

Towards FATEful Smart Contracts^{*}

Luigi Bellomarini¹, Marco Favorito¹, Eleonora Laurenza¹, Markus Nissl² and Emanuel Sallinger^{2,3}

¹Bank of Italy, Italy

²TU Wien, Austria

³University of Oxford, UK

Abstract

Achieving fair, accountable, transparent, and ethical decentralized finance requires activating enabling properties at the level of smart contracts, the executable scripts at its basis. In this vision paper, a joint effort of the Central Bank of Italy, TU Wien, and the University of Oxford, we leverage the vast amount of experience in this sense from the database community and propose a logic-based reasoning framework that captures smart contracts as a set of rules in DatalogMTL, a temporal language for querying databases. We present the high-level architecture of our framework and explain how the theoretical underpinnings of the reasoning of DatalogMTL convey important properties to the approach.

Keywords

Smart contracts, DatalogMTL, FATE principles

1. Introduction

The Artificial Intelligence and database communities are experiencing a growing infusion of the *FATE* (*Fairness, Accountability, Transparency, Ethics*) principles [2]. These principles are gaining prominence, drawing attention to the non-functional requirements of everyday AI-assisted and data-driven decision-making and catalyzing the discussion around regulatory bodies. Unfortunately, the same level of attention to these high-level concerns is not mirrored in developer circles, and recent studies underscore the scant regard machine learning developers have shown for FATE concerns in machine learning applications [3, 4].

FATE and DeFi. We see similar patterns emerging when assessing developers' awareness of FATE concerns within the industrial realm of *Decentralized Finance* (DeFi). DeFi entails financial transactions devoid of intermediaries, instead relying on software modules executed on a decentralized public ledger [5]. At the core of DeFi is the notion of *smart contracts* [6], which are machine-readable and executable agreements that establish and enforce the binding terms for the parties involved.

Supporting FATE. In the AI and data world, social forces have been effective in supporting FATE, for example, by means of third-party audits of the algorithms, either conducted by experts or by everyday users. As recently highlighted by Hong in CACM [2], prominent examples can be found in the fights against the racial bias of face-recognition systems [7], and commercial gender disparities in photo cropping or credit card algorithms [8]. These audits, spurred by social forces and the scientific community, have provided regulators with positive guidance.

In contrast to AI, DeFi boasts a substantial theoretical transparency advantage, thanks to its open-access code and the ability for anyone to inspect smart contract data on a public ledger. However, enforcing policies in a decentralized context remains an incredibly challenging task. The praiseworthy goal of establishing standards, taxonomies, compliance measures, quality controls, and upholding

AMW 2024: 16th Alberto Mendelzon International Workshop on Foundations of Data Management, September 30th–October 4th, 2024, Mexico City, Mexico

^{*}This is a short version of the paper [1] already presented at the 6th Distributed Ledger Technologies Workshop.

✉ luigi.bellomarini@bancaditalia.it (L. Bellomarini); marco.favorito@bancaditalia.it (M. Favorito); eleonora.laurenza@bancaditalia.it (E. Laurenza); nissl@dbai.tuwien.ac.at (M. Nissl); sallinger@dbai.tuwien.ac.at (E. Sallinger)

🆔 0000-0001-6863-0162 (L. Bellomarini); 0000-0001-9566-3576 (M. Favorito); 0000-0002-2786-8163 (E. Laurenza); 0000-0001-8196-5688 (M. Nissl); 0000-0001-7441-129X (E. Sallinger)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

The views and opinions expressed in this paper are those of the authors and do not necessarily reflect the official policy or position of Banca d'Italia.

ethical principles [9] can benefit from robust support from the social forces in the monitoring and enforcement of such policies. However, this must align with the specific technical attributes of these smart contracts. Yet, smart contracts have been criticised within the community due to their overly complex business logic, and limited explainability, resulting in a lack of transparency. Furthermore, they often prove challenging to describe and communicate, rendering them less user-friendly, particularly for non-technical users [10, 11, 12].

A Knowledge Representation and Reasoning (KRR) approach to smart contracts. In the area of deductive AI and ontological reasoning on databases, logic-based approaches built on top of KRR formalisms are gaining increasing attention in industrial settings, with many successful financial applications [13, 14, 15]. Modern logical languages manage to strike a good balance between expressive power and computational complexity, resulting in compact and efficiently executable formalizations of complex domains, for instance, being able to capture SPARQL under OWL 2 entailment regimes [16] and so enabling ontological reasoning. The *declarative* paradigm sustains simplicity, transparency, compactness, and understandability of code, which becomes algorithm-independent and closer to the high-level specifications, policies, and standards. The *well-defined semantics* of KRR languages fosters non-ambiguity, ease of use for non-technical users, and correctness. The intrinsic *step-by-step nature of logical reasoning* is conceptually close to notions of *explainability* and thus supports decision transparency.

The thesis of this vision paper is that a KRR framework for smart contracts that addresses FATE by design is both theoretically and practically viable. For the theory, we show that, by building on the underpinnings of logic-based reasoning, the features important to achieving FATE desiderata can be obtained and rigorously justified. In terms of application, we show that our framework is adaptable to serve as both an interpreted and a compiled execution mode for real-world contracts.

Contributions to industrial advances. In recent industrial EDBT work done by the Central Bank of Italy [17], they started to investigate the possibility of encoding complex smart contracts in DatalogMTL [18, 19], a temporal extension of the Datalog language [20] of databases. They obtained promising results, highlighting the potential of a KRR approach in the specific case of a derivative contract. In this work, a joint effort of the Central Bank of Italy, TU Wien, and the University of Oxford, we propose (i) a **full-fledged and general framework for smart contracts** that sustains FATE concerns (ii) by leveraging the vast amount of **experience from the database community to achieve enabling properties**; (iii) using our framework to study and implement **proof of concepts for many smart contracts** where FATE is a core desideratum of a central bank. More details can be found in the full version of the paper [1].

2. Overview of the Framework and Related Work

We use a form of *declarative logical object-oriented approach* and encode the behaviour of a *class of smart contracts* as a set Σ of reasoning rules—or programs—working on a database D of temporal facts. A temporal fact of D is such that it holds in a given time interval, for example, $price(123, 2)@[2023-09-01, 2023-09-02]$ defines the price 2 for the asset 123 in a two-day interval.

To model the rules of Σ , we introduce $DatalogMTL^S$, a variant of DatalogMTL with features of practical utility. A smart contract is then an *instance of a smart contract class*, whose time-dependent status is represented as a database D of temporal facts. Instances are *stateful* objects and the contract execution consists in invocations, akin to method calls, that are carried out by the involved parties. Calls result in updates to the status D through the addition of new facts. The semantics of a call is operationally described as the application of the rules of Σ (denoted as $\Sigma(D)$) to the temporal facts of D , extended with call-specific facts.

DatalogMTL^S rules are sets of *head* \leftarrow *body* logic implications where the body is a conjunction of atoms and the head is an atom. As a general guideline, whenever the body of a rule is satisfied by a conjunction of facts in D at a point in time t , the evaluation of the rule triggers the insertion in D of a new fact for the head atom, holding at t . For example, the rule $'position(x, w) \leftarrow buy(x, a, q), price(a, p), w = p * q'$

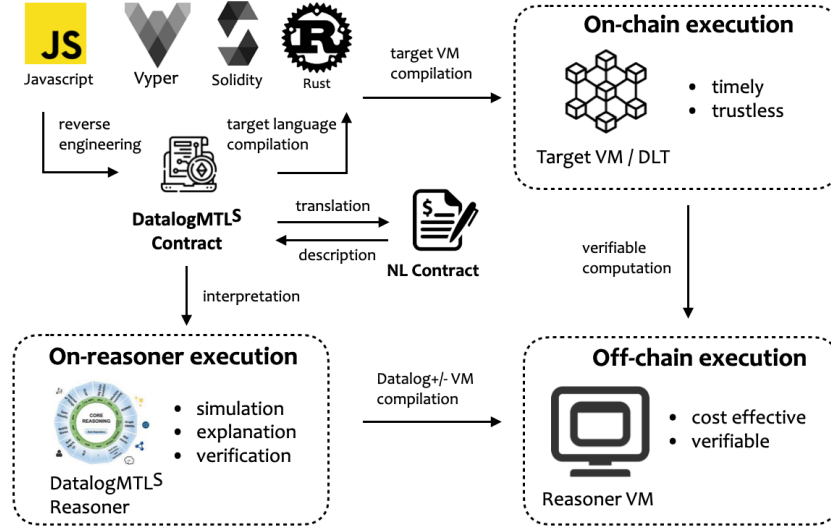


Figure 1: Overview of the DatalogMTL^S framework for smart contracts.

states that, for every point in time t , the position w of a trader x that buys an amount q of an asset a of price p is obtained as $p * q$. So, if D contains the price fact $price(123, 2)$ and $buy(0x241F, 123, 12)$, both holding at $[2023-09-01, 2023-09-02]$, a new $position(0x241F, 24)$ holding in the same time interval will be added to D .

In DatalogMTL^S, temporal operators can be used to either modify the temporal binding of body atoms to facts of D , or to alter the temporal validity of the generated facts. For instance, an expression of the form $\Diamond_{[0,1]} position(x, w)$ in the body holds at a point in time t , if the trader x had at least an open position w in the interval $[t - 1, t]$, while an expression of the form $\Box_{[0,1]} position(x, w)$ in the head, states that the position will be open in the interval $[t, t + 1]$ for every t . In the following, we show a smart contract class defining a simple financial market:

- R1: $accepted(\$sender, y) \leftarrow \#open(y), \neg marketClosed.$
R2: $\Box position(x, y, k) \leftarrow accepted(x, y), price(p), k = y * p.$
R3: $return(x, g), \Box position(x, y, 0) \leftarrow \#close(), price(p), position(\$sender, y, k), g = y * p - k.$

At a specific point in time in which the market is not closed, a trader tries to open a position by investing an amount y (Rule R1). The constant $\$sender$ is a call-level variable that at runtime binds to the invoking trader. When the transaction is accepted (Rule R2), the position of the trader x on the amount y is updated by multiplying by the current price p . The \Box operator stands for a new future temporal validity of the position fact. Finally, when the position is closed (Rule R3), the final profit g is computed based on the current price.

Execution modes. From a practical perspective, our framework implements $\Sigma(D)$ by enabling three execution alternatives, each offering specific properties, as reported in Figure 1: (i) *on-reasoner execution*: the rules are applied natively by a reasoning system supporting DatalogMTL or DatalogMTL^S such as *Temporal Vatalog* [21] or *MeTeoR* [22]; (ii) *on-chain execution*: the DatalogMTL^S programs are verifiably translated into the language of a target system, for instance, *Solidity* [23] or *Bitcoin Script* [24] and executed within the target systems; (iii) *off-chain execution*, the rules are applied with an on-reasoner execution with persistent effects on a blockchain and cryptographically verifiable computation.

The expert intention is substantiated as a natural language (NL) contract or directly encoded in a DatalogMTL^S program Σ , and the use of large language models (LLM) can bridge the gap between the natural language specification of the contract and the encoding of its DatalogMTL^S version [25].

In the on-reasoner execution mode, to evaluate Σ , reasoners use variants of the CHASE procedure [26] specialized for the temporal extensions [27]. They offer a good degree of *explainability* of the produced facts as a side effect of the inference process of the chase. Also, they are suited to be used for *simulation*

and *runtime verification* purposes as step-by-step debugging can be emulated by incrementally adding facts to D and monitoring the results entailed by the application of Σ . On the other hand, the execution relies on a trusted centralized system.

Conversely, on-chain execution offers a *trustless* paradigm, only requiring that the user acknowledges the translation of the DatalogMTL^S code into the target language, which can be verified through open-sourced translators. This trustlessness is underpinned by the validation properties ensured by the distributed consensus protocol embraced by the blockchain, guaranteeing the integrity of the mined blocks. What is more, on-chain execution is characterized by its *timeliness*, as results are promptly included in the first mined block.

However, on-chain execution comes with a high cost and is ill-suited for complex applications. In contrast, off-chain execution is widely regarded as a practical and efficient alternative, and there is a large body of related work such as state-channels [28], Plasma [29], and Zero-Knowledge Rollups [30]. In particular, specific protocols have been proposed, that help attest the integrity of off-chain computation (i.e., *verifiable computation*), such as ZK-SNARK [31] and ZK-STARK [32]. Towards this direction, the construction of specialized virtual machines compiling succinct ZK proofs for DatalogMTL^S executions is envisaged here, but beyond the scope of this vision paper and a matter of future work.

Termination and Complexity. Fact entailment in DatalogMTL is a decidable task, in particular PSPACE in data complexity [18]; therefore we have *guaranteed termination* and *guaranteed computational complexity*. Moreover, the rules modelling real-world smart contracts need to allow for the derivation of facts into present and future time points, while the propagation towards the past is almost never required. Under this condition, the set Σ belongs to the *forward-propagating fragment*, namely, DatalogMTL^{FP} [18], for which a *finite representation* of infinite models is always possible [19]. It is important to point out that in DatalogMTL the use of arithmetic and recursion can, in general, lead to undecidability [33] and a comprehensive study of arithmetic in DatalogMTL has not been provided yet. However, our framework conditions the activations of the rules on the specific smart contract functions being called, which reduces the cases of potentially harmful recursion.

3. The Framework in Action

In this section, we show the usefulness of our framework with a smart contract of industrial relevance.

ERC-20. The ERC-20 [34] is a well-known and widely adopted Token Standard that implements an API for tokens within smart contracts. The following DatalogMTL^S smart contract implements a simple ERC-20 contract with a fixed supply S .

$$\begin{aligned}
R1: & \text{totalSupply}(S), \overline{\boxplus} \text{balanceOf}(\$sender, S) \leftarrow \#init(S). \\
R2: & \overline{\boxplus} \text{balanceOf}(\$sender, 0) \leftarrow \neg \text{balanceOf}(\$sender, X), \#create(). \\
R3: & \overline{\boxplus} \text{balanceOf}(\$sender, B_s - A), \\
& \overline{\boxplus} \text{balanceOf}(to, B_r + A) \leftarrow \text{balanceOf}(\$sender, B_s), \text{balanceOf}(to, B_r), \\
& B_s \geq A, \#transfer(to, A).
\end{aligned}$$

Atoms with predicates $\#init$, $create$, and $\#transfer$ are called *trigger atoms* and represent contract functions, while atoms with $balanceOf$ and $totalSupply$ are *status atoms* and they persist across state transitions. Other atoms are *transient* and they are discarded at the end of each function evaluation. Rule $R1$ initializes the smart contract state by adding the facts $\boxplus \text{totalSupply}(S)$ and $\overline{\boxplus} \text{balanceOf}(\$sender, S)$. Note that using the $\overline{\boxplus}$ operator allows to overwrite the balance in case of a transfer. Rule $R2$ allows the sender to initialize a balance, if not already done earlier. Rule $R3$ implements the “transfer” function from the sender address $\$sender$ to the recipient address to of amount A . Atoms of the form $balanceOf(address, X)$ in the body are used to query the balance X of $address$ (a common pattern in logic programming), the condition $B_s \geq A$ imposes that there is enough balance from the sender account to complete the transfer, and the head of the rule uses the $\overline{\boxplus}$ to update the balances accordingly. We omitted allowances and the *transferFrom* and *approve* functions.

<pre> 1 contract SimpleERC20Contract { 2 uint256 immutable totalSupply; 3 mapping(address => uint256) public balanceOf; 4 5 constructor(uint256 _S) { 6 totalSupply = _S; 7 balanceOf[msg.sender] = _S; 8 } 9 10 function create() public { 11 if (balanceOf[msg.sender] != 0) { revert(); } 12 balanceOf[msg.sender] = 0; 13 } </pre>	<pre> 14 15 function transfer(address _to, uint256 _A) public { 16 uint256 Bs = balanceOf[msg.sender]; 17 uint256 Br = balanceOf[_to]; 18 if (!(Bs >= _A)) { revert(); } 19 balanceOf[msg.sender] = Bs - _A; 20 balanceOf[_to] = Br + _A; 21 } 22 } </pre>
--	---

Figure 2: The Solidity code generated from the ERC-20 DatalogMTL^S program example.

Compilation to Solidity. In [1] we outlined an approach to compile DatalogMTL^S programs into Solidity code that implements the same contract. Intuitively, rules with trigger atoms in the body are translated into functions, and status atoms are translated into contract state variables and data structures. Figure 2 shows an example of how we can translate the ERC-20 DatalogMTL^S program into Solidity.

Formal Verification. Our framework can enable formal verification of smart contracts written in DatalogMTL^S. An example of how formal properties can be verified is by additional program rules that formalize *invariants* in the DatalogMTL^S language. These are formalized using rules of the form $\perp \leftarrow A_1, \dots, A_m$, meaning that, if all expressions A_1, \dots, A_m are true (that is, the invariant does not hold), then the function call that triggered that rule must be reverted. These rules are compiled into assert instructions and can be seen as a runtime verification technique that rejects all transactions that violate the invariants, as in [35]. For example, in the ERC-20 smart contract, we might add some invariant conditions that must be true at any time during the lifetime of the smart contract, such as:

- The sum of user balances is equal to *totalSupply* [36] (the operator *msum* is an aggregation operator that computes the sum, see [37, 27])

$$\begin{aligned}
 & actualTotalSupply(msum(\langle(B)\rangle)) \leftarrow balanceOf(_, B). \\
 & \perp \leftarrow actualTotalSupply(N_1), totalSupply(N_2), N_1 \neq N_2.
 \end{aligned}$$

- The sums of sender and receiver balances before and after the transfer are equal [38]:

$$\begin{aligned}
 & \perp \leftarrow \Diamond_{[1,1]} \#transfer(), address(to), \Diamond_{[1,1]} balanceOf(\$sender, B_s), \Diamond_{[1,1]} balanceOf(to, B_r), \\
 & balanceOf(\$sender, B'_s), balanceOf(\$sender, B'_r), B_s + B_r = B'_s + B'_r.
 \end{aligned}$$

4. Conclusion

This preliminary work proposes a framework for expressing and evaluating smart contracts, focusing on improving the explainability and transparency of traditional smart contract development. We achieve this by building on decades of research in KRR, particularly in declarative logic-based programming, leveraging the expressive power of DatalogMTL. Inspired by a recent application of such formalism for modelling smart contracts [39, 17], we generalize those approaches and develop a foundational framework which supports the formalization and evaluation of arbitrary smart contracts. While in this paper we focused more on the vision and its industrial applications in the financial sector, we can already foresee many research avenues as future works. First, we aim to develop novel techniques (e.g., zero-knowledge techniques for off-chain execution) and implementations for realizing each execution mode. Next, it would be interesting to investigate the impact of developing smart contracts in DatalogMTL^S, in terms of ease of use, explainability, and code quality. Finally, we want to devise and apply formal verification techniques for DatalogMTL^S smart contracts.

References

- [1] L. Bellomarini, M. Favorito, E. Laurenza, M. Nissl, E. Sallinger, Toward fateful smart contracts, in: DLT, CEUR Workshop Proceedings, CEUR-WS.org, 2024.
- [2] J. I. Hong, Teaching the fate community about privacy, *Commun. ACM* 66 (2023) 10–11.
- [3] Á. A. Cabrera, E. Fu, D. Bertucci, K. Holstein, A. Talwalkar, J. I. Hong, A. Perer, Zeno: An interactive framework for behavioral evaluation of machine learning, in: CHI, ACM, 2023, pp. 419:1–419:14.
- [4] T. Li, Y. Agarwal, J. I. Hong, Coconut: An IDE plugin for developing privacy-friendly apps, *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2 (2018) 178:1–178:35.
- [5] E. Napoletano, B. Curry, What is defi? understanding decentralized finance, <http://bitly.ws/xd7Y>, 2021. Last accessed on 2022-11-28.
- [6] I. Swaps, D. Association, Legal guidelines for smart derivatives contracts: the isda master agreement, 2019.
- [7] P. E. Naeini, J. Dheenadhayalan, Y. Agarwal, L. F. Cranor, An informative security and privacy “nutrition” label for internet of things devices, *IEEE Secur. Priv.* 20 (2022) 31–39.
- [8] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: FAT, volume 81 of *Proceedings of Machine Learning Research*, PMLR, 2018, pp. 77–91.
- [9] Bank of Italy, Public consultation on the working document relating to phase one of the research project on smart contracts, <https://bit.ly/3xQKRL8>, 2023. [Online; accessed 3-May-2024].
- [10] G. Ciatto, R. Calegari, S. Mariani, E. Denti, A. Omicini, From the blockchain to logic programming and back: Research perspectives, in: WOA, 2018, pp. 69–74.
- [11] M. Li, J. Weng, A. Yang, J. Weng, Y. Zhang, Towards interpreting smart contract against contract fraud: A practical and automatic realization, *Cryptology ePrint Archive*, Paper 2020/574, 2020. URL: <https://eprint.iacr.org/2020/574>, <https://eprint.iacr.org/2020/574>.
- [12] E. Regnath, S. Steinhorst, Smaconat: Smart contracts in natural language, in: 2018 Forum on Specification and Design Languages (FDL), 2018, pp. 5–16. doi:10.1109/FDL.2018.8524068.
- [13] A. Hogan, E. Blomqvist, M. Cochez, C. d’Amato, G. de Melo, C. Gutierrez, J. E. L. Gayo, S. Kirrane, S. Neumaier, A. Polleres, R. Navigli, A. N. Ngomo, S. M. Rashid, A. Rula, L. Schmelzeisen, J. F. Sequeda, S. Staab, A. Zimmermann, Knowledge graphs, *CoRR abs/2003.02320* (2020).
- [14] T. Baldazzi, L. Bellomarini, E. Sallinger, Reasoning over financial scenarios with the vatalog system, in: EDBT, OpenProceedings.org, 2023, pp. 782–791.
- [15] L. Bellomarini, D. Fakhoury, G. Gottlob, E. Sallinger, Knowledge graphs and enterprise AI: the promise of an enabling technology, in: ICDE, IEEE, 2019, pp. 26–37.
- [16] G. Gottlob, A. Pieris, Beyond SPARQL under OWL 2 QL entailment regime: Rules to the rescue, in: IJCAI, 2015, pp. 2999–3007.
- [17] A. Colombo, L. Bellomarini, S. Ceri, E. Laurenza, Smart derivative contracts in datalogmtl, in: EDBT, OpenProceedings.org, 2023, pp. 773–781.
- [18] P. A. Walega, B. C. Grau, M. Kaminski, E. V. Kostylev, Datalogmtl: Computational complexity and expressive power, in: IJCAI, ijcai.org, 2019, pp. 1886–1892.
- [19] L. Bellomarini, M. Nissl, E. Sallinger, Query evaluation in datalogmtl - taming infinite query results, *CoRR abs/2109.10691* (2021).
- [20] S. Ceri, G. Gottlob, L. Tanca, What you always wanted to know about datalog (and never dared to ask), *TKDE* 1 (1989) 146–166.
- [21] L. Bellomarini, L. Blasi, M. Nissl, E. Sallinger, The temporal vatalog system, in: RuleML+RR, volume 13752 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 130–145.
- [22] D. Wang, P. Hu, P. A. Walega, B. C. Grau, Meteor: Practical reasoning in datalog with metric temporal operators, in: AAI, AAAI Press, 2022, pp. 5906–5913.
- [23] T. S. Authors, Solidity documentation, <https://docs.soliditylang.org/en/v0.8.21/>, 2023. [Online; accessed 12-Oct-2023].
- [24] B. Wiki, Script, <https://en.bitcoin.it/wiki/Script>, 2023. [Online; accessed 12-Oct-2023].
- [25] T. Baldazzi, L. Bellomarini, S. Ceri, A. Colombo, A. Gentili, E. Sallinger, Fine-tuning large enterprise

- language models via ontological reasoning, in: International Joint Conference on Rules and Reasoning, Springer, 2023, pp. 86–94.
- [26] D. Maier, A. O. Mendelzon, Y. Sagiv, Testing implications of data dependencies, *ACM Transactions on Database Systems* 4 (1979) 455–468.
 - [27] L. Bellomarini, M. Nissl, E. Sallinger, Monotonic aggregation for temporal datalog, in: RuleML+RR (Supplement), volume 2956 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2021.
 - [28] A. Miller, I. Bentov, S. Bakshi, R. Kumaresan, P. McCorry, Sprites and state channels: Payment networks that go faster than lightning, in: Financial Cryptography, volume 11598 of *Lecture Notes in Computer Science*, Springer, 2019, pp. 508–526.
 - [29] M. Harishankar, D. Akestoridis, S. V. Iyer, A. Laszka, C. Joe-Wong, P. Tague, Plasma go: A scalable sidechain protocol for flexible payment mechanisms in blockchain-based marketplaces, *CoRR abs/2003.06197* (2020).
 - [30] Ethereum.org, Zero-knowledge rollups, <https://shorturl.at/wIYZ4>, 2023. Last accessed on 2022-11-28.
 - [31] T. Chen, H. Lu, T. Kunpittaya, A. Luo, A review of zk-snarks, *CoRR abs/2202.06877* (2022).
 - [32] E. Ben-Sasson, I. Bentov, Y. Horesh, M. Riabzev, Scalable, transparent, and post-quantum secure computational integrity, *IACR Cryptol. ePrint Arch.* (2018) 46.
 - [33] B. C. Grau, I. Horrocks, M. Kaminski, E. V. Kostylev, B. Motik, Limit datalog: A declarative query language for data analysis, *SIGMOD Rec.* 48 (2019) 6–17.
 - [34] Ethereum Improvement Proposals, ERC-20: Token Standard, 2015. URL: <https://eips.ethereum.org/EIPS/eip-20>.
 - [35] A. Li, J. A. Choi, F. Long, Securing smart contract with runtime validation, in: PLDI, ACM, 2020, pp. 438–453.
 - [36] Á. Hajdu, D. Jovanović, solc-verify: A modular verifier for solidity smart contracts, in: Verified Software. Theories, Tools, and Experiments: 11th International Conference, VSTTE 2019, New York City, NY, USA, July 13–14, 2019, Revised Selected Papers 11, Springer, 2020, pp. 161–179.
 - [37] A. Shkapsky, M. Yang, C. Zaniolo, Optimizing recursive queries with monotonic aggregates in deals, in: ICDE, 2015, pp. 867–878.
 - [38] L. Alt, C. Reitwießner, Smt-based verification of solidity smart contracts, in: ISoLA (4), volume 11247 of *Lecture Notes in Computer Science*, Springer, 2018, pp. 376–388.
 - [39] M. Nissl, E. Sallinger, Modelling smart contracts with datalogmtl, in: EDBT/ICDT Workshops, volume 3135 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022.