for Go blockchain software. Moreover, the evaluation shows that the execution time of the analyses is not impractical for real use cases.

To the best of our knowledge, GoLiSA is the first sound semantic-based static analyzer for blockchain software able to precisely detect critical non-determinism behaviors while scaling to real-world programs.

Summarizing, our contribution is threefold, as we provide:

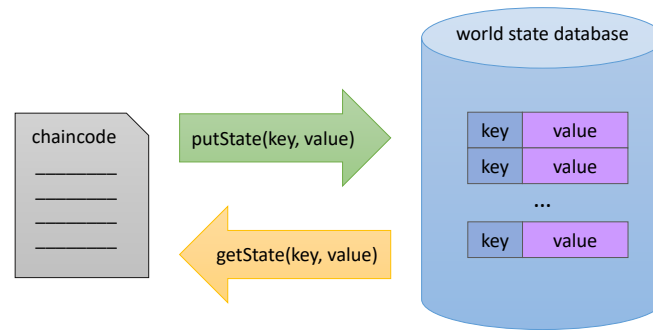
- a detailed investigation on the sources and the sinks that lead to non-determinism issues in the most popular blockchain frameworks;
- a flow-based static analysis for the detection of critical non-determinism behaviors, with two instantiations exploiting different formalizations;
- an open-source sound static analyzer for detecting critical non-deterministic behaviors in blockchain software written in Go.

Paper structure

Sect. 2 reports an overview about blockchain software using Go and the most popular frameworks to develop it. Sect. 3 discusses the problem of non-deterministic behavior in blockchain context. After reporting an overview on information flow analyses, Sect. 4 presents our core contribution for detecting non-deterministic behavior in blockchain software, that relies on GoLiSA. Sect. 5 reports our experimental results. Sect. 6 discusses the related work. Finally, Sect. 7 concludes the paper.

2 Preliminaries: Go and Blockchain

Go (<https://golang.org>) is a statically typed, compiled, open-source, and general-purpose high-level programming language designed by Google to speed up software development, and that is appreciated for its cross-compilation feature. Its versatility and performance



■ **Figure 2** The world state database of Hyperledger Fabric.

167 contributed to its diffusion in the blockchain environment: popular frameworks such as the
 168 Hyperledger Fabric³, Tendermint⁴ and the Cosmos SDK⁵ are written in Go. These rely on
 169 Go to develop efficient smart contracts and DApps, exploiting its high performances.

170 2.1 Blockchain Environments

171 *Hyperledger Fabric* (HF) is a permissioned blockchain framework designed to be adopted
 172 in enterprise contexts, supported by the Linux Foundation and other contributors such as
 173 IBM, Cisco, and Intel. In HF, smart contracts and DApps are written in *chaincode* that
 174 can be implemented in several GPLs such as Go, JavaScript, and Java. In most cases, the
 175 chaincode interacts only with the world state database component of the ledger, and not
 176 with the transaction log [26]. Go is currently the most popular language on GitHub related
 177 to *chaincode*⁶, as Go smart contracts are the best performing ones [23].

178 *Tendermint Core*, recently rebranded as *Ignite*, is a platform for building blockchain
 179 nodes, supporting both public and permissioned *proof-of-stake* (PoS) networks. It is a
 180 Byzantine Fault Tolerant (BFT) middleware that separates the application logic from the
 181 consensus and networking layers, allowing one to develop blockchain applications written in
 182 any programming language, including Go, and replicate them on many machines [7].

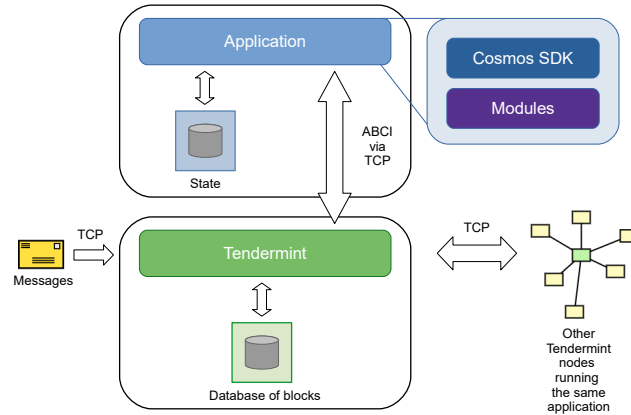
183 *Cosmos SDK* is an open-source Go framework that eases the development of blockchain
 184 applications while optimizing their execution by running them on Tendermint Core. As shown
 185 in Figure 3, Cosmos SDK abstracts all the boilerplate code needed to set up a Tendermint
 186 Core node, allowing for customized protocol configurations. The programming style follows
 187 the object-capability model, where the security of subcomponents is imperative, especially
 188 those belonging to the core library. Cosmos SDK is a framework for DApps, supporting
 189 different functionalities through highly customizable modules (that can also manage smart
 190 contracts).

³ <https://www.hyperledger.org/use/fabric>

⁴ <https://tendermint.com/>

⁵ <https://v1.cosmos.network/sdk>

⁶ Querying the keyword `chaincode` on GitHub (<https://github.com/search?q=chaincode>) results in more than 2100 repositories, about half of which are written in Go. Accessed: 01-12-2022.



■ **Figure 3** Cosmos SDK architecture

191 2.2 Blockchain Consensus

192 Consensus protocols ensure the validity and authenticity of transactions performed in the
 193 blockchain, as they check results of smart contracts or DApps computations through the
 194 state of the network's nodes. If a given number of nodes agree on the final state, consensus is
 195 reached and the transaction is validated. Otherwise, it is discarded and the nodes proposing
 196 spurious states are excluded from the network. When consensus cannot be reached, the
 197 blockchain either forks or halts. Deterministic execution is thus required for software that
 198 runs in a blockchain, as it guarantees that, when starting from a common state, the same
 199 result is reached in any distinct blockchain node, avoiding inconsistencies among peers and
 200 consensus failures. Nevertheless, GPLs provide several components that can explicitly lead
 201 to non-determinism, such as (pseudo-)random number generators or external computations.
 202 Furthermore, even methods that are explicitly sequential and deterministic pose a threat
 203 when executed on different nodes, such as the `time.Now()` call from Figure 1. Despite
 204 these threats, popular blockchain frameworks such as HF and Cosmos SDK do not enforce
 205 particular restrictions on the usage of non-deterministic methods and components.

206 3 Non-Deterministic Behaviors in Blockchain Software: Sources and 207 Sinks

208 When trying to prevent non-deterministic vulnerabilities, a first solution is to limit the
 209 expressiveness of the GPL by either black- or white-listing APIs and constructs. Consider
 210 the Go snippets reported in Figure 4. Both fragments rely on the `time` API to retrieve
 211 a timestamp from the host system. In general, the results of calls to the `time` API are
 212 subjective to the node executing them, and they might lead to blockchain non-determinism
 213 due to different system settings (e.g., time, date, time zones, ...) or due to nodes executing
 214 the code at slightly different times. Specifically, Figure 4a shows a safe usage of the `time` API:
 215 the timestamp is only used for logging with no observable consequences on the blockchain
 216 state or the execution result. Instead, Figure 4b reports a problematic usage of the API,
 217 as the timestamp is stored in the blockchain using `PutState`, an HF-specific function that
 218 updates the shared network state. Since timestamps could differ on each node, this potentially
 219 leads to inconsistent executions (i.e., different blockchain states or execution results), causing

```

1 func transfer(from, to Address, value int64, stub *shim.ChaincodeStub) {
2     start := time.Now()
3     //... transfer operations that takes some milliseconds ...
4     elapsed := time.Now().Sub(start)
5     log.Println("Time elapsed for the transfer operations: ", elapsed)
6 }

```

(a) Example of safe use of the time API

```

1 func transfer(from, to Address, value int64, stub *shim.ChaincodeStub) {
2     t := time.Now()
3     //... transfer operations ...
4     err := shim.PutState("transaction-time", t)
5     //... other operations ...
6 }

```

(b) Example of issue of non-determinism with the time API

■ **Figure 4** Examples of harmless and harmful non-determinism in blockchain

220 transaction failure⁷.

221 It should thus be evident that identifying sources of non-determinism and preventing
 222 their usage is not enough when we aim at discerning between harmful and harmless non-
 223 deterministic constructs. In fact, one should also recognize how these are used, determining if
 224 they can influence the shared blockchain state. In the rest of this section we discuss, for each
 225 blockchain framework presented in Section 2, (i) the constructs that generate potentially
 226 harmful non-determinism (that is, *sources* of non-deterministic values), and (ii) the blockchain
 227 state modifiers and response builders (i.e., statements that make a transaction succeed or fail),
 228 namely *sinks* that are sensitive to non-determinism⁸. This will prepare the ground for the
 229 core contribution of this paper: a static approach to detect critical usages of non-determinism
 230 in blockchain software, reported in Section 4.

231 3.1 Sources of Non-Determinism

232 The sources of non-determinism can be logically split in two families, the first being related
 233 to the combination of framework and GPL adopted to develop the software. This family
 234 comprises a set of constructs and APIs allowed by the framework that may break the
 235 consensus during the execution of smart contracts or DApps. In Go, these are:

- 236 ■ *iteration over maps* that, being the iteration order unspecified⁹, is not guaranteed to be
 237 deterministic;
- 238 ■ *parallelization and concurrency*, that can lead to race conditions on shared resources,
 239 thus creating non-determinism on the computed values;
- 240 ■ *global variables*, that may change innately and cause inconsistencies to the results, since
 241 they depend on the application state of a peer and not on that of the blockchain [32, 5].
- 242 ■ *random value generators*, that can potentially be allowed in smart contracts [9] to employ
 243 custom logic while being non-deterministic by-definition.

244 The second family instead involves statements related to the underlying environment,
 245 such as file systems, operating systems, databases, and Internet connections. While these are

⁷ In this case, the `GetTxTimestamp` method from the HF API should have been used instead of `time.Now`.

⁸ The complete list of sources and sinks of non-determinism is available at <https://github.com/lisa-analyzer/go-lisa/blob/master/go-lisa/sources-sinks.md>.

⁹ https://golang.org/ref/spec#For_statements

<i>Level</i>	<i>Category</i>	<i>Package</i>	<i>Statements/Methods</i>
<i>Framework/Language</i>	Map iteration	-	<code>range</code> on map
	Parallelization/concurrency	-	<code>go</code> (Go routine), <code><-</code> (channel)
	Random number generation APIs	<code>math/rand</code> , <code>crypto/rand</code>	*
	Global variables	-	-
<i>Environment</i>	File system APIs	<code>io</code> , <code>embed</code> , <code>archive</code> , <code>compress</code>	*
	OS APIs	<code>os</code> , <code>syscall</code> , <code>internal</code> , <code>time</code>	*
	Database APIs	<code>database</code>	*
	Internet APIs	<code>net</code>	*

■ **Table 1** Potential non-deterministic behaviors related to Go

246 not intrinsically non-deterministic, they become dangerous when their result is expected to
 247 be consistent on different environments. These comprise APIs handling:

- 248 ■ *file systems*, as the program might rely on files that are not present on all nodes, as they
 249 might have been deleted, edited, moved, or there might be insufficient disk space causing
 250 any operation to fail;
- 251 ■ *operating systems* (OS), since the blockchain might operate on various hosts and language
 252 APIs could return different results on each OS (e.g., time and date methods could return
 253 different values if nodes are not synchronized);
- 254 ■ *databases*, where records might be deleted, edited, or contain different data;
- 255 ■ *Internet connections*, as networking setup or errors could cause some addresses to be
 256 unreachable on few nodes of the network.

257 Table 1 summarizes the instructions and libraries of Go¹⁰ that we considered as cases of
 258 non-determinism, where * represents the entirety of the package. For the sake of simplicity,
 259 the table reports instructions and packages omitting the full signatures of each method.
 260 Note that only few methods within those packages lead to non-deterministic behaviors: for
 261 instance, most methods from package `time` handling dates and times do not pose a threat
 262 in smart contracts and DApps, and are in fact quite common. However, operations such
 263 as retrieving the current time of the OS (i.e., methods `Since`, `Now`, `Until`) are potentially
 264 dangerous.

265 3.2 Sinks of Non-Determinism

266 Sinks of non-determinism comprise constructs and APIs with the ability of both modifying
 267 the common state of the blockchain or having an impact on the response of blockchain
 268 networks. While the former is inherently involved in consensus protocols, the execution
 269 of code within the blockchain does not necessarily change the shared state (e.g., functions
 270 that simply read a value). However, the execution may lead to non-deterministic responses,
 271 compromising the consensus of the network, as in the simple example reported in Figure 5.
 272 Table 2, where the **Critical point** column identifies what part of the API should not receive
 273 non-deterministic values, summarizes the main instructions and components that we consider
 274 as sinks for non-determinism.

¹⁰The full list of Go APIs sources considered in our analyses is available at https://github.com/lisa-analyzer/go-lisa/blob/master/go-lisa/src/main/resources/for-analysis/nondeterm_sources.txt. The list consider API until Go version 1.17.


```

1 func (s *SmartContract) transaction(APIStub shim.ChaincodeStubInterface) sc.Response {
2
3     if rand.Int() % 2 == 0 {
4         return shim.Error("Fail")
5     } else {
6         return shim.Success(nil)
7     }
8 }

```

■ **Figure 5** Example of issue of non-determinism related to the blockchain response

Framework	Package	Type/Interface	Statements/Methods	Critical point
<i>Hyperledger Fabric</i>	shim	ChaincodeStubInterface	PutState DelState PutPrivateData DelPrivateData Success Error	parameters parameters parameters statement statement
<i>Tendermint Core</i>	abci/types	Application	ResponseBeginBlock ResponseDeliverTx ResponseEndBlock ResponseCommit ResponseCheckTx	instance returned instance returned instance returned instance returned
<i>Cosmos SDK</i>	types	KVStore	Set Delete	parameters parameters
	kv, dbadapter, gaskv, iavl, listenkv, prefix, tracekv,	Store	Set Delete	parameters parameters
	types/errors		ABCIError Redact ResponseDeliverTx ResponseCheckTx WithType Wrap Wrapf	statement statement statement statement statement statement

■ **Table 2** Main sinks for blockchain software written in Go

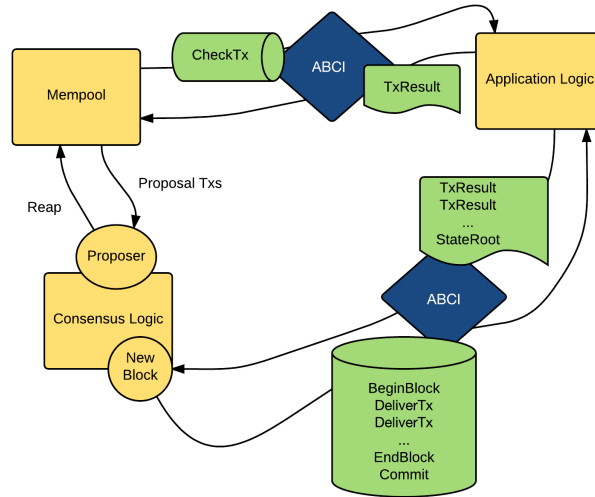
275 3.2.1 Hyperledger Fabric APIs for Go

276 In HF, chaincode executes transaction proposals against world state data that may change its
 277 state. Programmatically, interface `ChaincodeStubInterface` from the HF Go APIs enables
 278 access and modification of the blockchain state. Table 2 reports the current components (as
 279 of version 2.4) involved in the data-write proposal. The semantics of these components does
 280 not affect the blockchain state until the transaction is validated and successfully committed.
 281 Hence, if these components lead to different results (i.e., changes to the shared state) due
 282 to non-determinism, consensus will not validate the transaction and no new state will be
 283 committed. Regarding the response statements, HF provides the `Success` and `Error` methods
 284 to yield successful and failed transaction responses, respectively.


285 3.2.2 Tendermint Core APIs for Go

286 Tendermint Core is a middleware with no explicit access to application state by design,
 287 enabling communication through the *Application BlockChain Interface* (ABCI¹¹). Figure 6
 288 depicts the consensus process used to validate and store a transaction using the ABCI
 289 methods. As reported in the official documentation [27] of Tendermint Core v. 0.35.1, only

¹¹ <https://github.com/tendermint/tendermint/blob/7983f9cc36c31e140e46ae5cb00ed39f637ef283/docs/introduction/what-is-tendermint.md#abci-overview>.



■ **Figure 6** ABCI methods and consensus flow



■ **Figure 7** Main store of Cosmos SDK

290 `BeginBlock`, `DeliverTx`, `EndBlock`, and `Commit` must be strictly deterministic to ensure
 291 consensus. Although the logic of these methods is different, they possess similar structure:
 292 they all accept a request and return a response (`ResponseBeginBlock`, `ResponseDeliverTx`,
 293 `ResponseEndBlock`, `ResponseCommit`), with the latter that must be deterministic.

294 3.2.3 Cosmos SDK APIs

295 Cosmos SDK handles both the application and the blockchain state through the *store*¹². At
 296 a high level, the store is a set of key-value pairs used to store and retrieve data, implemented
 297 by default as a *multistore* (i.e., a store of stores), as shown in Figure 7. The multistore
 298 encapsulation enables modularity of the Cosmos SDK, as each module declares and manages
 299 its own subset of the state using specific keys. Keys are typically held by *keepers*, a Cosmos
 300 SDK abstraction with the role of managing access to the multistore’s subset defined by each

¹² <https://github.com/cosmos/cosmos-sdk/blob/2b24afad075894dd1727d057f87e2be24238016f/docs/core/store.md>.

```

1  var l, h
2  h := 1

```

(a)

```

1  var l, h
2  if l = true then
3    h := 3
4  else
5    h := 42

```

(b)

```

1  var l, h
2  if h = 1 then
3    (* time-consuming work *)
4  l := 0

```

(c)

■ **Figure 8** Example of (a) explicit, (b) implicit, and (c) side channel flows, where `h` and `l` represent secret and public variables respectively

301 module. The `Store` type is declared in several packages (e.g., `kv`, `tracekv`, `gaskv`, `ival`),
 302 with all definitions implementing the `KVStore` interface. The latter provides common APIs
 303 to access and modify the state of the blockchain using methods such as `Set` and `Del`. As for
 304 responses, Cosmos provides several methods (such as `ABCIError`, `Wrap`, `ResponseDeliverTx`)
 305 in package `types/errors` to return failed transaction responses.

306 4 Information Flow Analysis for Non-Determinism Detection

307 In this section we introduce and discuss our approach for detecting non-deterministic behaviors
 308 in blockchain software. In particular, we consider non-determinism as *critical* only if a non-
 309 deterministic value can affect the blockchain state, either directly (i.e., being stored inside
 310 the state) or indirectly (e.g., guarding the execution of state updates). Any other usage of
 311 non-determinism is considered safe, as it does not affect the blockchain state or response. As
 312 such, when mentioning non-determinism in the remainder of the chapter, we implicitly refer
 313 to its critical version. We rely on information flow analysis for detecting values originating
 314 from sources of non-determinism that can affect the state of the blockchain. We only focus
 315 on static analyses, since they soundly over-approximate all possible behaviors of target
 316 programs and can thus give guarantees about the absence of such behaviors. We instantiate
 317 two types of analyses: a *Taint* analysis, able to capture the so-called *explicit flows*, and a
 318 *Non-interference* analysis, that can also detect *implicit flows*.

319 4.1 An Overview on Information Flow

320 Information flow analyses [11, 38] address the problem of understanding how information
 321 *flows* from one variable to another during a program’s execution. These analyses usually
 322 partition the space of program variables into *private* (or secret) and *public*, with the latter
 323 being accessible to — and in some cases also modifiable by — an external attacker. The
 324 goal of these analyses is then to find program executions where information *flows* from one
 325 partition to the other, that is, where values of variables from one partition can affect the
 326 values of variables from the other one. Figure 8 reports examples¹³ of the three main types
 327 of flows, namely:

- 328 ■ *explicit flow*: when a secret variable is assigned to a value obtained starting from public
 329 variables;
- 330 ■ *implicit flow*: when an assignment to a secret variable is conditionally executed depending
 331 on values of public variables;
- 332 ■ *side channel*: where some observable properties of the execution, e.g., the amount of
 333 computational resources used, depends on the values of some secret variables.

¹³[https://en.wikipedia.org/wiki/Information_flow_\(information_theory\)](https://en.wikipedia.org/wiki/Information_flow_(information_theory)).

334 In general, the term *source* is traditionally used for variables holding values that one
 335 wants to track along program executions, while *sink* is used to describe locations where
 336 values coming from sources should not flow. Using this terminology, when the property of
 337 interest ensures the *integrity* of secret variables, information flow analyses can be instantiated
 338 using public variables as sources and private ones as sinks, exactly as in Fig. 8 and in the list
 339 above. These are able to detect situations where (i) a possibly corrupted value provided by a
 340 malicious attacker could be stored into variables whose content is supposed to be reliable, or
 341 (ii) such a value governs the update to private variables. If, however, one wants to ensure the
 342 *confidentiality* of secret variables, the same analyses can be recasted with private variables
 343 acting as sources and public ones as sinks, thus searching for flows in the opposite direction.
 344 The target of the analysis is then to find disclosures of private data to external entities.

345 In the context of non-deterministic behaviors in blockchain environments, information
 346 flow analyses can be used to detect when non-deterministic values end up or affect the
 347 blockchain’s state, thus checking the *integrity* of that state w.r.t. non-deterministic values.
 348 As such, we are interested in information flowing from public to private variables, and
 349 we will use *sources* to identify ones that are initialized to non-deterministic values and
 350 *sinks* to identify all variables that have an effect on the blockchain’s state. Moreover, we
 351 will focus on explicit and implicit flows. In fact, side channels are typically studied to
 352 detect secret information leaking through, for instance, execution time, thus violating the
 353 confidentiality of that information instead of its integrity. On the other hand, explicit and
 354 implicit flows identify non-deterministic values that are either used to update the blockchain’s
 355 state or a transaction’s result, or that govern their execution. As a concrete example, recall
 356 the code from Figure 1: the vulnerability presented there is an implicit flow since the
 357 blockchain’s state is not directly updated with non-deterministic values, but the execution of
 358 the update (i.e., the `return` statement) is conditional to some non-deterministic value (i.e.,
 359 `g.Expiration.Unix() < time.Now().Unix()`).

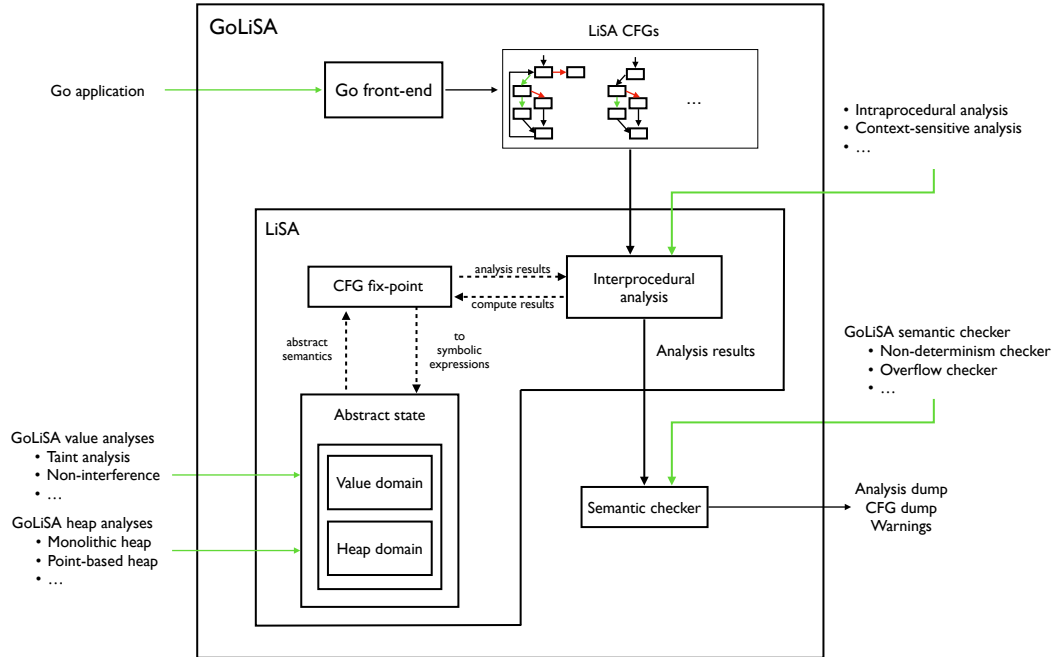
360 In the following, we introduce two well-established information flow analyses that we will
 361 use for non-determinism detection.

362 4.1.1 Non-Interference

Non-interference [24, 25] is a notion of security capturing the idea that if computations over
 private information are independent from public information, then no leakage of the former
 can happen. In simple terms, after partitioning the space of inputs of a program P into *low*
 (private or secret, denoted by \mathbb{L}), and *high* (public or available to anyone, denoted by \mathbb{H}),
Non-interference is satisfied if changes in the high input do not affect the observable (i.e.,
 low) output of the program:

$$\forall i_{\mathbb{L}} \in \mathbb{L}, \forall i_{\mathbb{H}}, i'_{\mathbb{H}} \in \mathbb{H} . P(i_{\mathbb{L}}, i_{\mathbb{H}})_{\mathbb{L}} = P(i_{\mathbb{L}}, i'_{\mathbb{H}})_{\mathbb{L}}$$

363 This notion is often instantiated in language-based security by partitioning the space of
 364 program variables between \mathbb{L} and \mathbb{H} , and finding instances of explicit or implicit flows between
 365 these partitions. Such analysis computes, for each program point, a mapping from variables
 366 to the information level they hold (low or high), while also keeping track of an execution state
 367 depending on the information level of the Boolean conditions that guard the program point.
 368 Violations of *Non-interference* for integrity can then be detected whenever an assignment to
 369 a variable in \mathbb{H} either (i) assigns a low value (that is, an expression involving variables in \mathbb{L}),
 370 or (ii) happens with a low execution state (that is, guarded by at least a Boolean condition
 371 that involves variables in \mathbb{L}), thus identifying both explicit and implicit flows. This can be
 372 formalized as a type system for security [38].



■ **Figure 9** GoLiSA overall execution

373 4.1.2 Taint Analysis

374 *Taint* analysis [43, 14] is an instance of information flow analysis that can be seen as
 375 simplification of *Non-interference* considering only explicit flows. In this context, variables
 376 are partitioned into *tainted* and *untainted* (or *clean*), with the former representing variables
 377 that can be tampered with by an attacker and the latter representing variables that should
 378 not contain tainted values across all possible program executions. Roughly, *Taint* analysis
 379 corresponds to the language-based *Non-interference* instantiation without the execution state,
 380 thus unable to detect implicit flows. *Taint* has been instantiated to detect many defects
 381 in real-world software, such as web-application vulnerabilities [16], privacy issues [22] (also
 382 related to GDPR compliance [20]), and vulnerabilities of IoT software [17].

383 4.2 The GoLiSA Static Analyzer

384 GoLiSA¹⁴ is an abstract interpretation [10] based static analyzer for Go applications, on
 385 which we will rely for the rest of the paper for reasoning about blockchain software written
 386 in Go. In this section, we present its architecture and its main feature. GoLiSA relies
 387 on LiSA [19, 34] (Library for Static Analysis¹⁵), a Java library that provides a complete
 388 infrastructure for the development of static analyzers based on abstract interpretation. In
 389 particular, LiSA implements several standard components of abstract interpretation-based
 390 analyzers, such as an extensible control-flow graph representation (CFG), a common analysis
 391 framework for the development of new static analyses, and fixpoint algorithms on LiSA
 392 CFGs.

¹⁴ Available at <https://github.com/lisa-analyzer/go-lisa>

¹⁵ LiSA project and documentation available at <https://github.com/lisa-analyzer/lisa>

393 The high-level analysis process of GoLiSA is reported in Fig. 9. The analysis starts with
 394 the *Go front-end* (a sub-component of GoLiSA) that compiles Go source code into LiSA
 395 CFGs and defines the semantics, types and language-specific algorithms that implement the
 396 Go execution model, capturing the peculiarities of Go in order to make them understandable
 397 to LiSA (e.g., scoping and shadowing of variables¹⁶). These CFGs are then passed to LiSA,
 398 that analyzes them in a generic language-independent fashion. Roughly, CFGs are passed
 399 to an *interprocedural analysis*, a component that cooperates with a *call graph* to resolve
 400 calls and compute their results. The interprocedural analysis computes fixpoints over CFGs
 401 according to some implementation-specific logic (e.g., modularly, relying on call chains, ...).
 402 Each individual fixpoint relies on language-specific analysis-independent semantics for CFG
 403 nodes, that is directly provided by GoLiSA: each node is rewritten into a sequence of *symbolic*
 404 *expressions*, modelling the effects that executing a high-level instruction has on the program
 405 state through low-level atomic semantic operations. Each of these symbolic expressions is fed
 406 to an *abstract state* [15], a combination of an abstract domain modelling the dynamic memory
 407 of the program (*heap domain*, e.g., point-based heap analysis [1]) and one for tracking values
 408 of program variables and memory locations (*value domain*, e.g., intervals [10]). The abstract
 409 state and its underlying domains compute a sound over-approximation of the expression's
 410 effects according to their specific logic, and this can later be exploited by *semantic checks* to
 411 issue warnings that are of interest for the user. All analysis components (interprocedural
 412 analysis, call graph, abstract state, heap domain, value domain and semantic checks) are
 413 part of LiSA's configuration, enabling modular composition and implementation of each
 414 component.

415 4.3 GoLiSA for Non-Deterministic Behaviors Detection

416 At this point, we are in position to instantiate GoLiSA for the static detection of non-
 417 deterministic behaviors in blockchain software. The core idea of our solution is to track the
 418 values generated by the hotspots identified in Section 3.1 during the execution of a program
 419 using either *Taint* analysis or *Non-interference*. Similarly, after the analysis completes, we
 420 can use a semantic checker to exploit the abstract information provided by the domain of
 421 choice, checking if any of the sinks specified in Section 3.2 receives one such non-deterministic
 422 value as parameter or, in the case of *Non-interference*, if the sink is found in a *low* execution
 423 state.

424 GoLiSA's analysis is instantiated as follows:

- 425 ■ *Taint* analysis and *Non-interference* are implemented as value domains, both of them being
 426 non-relational domains (i.e., mapping from variables to abstract values — taintedness
 427 and integrity level respectively — with no relations between different variables), with
 428 *Non-interference* keeping track of the abstractions for each guard;
- 429 ■ field-insensitive program point-based heap domain (Section 8.3.4 of [37]), where any
 430 concrete heap location allocated at a specific program point is abstracted to a single
 431 abstract heap identifier;
- 432 ■ context-sensitive [39, 28] interprocedural analysis, abstracting full call-chain results until
 433 a recursion is found;
- 434 ■ runtime types-based call graph, using the runtime types of call receivers to determine
 435 their targets;

¹⁶https://go.dev/ref/spec#Declarations_and_scope

436 ■ two semantic checkers, for *Taint* analysis and *Non-interference*, that scan the code in
437 search for sinks, checking the taintedness or integrity level of each sink.

438 The analysis begins by visiting the input program to detect the statements annotated as
439 sources and propagating the information from them. The analyses produce, for each program
440 point, a mapping stating if each variable is the result of a non-deterministic computation.
441 These mappings are then used by our semantic checkers, that visit the program in search for
442 statements annotated as sinks. When one is found, the mappings are used to determine if
443 values used as parameters of the call are critical or, in the case of *Non-interference*, if the
444 call happens on a critical state. The choice of the analysis to run (and thus of the checker to
445 execute) is left to the user.

446 For instance, let us consider the fragment reported in Figure 4a. At line 5, despite
447 variable `elapsed` being marked as tainted, no warning is raised by GoLiSA regardless of the
448 chosen analysis, as it does not reach any sensitive sink. Instead, analyzing the fragment from
449 Figure 4b results in the following alarm:

```
450 The value passed for the 2nd parameter of this call is tainted,  
451 and it reaches the sink at parameter 'value'
```

454 The warning is issued with both analyses, since variable `t` is marked as tainted and reaches a
455 blockchain state modifier through an explicit flow.

456 Consider now the example reported in Figure 1. Here, no explicit flow happens at
457 line 3, that contains the blockchain state modifier `Wrap`, but its execution depends on the
458 non-deterministic value used in the condition at line 2, that is, `time.Now().Unix()`. As this
459 is an implicit flow, the *Taint* analysis is not able to detect it. GoLiSA will however discover
460 it with *Non-interference*, raising the following alarm:

```
461 The execution of this call is guarded by a tainted condition,  
462 resulting in an implicit flow
```

465 4.4 Detection of Sources and Sinks in GoLiSA

466 To exploit information flow analyses, the analyzer must know which are the sources and sinks
467 of the program. In this regard, GoLiSA provides a solution based on annotations, marking
468 the corresponding statements as sources and sinks. In the following, we describe how GoLiSA
469 annotates sources (Table 1) and sinks (Table 2) depending on their types.

470 Methods and functions

471 As shown in Tables 1 and 2, all sinks and several sources correspond to functions and
472 methods of APIs from either the Go runtime or the blockchain frameworks. GoLiSA contains
473 a list of the signature of these functions and methods and it automatically annotates the
474 corresponding calls in the program by syntactically matching them. While we rely on manual
475 annotations, they can also be generated using automated tools (e.g., SARL [18]). For instance,
476 when GoLiSA iterates over the following snippet, it is able to discover the call to `time.Now`,
477 that gets annotated as source, and the one to `PutState`, whose parameters get annotated as
478 sinks:

```
479 key := "key"  
480 tm := time.Now()  
481 stub.PutState(key, []byte(tm))
```

484 Then, the information flow analysis propagates taintedness from the return value of `time.Now`
485 to the second parameter of `PutState`, thus issuing an alarm at line 3.

486 **Map iterations**

487 To detect iterations over maps, one needs to reason about typing. GoLiSA exploits runtime
 488 types inferred by the analysis to identify **range** statements happening over maps. If a map
 489 iteration occurs, that is, if the object in a **range** statement is inferred to be a map, then
 490 GoLiSA marks as sources the variables used to store keys and values of the map. Consider
 491 as an example the following code snippet:

```
492 s := ""
4931 kvs := map[string]string{"a": "hello", "b": "world!"}
4942 for k, v := range kvs {
4953     s += v
4964 }
4975 stub.PutState("key", []byte(s))
4986
```

500 While analyzing the code, **range** statements are checked for the types of their parameter.
 501 GoLiSA annotates as sources both **k** and **v**, as **kvs** is inferred to be a map, while the sink
 502 at line 6 is detected through already discussed annotations. Information flow analyses can
 503 then propagate the taintedness from **v** to **s**, that in turn flows to the second parameter of
 504 **PutState**, issuing an alarm at line 6.

505 **Global variables**

506 GoLiSA syntactically annotates every global variable appearing in the program as a source
 507 of non-determinism, as their value could be modified independently on each peer. For
 508 instance, in the following code, the value of global variable **glob** could differ from peer to
 509 peer depending on the number of times function **inc** has been executed. This can happen as
 510 not all peers simulate the same transaction, for instance due to differences in the endorsement
 511 policy of each peer [32].

```
512 var glob string
5131 func inc() {
5142     glob += "a"
5153 }
5164 func (s *SmartContract) transaction(stub shim.ChaincodeStubInterface) sc.Response {
5175     stub.PutState("key", []byte(glob))
5186 }
5197
```

521 Before the analysis, GoLiSA iterates over all program components, annotating **glob** as a
 522 source. The sink at line 6 is annotated as sink as previously discussed. Then, the information
 523 flow analysis propagates taintedness from **glob** to the second parameter of the call to
 524 **PutState**, raising an alarm at line 6.

525 **Go routines**

526 GoLiSA inspects the code of Go routines, checking the scope of variables they use. If these are
 527 defined outside the routine using them, they are effectively shared among threads, potentially
 528 leading to race conditions or non-deterministic behaviors. Hence, GoLiSA annotates the
 529 such variables as sources. As an example, the following snippet defines and invokes a simple
 530 Go routine that modifies a variable defined in an enclosing scope:

```
531 s := ""
5321 go func(){
5332     for i := 1; i <= 10000; i++ {
5343         s += "0"
5354     }
5365 }
5376 stub.PutState("key", []byte(s))
5387
```


540 When GoLiSA finds the Go routine, it checks the scopes of each variable, inferring that `s`
 541 is declared outside the routine itself. Hence, GoLiSA annotates `s` at line 1 as source, while
 542 the sink at line 7 is annotated as previously discussed. Then, the information flow analysis
 543 propagates taintedness from `s` to the second parameter of `PutState`, issuing an alarm at line
 544 7 since the value of `s` depends how many times the Go routine has executed the loop body.

545 Go channels

546 Channels are pipes that connect concurrent Go routines. Operator `<-` allows interaction with
 547 channels to retrieve a value from them, blocking until one is available. GoLiSA annotates as
 548 sources the instructions reading values from channels, as the order in which these are written
 549 is intrinsically non-deterministic. Consider the following example:

```
550 c := make(chan int)
551 go myroutine1(c)
552 go myroutine2(c)
553 x, y := <- c, <- c
554 stub.PutState("key", []byte(x))
555
```

557 GoLiSA iterates over the program searching for occurrences of the operator `<-`. It then
 558 annotates variables `x` and `y` as sources, as they receive a value from channel `c`. The sink at
 559 line 5 is detected as previously discussed. The information flow analysis can then propagate
 560 taintedness from `x` to the second parameter of `PutState`, resulting in an alarm at line 5.

561 5 Experimental Evaluation

562 In this section, we discuss the experimental evaluation of the information flow analyses
 563 implemented in GoLiSA to detect non-determinism issues in real-world blockchain software.
 564 First, we study them from a quantitative point of view, on a set of 651 real-world HF
 565 smart contracts retrieved from public GitHub repositories. The evaluation focuses on the
 566 HF framework since, to the best of our knowledge, it is the only framework supported by
 567 several static analyzers detecting non-determinism issues. This will allow us to compare
 568 GoLiSA against state-of-the-art static analyzers in this domain. Furthermore, HF is currently
 569 the most popular and widespread blockchain framework among public GitHub repositories,
 570 with most smart contracts written in Go. Nevertheless, GoLiSA provides support also for
 571 detecting non-determinism behaviors for Cosmos SDK and Tendermint Core smart contracts
 572 and DApps.¹⁷

573 We compare GoLiSA with two open-source static analyzers for chaincodes, namely
 574 `Revive^CC` and `ChainCode Analyzer`. The experiments show that GoLiSA produces
 575 more precise results in detecting non-deterministic behaviors, outperforming existing static
 576 analyzers.

577 Then, we evaluate the quality of our results on two specific real-world applications, to show
 578 how the static analyses discussed in Section 4 work and how the information is propagated
 579 in smart contracts. In particular, we selected the first application from the HF benchmark,
 580 while the second one is a Cosmos SDK application.¹⁸

581 All the experiments was performed on a HP EliteBook 850 G4 equipped with an Intel
 582 Core i7-7500U at 2,70/2,90 GHz and 16 GB of RAM running Windows 10 Pro 64bit, Oracle
 583 JDK version 13, and Go version 1.17.

¹⁷ An industrial application of GoLiSA for detecting non-determinism in Cosmos SDK can be found here [36].

¹⁸ The example reported in Figure 1 contains a snippet of code of this application

Analysis	#A	#U	ET	AT	#W	#TP	#FP	#FN
<i>Taint</i>	68	583	2h:15m:03s	12.45s	173	118	55	7
<i>Non-interference</i>	69	582	2h:25m:18s	13.39s	195	124	71	0

■ **Table 3** Analysis evaluation

584 5.1 Quantitative Evaluation

585 The experimental artifact set has been retrieved from 954 GitHub repositories, by querying
586 for the *chaincode* keyword, as smart contracts are called in HF, and selecting chaincodes
587 from unforked Go repositories only¹⁹, that include the *Invoke* and *Init* methods: these are
588 the transaction requests' entry points for chaincodes.²⁰ Then, we filtered out files unrelated
589 to smart contracts and removed chaincodes not analyzable due of failures either GoLiSA
590 or the tools discussed in Sect. 5.1.1. In particular, GoLiSA failures on such chaincodes are
591 due to current missing support of high-order functions, recursion, and C code invocation via
592 the built-in Go `cmd/go` package.²¹ This resulted in a benchmark consists of 651 chaincodes
593 only (~ 167391 LoCs), that, from here on, we refer to as $\mathbb{H}\mathbb{F}$. Then, each chaincode has been
594 manually inspected before applying GoLiSA to search for critical non-deterministic behavior.
595 In particular, for each chaincode, we manually searched for sources of non-determinism (if
596 present) and checked if the result of the corresponding instructions could have an impact
597 (i.e., an update) on the blockchain global state or on the response. If so, we classified this
598 behaviour as critical/harmful. On the selected benchmark, we have found a total of 124
599 critical/harmful non-deterministic behaviours. In our evaluation, a warning raised by an
600 analyzer has been classified as true positive (TP) if it was part of the 124 critical behaviours
601 mentioned above, and as false positive (FP) if not. All the critical behaviours, part of the
602 124 manually detected, for which there was no warning, have been marked as false negative
603 (FN).

604 Table 3 reports the results of the experimental evaluation of GoLiSA over the benchmark
605 $\mathbb{H}\mathbb{F}$, where **#A** is the number of affected chaincodes (i.e., chaincodes where at least a warning
606 was issued), **#U** is the number of unaffected chaincodes (i.e., chaincodes where no warning
607 was raised), **ET** is the total execution time, **AT** is the average execution time, **#W** is the
608 total number of warnings issued, **#TP** is the number of true positives among the raised
609 warnings, **#FP** is the number of false positives among the raised warnings, and **#FN** is the
610 number of false negatives. In terms of execution time, the analyses performed averagely in
611 around 15 seconds per chaincode. The experiments shows that *Non-interference* performs
612 better than *Taint* in terms of precision, being able to detect all the true positives contained
613 in $\mathbb{H}\mathbb{F}$, with a ratio of false positives less than 40%. This was expected since, as we have
614 already discussed in Section 4 and unlike *Non-interference*, *Taint* is only able to track explicit
615 information flows. In fact, the 7 false negatives (column **#FN** of Table 3) produced by *Taint*
616 correspond to implicit non-deterministic behaviors.

¹⁹ <https://api.github.com/search/repositories?q=chaincode+fork:false+language:Go+archived:false&sort=stars&order=desc>. Accessed: 17-10-2022.

²⁰ See <https://pkg.go.dev/github.com/hyperledger/fabric-chaincode-go/shim>.

²¹ We decided not to implement those standard features since this would have required a relevant effort to support only a few more chaincodes.

Tools	# W	# TP	# FP	# FN
GoLiSA - <i>Taint</i>	173	118	55	7
GoLiSA - <i>Non-interference</i>	195	124	71	0
ChainCode Analyzer	203	68	135	53
Revive [^] CC	351	79	272	1

■ **Table 4** Warnings triggered by the analyzers on $\mathbb{H}\mathbb{F}$

5.1.1 Comparison

We compared GoLiSA with the open-source static analyzers for Go chaincode described in Section 6, namely ChainCode Analyzer and Revive[^]CC. Table 4 reports the comparison between GoLiSA and these tools over the same benchmark $\mathbb{H}\mathbb{F}$ discussed in Section 5.1.

The comparison shows that GoLiSA - *Non-interference* finds all the true issues contained in the benchmark, achieving the best and most accurate result in terms of precision with a 36.41% false positives ratio. Instead, although it has some false negatives, GoLiSA - *Taint* is the analysis with the lowest percentage of false positives with the 31.79% .

Revive[^]CC triggers 351 warnings out of which 77.49% are false positives. The only non-deterministic behaviour not detected by Revive[^]CC (last column) is due to the fact that it considers the `ioutil.ReadFile` API as safe, although reading a file should be considered non-deterministic in the blockchain context. Finally, ChainCode Analyzer is more precise w.r.t. Revive[^]CC, with 66.50% of false positives, but it has also the greatest number of false negatives, failing to detect a huge number of critical non-deterministic behaviors. This can be attributed to the fact that ChainCode Analyzer does not consider several APIs leading to non-determinism as critical and it fails to soundly detect iteration over maps.

Note that the amount of true positives discovered by GoLiSA analyses differs from the ones of other tools. In fact, GoLiSA is the only tool involved in our comparison that issues warnings on sinks rather than sources. This translates to fewer alarms being issued whenever values of multiple sources flow to the same sink (here, GoLiSA issues a single warning, while other tools issue one for each source), and to more alarms being raised whenever the value of a single source flows to multiple sinks (here instead, other tools issue a single warning, while GoLiSA issues one for each sink).

5.2 Qualitative Evaluation

5.2.1 Explicit Flow: the Boleto Contract

The *boleto* contract²², taken from $\mathbb{H}\mathbb{F}$, comes with a real non-determinism issue that can be found with explicit flows, and that was also detected by other tools during the comparison of Section 5.1.1. The *boleto* contract (Figure 10) seems to be a proof of concept application handling tickets in an e-commerce store, with the method `registrarBoleto` used to register a ticket.

Analyzing *boleto*, GoLiSA detects the explicit flow leading to a non-deterministic behavior with both *Taint* and *Non-interference*. Method `registrarBoleto` contains two different sources of non-determinism that directly flow into the same sink. The first source detected

²²<https://github.com/arthurmsouza/boleto/blob/master/boleto-chaincode/boleto.go>

```

1 func (s *SmartContract) registrarBoleto(APIStub shim.ChaincodeStubInterface, args []
  string) sc.Response {
2 // [...]
3 objBoleto.CodigoBarra = strconv.Itoa((rand.Intn(5) + 10000000 + // [...]
4 var notExpiredDate = time.Now()
5 objBoleto.DataVencimento = notExpiredDate.Format("02/01/2006")
6 // [...]
7 boletoAsBytes, _ := json.Marshal(objBoleto)
8 APIStub.PutState(args[0], boletoAsBytes)
9 // [...]
10 }

```

■ **Figure 10** Method `registrarBoleto` of *boleto* contract

650 by GoLiSA is the usage of the *Random API* to generate a barcode at line 3. Instead, the
651 second source is the usage of the *OS API* that retrieves the local machine’s time to set a
652 date at line 4. As values from both sources are used to update fields of `objBoleto`, the latter
653 is marked as tainted by the analysis, resulting in `boletoAsBytes` being tainted as well. As
654 reported in Table 2, `PutState`’s parameters are considered as sinks by GoLiSA’s analyses.
655 According to the official documentation of HF²³, the `PutState` method does not affect the
656 ledger until the transaction is validated and successfully committed. However, a transaction
657 needs to produce the same results among different peers to be validated. Hence, as passing
658 non-deterministic values to `PutState` will cause the transaction to fail, GoLiSA raises a
659 warning on line 8.

660 5.2.2 Implicit Flow: Cosmos SDK v.43

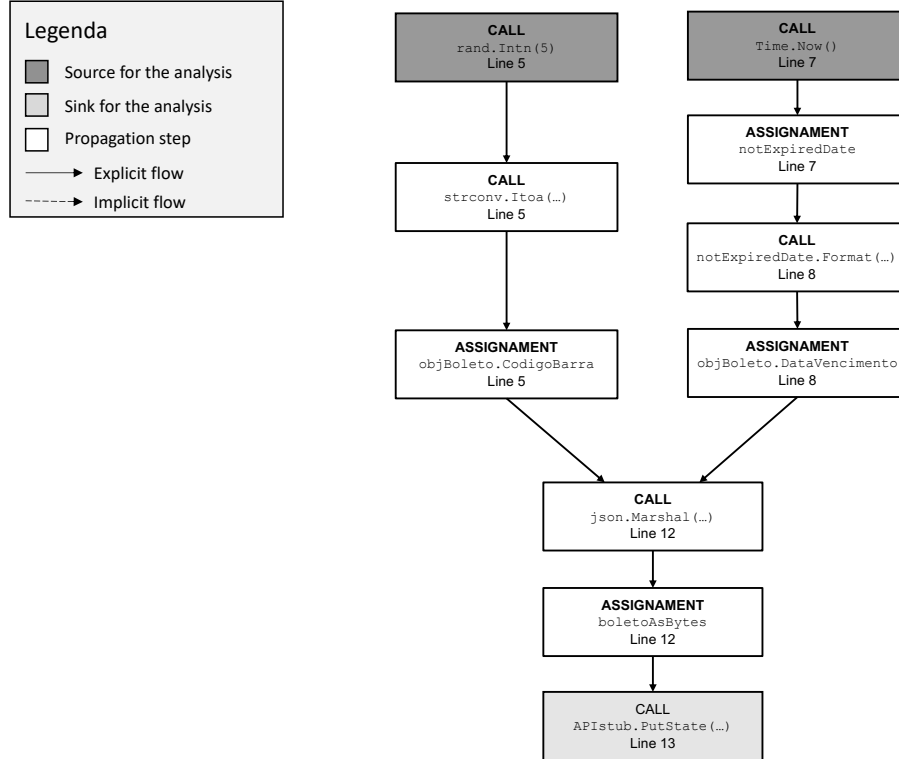
661 Analyzing the code in Figure 1, GoLiSA is able to detect an implicit flow that leads
662 to a non-deterministic behavior, that can only be detected using *Non-interference*. The
663 `ValidateBasic` method of Cosmos SDK v. 0.43.x and v. 0.44.{0,1} was designed to validate
664 a grant to ensure it has not yet expired. In this case, the source detected by GoLiSA is
665 the *OS API* used to retrieve the local machine time involved in the expiration check of
666 the grant time at line 2 of Figure 1. By propagating the information, GoLiSA detects that
667 the expiration check governs the execution of return statement. Since the `Wrap` method is
668 annotated as a sink, GoLiSA triggers an alarm at line 3 of Figure 1 as the sink is contained
669 in a block whose guard depends on non-deterministic values.

670 5.3 Limits

671 Unlike some frameworks and GPLs used in other blockchains, frameworks targeted by this
672 paper are used to develop *permissioned*, and often *private*, blockchains, meaning that the
673 related software is not publicly available or released with open-source licenses. This is also
674 the reason why the benchmark `HF` crawled from GitHub consists of 651 chaincodes, a number
675 that is not comparable with smart contract benchmarks obtained investigating other (public
676 and permissioned) blockchains. For instance, [44] collects 3075 distinct smart contracts from
677 the Ethereum blockchain, resulting in a wider benchmark.

678 The proposed solution for detecting non-deterministic behaviors is fully static. It is well
679 known that static analysis is intrinsically conservative and may produce false positives. Even

²³<https://github.com/hyperledger/fabric-chaincode-go/blob/1476cf1d3206f620db7eea12312c98669d39fa22/shim/interfaces.go>.

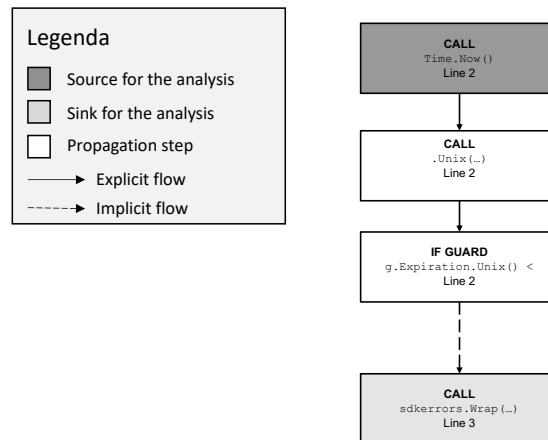


■ **Figure 11** Simplified view of explicit flow computed by GoLiSA during the analysis of registrarBoleto

680 if few have been raised by GoLiSA on the selected benchmark, one should expect more false
681 positives when applying our approach to arbitrary DApps.

682 **6 Related Work**

683 The non-determinism of smart contracts written in GPLs is a well-known issue [32, 45].
684 Frameworks such as Takamaka [41, 42] enforce determinism adopting a conservative approach
685 that limits the set of instructions and APIs of the target language, avoiding unsafe statements
686 that might lead to non-deterministic behaviors through white-listing fully deterministic APIs.
687 This approach ensures safe development while preventing that API extensions coming with
688 new language versions can bypass the check. However, it also severely limits the exploitable
689 features of the GPL. On the other hand, black-listing undesired APIs is a much harder
690 approach to maintain, but it seems the most widespread technique in Go analyzers. For
691 instance, ChainCode Analyzer [31] and Revive^{CC} [40] detect mainly black-listed imports
692 related non-deterministic APIs using a syntactical approach. Besides, they can detect non-
693 deterministic map iterations by AST traversal with minimal syntactic reasoning. Signature of
694 invoked functions can also be black-listed instead of imports [32]. These tools and frameworks
695 inherently limit API usage, sensibly reducing the benefits of adopting a GPL even when the



■ **Figure 12** Simplified view of implicit flow computed by GoLiSA during the analysis of Figure 1

696 code poses no harm to the blockchain. The problem of detecting non-determinism has also
 697 been covered for parallel applications, suggesting that non-determinism is “*most often the*
 698 *result of a mistake on the part of the programmer*” [13].

699 **7 Conclusion**

700 In this paper, we proposed a flow-based approach for detecting critical non-deterministic
 701 behaviors, namely the ones affecting the blockchain state. Our proposal has been implemented
 702 in GoLiSA, a static analyzer for Go applications. To the best of our knowledge, GoLiSA is the
 703 first semantic-based static analyzer for blockchain software able to detect non-deterministic
 704 behaviors, with an extremely low false alarm prevision. In the context of smart contracts, the
 705 proposed approach is placed in an off-chain architecture, i.e., the analysis is done before smart
 706 contracts are deployed in the blockchain, and it is not mandatory. As future work, besides
 707 supporting the missing Go features discussed in Section 5 to enhance the analysis coverage,
 708 we plan to test GoLiSA in an on-chain architecture [35], making the non-determinism checker
 709 part of the consensus protocol, with the goal of keeping the code stored within the blockchain
 710 deterministic. The analysis could be enriched with a context-sensitive flow reconstructor,
 711 such as BackFlow [21], that starting from the results of a information flow engine, reconstructs
 712 how the information flows inside the program and builds paths connecting sources to sinks.
 713 Moreover, we have focused on the non-determinism problem, but our future research will
 714 address the problem of detecting other and equally critical vulnerabilities that can affect
 715 blockchain software written using general-purpose languages, such as numerical overflow.

716 Our proposal follows a fully static approach, justified by the fact that we aim at proving
 717 the determinism of blockchain software, regardless of the possible executions. However,
 718 even if the evaluation on the selected benchmark shows optimal results, the risk of getting
 719 false alarms analyzing other applications is still present, being our approach based on
 720 over-approximating possible executions via abstract interpretation. In future works, hybrid
 721 approaches between static and dynamic analyses will be investigated to get the benefits of
 722 both techniques.

723 Finally, in order to assess the effectiveness of our proposal, we have conducted our

724 evaluation on Hyperledger Fabric blockchain software, mostly because it is the most popular
725 framework among those cited in the paper. To give a larger coverage to GoLiSA of the
726 blockchain software that can analyze, the next step will be to design significant benchmarks
727 also for the other frameworks, such as Tendermint core and Cosmos SDK, on which we can
728 experiment our static analyzer.

729 — References —

- 730 1 Lars Ole Andersen. Program Analysis and Specialization for the C Programming Language,
731 1994. Accessed: 01-12-2022. URL: [https://www.cs.cornell.edu/courses/cs711/2005fa/
732 papers/andersen-thesis94.pdf](https://www.cs.cornell.edu/courses/cs711/2005fa/papers/andersen-thesis94.pdf).
- 733 2 Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis,
734 Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich,
735 Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith
736 Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolic, Sharon Weed Cocco,
737 and Jason Yellick. Hyperledger Fabric: A Distributed Operating System for Permissioned
738 Blockchains. In *Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, Porto,
739 Portugal, April 23-26, 2018*, pages 30:1–30:15. ACM, 2018. doi:10.1145/3190508.3190538.
- 740 3 A. M. Antonopoulos. *Mastering Bitcoin: Programming the Open Blockchain*. O'Reilly, 2nd
741 edition, 2017.
- 742 4 A. M. Antonopoulos and G. Wood. *Mastering Ethereum: Building Smart Contracts and Dapps*.
743 O'Reilly, 2018.
- 744 5 Sotirios Brotsis, Nicholas Kolokotronis, Konstantinos Limniotis, Gueltoum Bendiab, and
745 Stavros Shiaeles. On the security and privacy of hyperledger fabric: Challenges and open
746 issues. In *2020 IEEE World Congress on Services (SERVICES)*, pages 197–204, 2020. doi:
747 10.1109/SERVICES48979.2020.00049.
- 748 6 E. Buchman. *Tendermint: Byzantine Fault Tolerance in the Age of Blockchains*. PhD thesis,
749 2016.
- 750 7 Ethan Buchman. Byzantine Fault Tolerant State Machine Replication in Any Programming
751 Language. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*,
752 PODC '19, page 546, New York, NY, USA, 2019. Association for Computing Machinery.
- 753 8 V. Buterin. Ethereum Whitepaper, 2013. Available at [https://ethereum.org/en/
754 whitepaper/](https://ethereum.org/en/whitepaper/).
- 755 9 Krishnendu Chatterjee, Amir Kafshdar Goharshady, and Arash Pourdamghani. Probabilistic
756 Smart Contracts: Secure Randomness on the Blockchain. In *IEEE International Conference
757 on Blockchain and Cryptocurrency, ICBC 2019, Seoul, Korea (South), May 14-17, 2019*, pages
758 403–412. IEEE, 2019. doi:10.1109/BL0C.2019.8751326.
- 759 10 Patrick Cousot. *Principles of Abstract Interpretation*. MIT Press, 2021.
- 760 11 Dorothy E. Denning. A Lattice Model of Secure Information Flow. *Commun. ACM*, 19(5):236–
761 243, 1976. doi:10.1145/360051.360056.
- 762 12 ebuchman. Cosmos-SDK Vulnerability Retrospective: Security Advisory Jackfruit,
763 October 12, 2021, 2021. Accessed: 01-12-2022. URL: [https://forum.cosmos.network/t/
764 cosmos-sdk-vulnerability-retrospective-security-advisory-jackfruit-october-12-2021/
765 5349](https://forum.cosmos.network/t/cosmos-sdk-vulnerability-retrospective-security-advisory-jackfruit-october-12-2021/5349).
- 766 13 Perry A. Emrath and David A. Padua. Automatic Detection of Nondeterminacy in Parallel
767 Programs. In *Proceedings of the 1988 ACM SIGPLAN and SIGOPS Workshop on Parallel
768 and Distributed Debugging*, PADD '88, page 89–99, New York, NY, USA, 1988. Association
769 for Computing Machinery. doi:10.1145/68210.69224.
- 770 14 Michael D. Ernst, Alberto Lovato, Damiano Macedonio, Ciprian Spiridon, and Fausto Spoto.
771 Boolean Formulas for the Static Identification of Injection Attacks in java. In *Logic for*

- 772 *Programming, Artificial Intelligence, and Reasoning - 20th International Conference, LPAR-*
773 *20 2015, Suva, Fiji, November 24-28, 2015, Proceedings*, volume 9450 of *Lecture Notes in*
774 *Computer Science*, pages 130–145. Springer, 2015. doi:10.1007/978-3-662-48899-7\10.
- 775 15 Pietro Ferrara. A generic framework for heap and value analyses of object-oriented programming
776 languages. *Theor. Comput. Sci.*, 631:43–72, 2016. doi:10.1016/j.tcs.2016.04.001.
- 777 16 Pietro Ferrara, Elisa Burato, and Fausto Spoto. Security Analysis of the OWASP Benchmark
778 with Julia. In *Proceedings of the First Italian Conference on Cybersecurity (ITASEC17), Venice,*
779 *Italy, January 17-20, 2017*, volume 1816 of *CEUR Workshop Proceedings*, pages 242–247.
780 CEUR-WS.org, 2017. <http://ceur-ws.org/Vol-1816/paper-24.pdf> Accessed: 01-12-2022.
- 781 17 Pietro Ferrara, Amit Kr Mandal, Agostino Cortesi, and Fausto Spoto. Static analysis for
782 discovering IoT vulnerabilities. *Int. J. Softw. Tools Technol. Transf.*, 23(1):71–88, 2021.
783 doi:10.1007/s10009-020-00592-x.
- 784 18 Pietro Ferrara and Luca Negrini. SARL: Oo framework specification for static analysis. In
785 Maria Christakis, Nadia Polikarpova, Parasara Sridhar Duggirala, and Peter Schrammel,
786 editors, *Software Verification*, pages 3–20, Cham, 2020. Springer International Publishing.
- 787 19 Pietro Ferrara, Luca Negrini, Vincenzo Arceri, and Agostino Cortesi. Static analysis for
788 dummies: experiencing lisa. In Lisa Nguyen Quang Do and Caterina Urban, editors,
789 *SOAP@PLDI 2021: Proceedings of the 10th ACM SIGPLAN International Workshop on*
790 *the State Of the Art in Program Analysis, Virtual Event, Canada, 22 June, 2021*, pages 1–6.
791 ACM, 2021. doi:10.1145/3460946.3464316.
- 792 20 Pietro Ferrara, Luca Olivieri, and Fausto Spoto. Tailoring Taint Analysis to GDPR. In *Privacy*
793 *Technologies and Policy - 6th Annual Privacy Forum, APF 2018, Barcelona, Spain, June*
794 *13-14, 2018, Revised Selected Papers*, volume 11079 of *Lecture Notes in Computer Science*,
795 pages 63–76. Springer, 2018. doi:10.1007/978-3-030-02547-2\4.
- 796 21 Pietro Ferrara, Luca Olivieri, and Fausto Spoto. Backflow: Backward context-sensitive
797 flow reconstruction of taint analysis results. In *Verification, Model Checking, and Abstract*
798 *Interpretation*, pages 23–43, Cham, 2020. Springer International Publishing.
- 799 22 Pietro Ferrara, Luca Olivieri, and Fausto Spoto. Static Privacy Analysis by Flow Reconstruction
800 of Tainted data. *Int. J. Softw. Eng. Knowl. Eng.*, 31(7):973–1016, 2021. doi:10.1142/
801 S0218194021500303.
- 802 23 Luca Foschini, Andrea Gavagna, Giuseppe Martuscelli, and Rebecca Montanari. Hyperledger
803 Fabric Blockchain: Chaincode Performance Analysis. In *2020 IEEE International Conference*
804 *on Communications, ICC 2020, Dublin, Ireland, June 7-11, 2020*, pages 1–6. IEEE, 2020.
805 doi:10.1109/ICC40277.2020.9149080.
- 806 24 Joseph A. Goguen and José Meseguer. Security Policies and Security Models. In *1982 IEEE*
807 *Symposium on Security and Privacy, Oakland, CA, USA, April 26-28, 1982*, pages 11–20.
808 IEEE Computer Society, 1982. doi:10.1109/SP.1982.10014.
- 809 25 Joseph A. Goguen and José Meseguer. Unwinding and Inference Control. In *Proceedings of*
810 *the 1984 IEEE Symposium on Security and Privacy, Oakland, California, USA, April 29 -*
811 *May 2, 1984*, pages 75–87. IEEE Computer Society, 1984. doi:10.1109/SP.1984.10019.
- 812 26 Hyperledger. Hyperledger fabric documentation. [https://hyperledger-fabric.readthedocs.](https://hyperledger-fabric.readthedocs.io/en/release-2.2/blockchain.html#what-is-hyperledger-fabric)
813 [io/en/release-2.2/blockchain.html#what-is-hyperledger-fabric](https://hyperledger-fabric.readthedocs.io/en/release-2.2/blockchain.html#what-is-hyperledger-fabric).
- 814 27 Tendermint Inc. What is Tendermint: A Note on Determinism, 2022.
815 Accessed: 01-12-2022. URL: [https://github.com/tendermint/tendermint/blob/](https://github.com/tendermint/tendermint/blob/7983f9cc36c31e140e46ae5cb00ed39f637ef283/docs/introduction/what-is-tendermint.md#a-note-on-determinism)
816 [7983f9cc36c31e140e46ae5cb00ed39f637ef283/docs/introduction/what-is-tendermint.](https://github.com/tendermint/tendermint/blob/7983f9cc36c31e140e46ae5cb00ed39f637ef283/docs/introduction/what-is-tendermint.md#a-note-on-determinism)
817 [md#a-note-on-determinism](https://github.com/tendermint/tendermint/blob/7983f9cc36c31e140e46ae5cb00ed39f637ef283/docs/introduction/what-is-tendermint.md#a-note-on-determinism).
- 818 28 Uday P. Khedker and Bageshri Karkare. Efficiency, precision, simplicity, and generality
819 in interprocedural data flow analysis: Resurrecting the classical call strings method. In
820 *Compiler Construction*, pages 213–228, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
821 doi:10.1007/978-3-540-78791-4\15.
- 822 29 J. Kwon. Tendermint: Consensus without mining. 2014.
- 823 30 J. Kwon and E. Buchman. Cosmos whitepaper, 2019.

- 824 31 kzhry. Chaincode Analyzer, 2021. Accessed: 01-12-2022. URL: <https://github.com/hyperledger-labs/chaincode-analyzer>.
825
- 826 32 Penghui Lv, Yu Wang, Yazhe Wang, and Qihui Zhou. Potential Risk Detection System
827 of Hyperledger Fabric Smart Contract based on Static Analysis. In *IEEE Symposium on*
828 *Computers and Communications, ISCC 2021, Athens, Greece, September 5-8, 2021*, pages 1–7.
829 IEEE, 2021. doi:10.1109/ISCC53001.2021.9631249.
- 830 33 S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. Available at <https://bitcoin.org/bitcoin.pdf>, 2008.
831
- 832 34 Luca Negrini. *A generic framework for multilanguage analysis*. PhD thesis, Università Ca’
833 Foscari Venezia, 2023.
- 834 35 Luca Olivieri, Fausto Spoto, and Fabio Tagliaferro. On-Chain Smart Contract Verification over
835 Tendermint. In *Financial Cryptography and Data Security. FC 2021 International Workshops*
836 *- CoDecFin, DeFi, VOTING, and WTSC, Virtual Event, March 5, 2021, Revised Selected*
837 *Papers*, volume 12676 of *Lecture Notes in Computer Science*, pages 333–347. Springer, 2021.
838 doi:10.1007/978-3-662-63958-0_28.
- 839 36 Luca Olivieri, Fabio Tagliaferro, Vincenzo Arceri, Marco Ruaro, Luca Negrini, Agostino Cortesi,
840 Pietro Ferrara, Fausto Spoto, and Enrico Talin. Ensuring determinism in blockchain software
841 with golisa: an industrial experience report. In Laure Gonnord and Laura Titolo, editors, *SOAP*
842 *’22: 11th ACM SIGPLAN International Workshop on the State Of the Art in Program Analysis,*
843 *San Diego, CA, USA, 14 June 2022*, pages 23–29. ACM, 2022. doi:10.1145/3520313.3534658.
- 844 37 Xavier Rival and Kwangkeun Yi. *Introduction to static analysis: an abstract interpretation*
845 *perspective*. Mit Press, 2020.
- 846 38 A. Sabelfeld and A.C. Myers. Language-based information-flow security. *IEEE Journal on*
847 *Selected Areas in Communications*, 21(1):5–19, 2003. doi:10.1109/JSAC.2002.806121.
- 848 39 Micha Sharir, Amir Pnueli, et al. *Two approaches to interprocedural data flow analysis*. New
849 York University. Courant Institute of Mathematical Sciences . . . , 1978.
- 850 40 sivachokkapu. Revivecc, 2021. Accessed: 01-12-2022. URL: <https://github.com/sivachokkapu/revive-cc>.
851
- 852 41 Fausto Spoto. A Java Framework for Smart Contracts. In *Financial Cryptography and Data*
853 *Security - FC 2019 International Workshops, VOTING and WTSC, St. Kitts, St. Kitts and*
854 *Nevis, February 18-22, 2019, Revised Selected Papers*, volume 11599 of *Lecture Notes in*
855 *Computer Science*, pages 122–137. Springer, 2019. doi:10.1007/978-3-030-43725-1_10.
- 856 42 Fausto Spoto. Enforcing Determinism of Java Smart Contracts. In *Financial Cryptography and*
857 *Data Security - FC 2020 International Workshops, AsiaUSEC, CoDeFi, VOTING, and WTSC,*
858 *Kota Kinabalu, Malaysia, February 14, 2020, Revised Selected Papers*, volume 12063 of *Lecture*
859 *Notes in Computer Science*, pages 568–583. Springer, 2020. doi:10.1007/978-3-030-54455-3\
860 _40.
- 861 43 Omer Tripp, Marco Pistoia, Stephen J. Fink, Manu Sridharan, and Omri Weisman. TAJ:
862 effective taint analysis of web applications. In *Proceedings of the 2009 ACM SIGPLAN*
863 *Conference on Programming Language Design and Implementation, PLDI 2009, Dublin,*
864 *Ireland, June 15-21, 2009*, pages 87–97. ACM, 2009. doi:10.1145/1542476.1542486.
- 865 44 Shuai Wang, Chengyu Zhang, and Zhendong Su. Detecting nondeterministic payment bugs
866 in ethereum smart contracts. *Proc. ACM Program. Lang.*, 3(OOPSLA):189:1–189:29, 2019.
867 doi:10.1145/3360615.
- 868 45 Kazuhiro Yamashita, Yoshihide Nomura, Ence Zhou, Bingfeng Pi, and Sun Jun. Potential
869 Risks of Hyperledger Fabric Smart Contracts. In *2019 IEEE International Workshop on*
870 *Blockchain Oriented Software Engineering (IWBOSE)*, pages 1–10, 2019. doi:10.1109/IWBOSE.
871 2019.8666486.