

Opportunistic Algorithmic Double-Spending: How I learned to stop worrying and hedge the Fork

Nicholas Stifter^{1,2}, Aljosha Judmayer^{1,2}, Philipp Schindler^{1,2}, and
Edgar Weippl^{1,2}

¹ University of Vienna, Vienna, Austria

² SBA Research, Vienna, Austria

(firstletterfirstname)(lastname)@sba-research.org

Abstract. In this paper, we outline a novel form of attack we refer to as Opportunistic Algorithmic Double-Spending (*OpAl*). *OpAl* attacks *avoid equivocation*, i.e., do not require conflicting transactions, and are *carried out automatically* in case of a fork. Algorithmic double-spending is facilitated through transaction semantics that dynamically depend on the context and ledger state at the time of execution. Hence, *OpAl* evades common double-spending detection mechanisms and can *opportunistically* leverage forks, even if the malicious sender themselves is not responsible for, or even actively aware of, any fork. Forkable ledger designs with expressive transaction semantics, especially stateful EVM-based smart contract platforms such as Ethereum, are particularly vulnerable. Hereby, the cost of modifying a regular transaction to opportunistically perform an *OpAl* attack is low enough to consider it a viable default strategy. While Bitcoin’s stateless UTXO model, or Cardano’s EUTXO model, appear more robust against *OpAl*, we nevertheless demonstrate scenarios where transactions are *semantically malleable* and thus vulnerable. To determine whether *OpAl*-like semantics can be observed in practice, we analyze the execution traces of 922 562 transactions on the Ethereum blockchain. Hereby, we are able to identify transactions, which may be associated with frontrunning and MEV bots, that exhibit some of the design patterns also employed as part of the herein presented attack.

1 Introduction

Double-spending attacks in cryptocurrencies are primarily considered in two general categories. In the first category, an adversary is either themselves capable, or is able to coerce others, to carry out an attack that undermines the expected security guarantees of the underlying consensus protocol [71,75,86]. Hereby, attack vectors such as information withholding [26,65] and information eclipsing [34,3], as well as exploiting the rational behavior of participants [8,53,39], have received particular attention. The second category of double-spending attacks leverages inadequately chosen security parameters by merchants, i.e. they provide goods or services while the probability of the payment transaction being reverted is non-negligible [42,75]. In this regard, the probabilistic consensus guarantees of

Nakamoto consensus [67] may be misunderstood in practice, which contributes to insecure behavior by its users [75,58]. Regardless of the attack category, it is predominantly assumed that the adversary proactively performs double-spending through *equivocation* [36], i.e., by creating mutually exclusive transactions.

We hereby challenge this status quo and discuss an alternative attack, which we refer to as *opportunistic algorithmic double-spending*, whereby the intent to double-spend is intentionally encoded as part of the transaction semantics. Algorithmic double-spending *does not require equivocating* transactions and is facilitated through distributed ledgers that exhibit two properties, namely i) the ability to define transaction semantics that dynamically depend on the ledger state or execution context, which we refer to as *semantic malleability*, and ii) *probabilistic consensus decisions*, i.e., protocols without finality, or where security failures have compromised the *safety* of consensus decisions.

If these two conditions are fulfilled, *OpAl* can be used as a *free riding gadget* to profit from any sufficiently deep blockchain fork. *OpAl* attacks do not stand in contradiction to the security guarantees and desirable properties [29,67] offered by Nakamoto-style distributed ledgers. The existence of state instability, i.e. forks, is abstracted away in *idealized ledgers* by waiting sufficiently long for the relevant actions, e.g. transactions, to be included in the common prefix with high probability [10]. However, determining the correct choice of security parameters for real-world system settings is difficult [75,37] and unforeseen technical failures, or attacks, that undermine a ledger’s security assumptions through deep forks, can happen in practice [54], (see App. 3). Especially during such extraordinary events the threat of *OpAl* attacks can prove particularly severe. Even under the assumption that the ledger’s security guarantees hold, algorithmic double-spending can be of concern in cases where users exhibit an insecure interaction model, referred to as *hasty players* [10], whereby actions are taken based on unstable state. We crucially note that such patterns are commonly encountered in real-world ledgers such as Ethereum, e.g., in the context of *decentralized finance* (DeFi), where hastiness can be financially advantageous [16,89]. Our empirical analysis of Ethereum transactions in Section 7 also reveals that *OpAl*-like semantics are being used by entities which, according to block explorers, may be associated with MEV (miner extractable value) bots.

1.1 Paper Structure

An introduction and executive summary that outlines the concept of *OpAl* and highlights the contributions of this paper is presented in Section 1. In Section 2 we discuss related work and background literature. Section 3 provides a definition of what is meant by *OpAl*. To gain a better understanding of the principles behind *OpAl*, we first define prerequisites and properties of semantic malleability in Section 4, and use them to investigate three different ledger designs (Sections 5 and 6). A proof-of-concept *OpAl* attack in the context of Ethereum is also presented in Section 6. In Section 7 we empirically analyze transaction traces from Ethereum to identify and characterize transactions where ledger context is

accessed. Finally, we consider possible mitigation strategies against algorithmic double-spending (Section 8) and highlight future research directions in Section 9.

2 Related Work

Beyond the related work on double-spending that was mentioned in the introduction, it is important to highlight that prior art has identified a range of security issues in distributed ledgers that tie-into the discussion of *OpAl*.

Luu et al. [55] outlines the problem of *transaction-order dependence* (TOD) in the context of smart contracts, whereby the *ordering* in which transactions are executed can lead to unintended state updates that may affect security. As a counter-measure, the utilization of *guard-functions* is suggested, which constrain transaction validity to desirable input states. Another security issue they cover is *timestamp dependence* in smart contracts. However, Luu et al. does not discuss the possibility that transactions could be designed to *intentionally* carry out adversarial actions based on the state upon which they are executing. More importantly, TOD and timestamp dependence is only analyzed in regard to adversarial ordering and the impact of forks by block producers, and not how regular network participants could also leverage such events to their advantage.

Natoli et al. [64] report an artifact related to semantic malleability that they observe in practice in private blockchain networks, which they refer to as the *blockchain anomaly* in analogy to the *paxos anomaly* [7]. The blockchain anomaly derives from the property that causality between transactions may be violated if there are no corresponding transaction ordering constraints imposed at the consensus level, leading to unintended outcomes.

Gazi et al. [30] describe a form of long-range attack in proof-of-stake systems referred to as *stake bleeding*, whereby an adversary may *replay* valid user transactions from the honest chain in a simulated attack chain in order to collect the transaction fees and increase their stake. This is done to render the simulation more plausible to a victim and avoid having to compromise old staking keys. One of the mitigation strategies they propose is *context-sensitive* transactions, whereby the transaction validity is tied to a particular ledger state through including a previous block hash.

Sergey and Hobor [74] introduce the analogy that smart contracts and accounts in Ethereum are like concurrent objects and threads accessing those objects. Following this analogy helps to explain why user interactions with smart contracts can behave like non-deterministic executions that exhibit issues otherwise known from parallel processing. As a remedy, the work suggests using *synchronization* methods such as *locks* within smart contract logic. The key difference to guard functions from Luu et al. is that this counter-measure is applied on a per-contract basis. Kolluri et al. [48] extends upon this analogy and provides an automated tool for discovering *event ordering (EO)* bugs in smart contracts.

McCorry et al. [59,60] presents constructions for (*hard-*) *fork oracles* which are used in the context of atomic cross-chain trading protocols for betting on the outcome of hardforks and trustless smart-contract based bribing.

Daian et al. [16] documents and outlines the widespread utilization of arbitrage bots for performing frontrunning and transaction reordering in decentralized exchanges, and highlights how the existence of *miner extractable value* (MEV) may lead to consensus instability, e.g., through *time-bandit* attacks. In a time-bandit attack an adversary may attempt to rewrite the suffix of a forkable blockchain through a 51% attack, for which it seeks to subsidize the cost for the required computational power by inserting profitable trades that harvest MEV from previous transactions and ledger states. At a high level, the effect of a successful *OpAl* attack can be comparable to the (retroactive) insertion of such a profitable trade, and it appears conceivable that *opportunistic* time-bandit attacks can also be realized through the herein presented mechanism.

Botta et al. [10] discusses MPC protocols in the context of forking blockchains. Specifically, the paper considers a model with *hasty* players that respond to actions which have not yet stabilized on the ledger by becoming part of the common prefix. They outline the issue of an adversary potentially being able to adaptively play different strategies on forks, gain knowledge from player moves in other forks, and the possibility of replaying insecurely created transactions from one fork in another. The mitigation strategy to the aforementioned issues that is presented is comparable to the replay protection in Gazi et al. [30], whereby the validity of transactions is tied to the state of a particular blockchain fork. We note that while Botta et al. appears closest to the topics discussed within the paper at hand, and effectively highlights the possible detrimental effects of blockchain forks as well as the practical implications of probabilistic finality, the notion of algorithmic double-spending is not considered at all, and double-spending only discussed in the context of the therein presented protocol.

To the best of our knowledge we are the first to identify and make the class of algorithmic double-spending attacks explicit, and also demonstrate the practicability and applicability of *OpAl* to current real-world systems.

3 What is Algorithmic Double-Spending?

In this section we revisit and define double-spending and propose that there exists the overlooked class of algorithmic double-spending, which does not necessitate conflicting actions, i.e., equivocation. We then discuss the implications, such as the possibility of *unintentional* double-spending, and raise the question whether double-spending requires economic damage. We observe that while research on double-spending provides concrete descriptions and formal analyses of particular instantiations of double-spending attacks, e.g., [42,2,43,36], a general definition of double-spending appears to be outstanding. A clearer definition may not only aid with classification efforts, but could also help identify new or overlooked attack forms. Motivated by this novel class of algorithmic double-spending attacks we present within this work, we hereby set out to propose such a more general definition:

Definition 1 (Double-Spending Attack). *In a double-spending attack, an adversary attempts to deceive a victim into performing an economic transaction*

directed at the adversary on the basis of a presumed valid system state, which is later revealed to be stale or invalid. Hereby, the adversary’s goal is to be able to reuse any of the resources that form the basis of the economic transaction for other purposes. We distinguish between the following double-spending attacks:

- **Equivocation-Based**, whereby the adversary issues multiple conflicting actions in the system, one of which is aimed at fooling the victim, and where at most one of the issued actions can eventually be performed in the system.
- **Algorithmic**, whereby the adversary performs a single action that can have different semantic meanings, depending on the system state in which they are interpreted, and where the interpretation of this action in some stale or invalid system states can be used to deceive the victim.

At the core of this work lies the insight, that double-spending may be facilitated through other means than the classical notion of requiring equivocation-based conflicting actions by an adversary. Algorithmic double-spending builds on a simple property that can be observed in various real-world distributed ledger designs with expressive transaction semantics:

Observation 1 (semantic malleability) *Given a transaction t , it may have different semantic outcomes, depending on the ledger state and environment upon which t is executed.*

We refer to this property as *semantic malleability* due to the fact that external factors, such as the consensus protocol and its ordering guarantees [45,87,51,44], as well as other actors in the system who may be *rushing* [49] e.g., in the context of frontrunning [23,16], and sandwich-attacks [89], are able to transition the state in a way that is able to *malleate* the intended semantics of transactions. From this observation, we can rather intuitively derive a basic strategy for an algorithmic double-spending attack: An adversary can encode both, the regular payment to the merchant, as well as an alternative malicious action, e.g., payment to herself, as different execution paths within a single transaction. The control flow of the transaction is designed to conditionally branch, depending on the ledger state σ at the time the transaction is processed by a miner. If the same transaction is included in a different state σ' , i.e., a fork, the “hidden” algorithmic double-spend is triggered without active participation from the attacker. Figure 1 illustrates this difference to equivocation-based double-spending.

The concept of algorithmic double-spending raises interesting challenges, two of which we outline in more detail. First, up until now *unintentional* double-spending, for example as a result of technical failures, did not appear of particular concern. Prior art identifies potential vulnerabilities that arise from order dependence in smart contracts [55,74,48] and violations of transaction causality in forks that can have unintended side-effects, which relate to the *Paxos anomaly* [64]. We expand upon these insights by highlighting that semantic malleability can lead to unintentional algorithmic double-spending as a result of unanticipated transaction reordering that causes state changes within a blockchain fork.

Hereby, it is difficult to distinguish between an intentional attack or unfortunate circumstances. Consider the following example: a user wishing to perform a payment with the stablecoin DAI on the Ethereum network first purchases the required amount in a preceding transaction through a decentralized exchange (DEX) [88], e.g., Uniswap [1]. If a blockchain fork occurs, replaying the above transaction sequence in the same order may not guarantee the same outcome e.g., if the necessary amount of DAI for the trade to be successfully executed is not (yet) available on the DEX because another user’s trade transaction that would have provided sufficient DAI was not yet processed due to reordering.

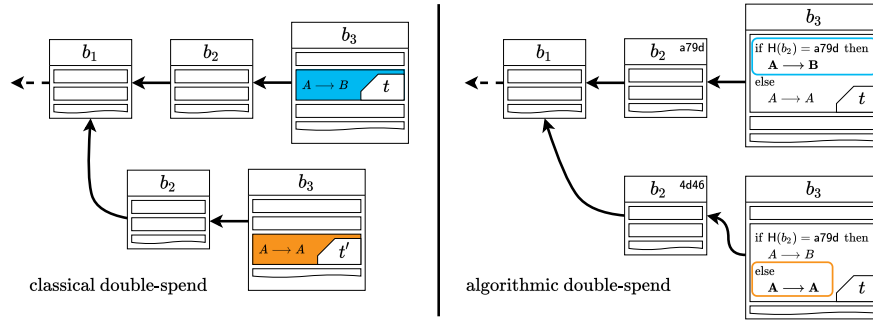


Fig. 1. Conceptual difference between equivocation- and algorithmic double-spending. Notice that in the former case $t \neq t'$ while in the latter case $t = t$.

Second, in stateful smart contract systems double-spending may not only be performed solely at the economic level through coin-reuse. For example, Botta et al. [10] highlights the need for mitigation strategies against an adversary leveraging forks in MPC protocols with hasty players. In this regard, double-spending attacks can be aimed at biasing the outcome of a MPC, which may not be quantifiable in terms of economic gain. Similarly, increasing the miner fee of a transaction may require a user to equivocate, raising the question if such behavior should be subsumed under the notion of double-spending. This presents the interesting problem how any divergent system behavior within forks, be it through equivocation- or algorithmic double-spending, should be addressed if it is not immediately apparent that they were intended for unjust economic gain. Notice that in our Definition 1 for double-spending, we assume some economic transaction from the victim to the adversary.

4 System Model and Assumptions

Within this section, we identify prerequisites and underlying properties that enable algorithmic double-spending. Our analysis is based on an intentionally simple system model to accommodate different ledger designs. We define the concept of *semantic malleability* that we introduced in Section 3 and argue that

ledgers with semantically malleable transactions are vulnerable to algorithmic double-spending, and thus *OpAl* attacks. In our analysis, we show that any distributed ledger that is robust to semantic malleability must satisfy two necessary properties, namely, *eventual replay validity* and *replay equivalence*.

Following Luu et al. [55], we conceptually view a blockchain as a transaction-based RSM, where its state is updated after every transaction. We denote \mathcal{S} the set of all possible system states and $\sigma \in \mathcal{S}$ a single system state. The initial starting state of a blockchain is defined as σ_0 . A valid transition from state σ to σ' , via transaction t , is denoted as $\sigma \xrightarrow{t} \sigma'$. $\text{PAST}(\sigma_n)$ is defined as the ordered list of transactions $\mathcal{T} = (t_1, t_2, \dots, t_n)$, that, when applied to σ_0 , lead to state σ_n . If there exists a non-empty sequence of transactions starting from state σ_a to state σ_b , we call σ_a a predecessor of σ_b , in short $\sigma_a < \sigma_b$. The predicate $\text{VALID}(t, \sigma)$ represents the transaction validation rules of the underlying protocol and returns TRUE iff the transaction t is considered valid (executable) in state σ . We only consider transactions that are non-trivial, and that no equivocation of transactions happens. A transaction t is non-trivial if it produces a different output state when executed on some valid input state, i.e., $\exists \sigma, \sigma' \in \mathcal{S}, \text{VALID}(\sigma \xrightarrow{t} \sigma') \wedge \sigma \neq \sigma'$. We assume that block producers, e.g., miners, adhere to protocol rules and *transaction liveness* is guaranteed, i.e., any valid transaction will eventually be executed.

Executing a transaction t in state σ alters (part of) the state σ and thus results in a new state σ' . The changes are captured by the function $\text{DIFF}(t, \sigma)$. For example, consider a state $\sigma = \{\text{Alice: 6, Bob: 5, Carol: 4}\}$ represented as an account-value mapping, and a transaction t , where Alice gives 2 coins to Bob. Then $\text{DIFF}(t, \sigma) = \{\text{Alice: -2, Bob: +2}\}$ captures the balance changes of Alice and Bob while other parts of the state (Carol's balance) remain unaffected. In this example a single account-value mapping is called a *substate*. Note that it is possible that the effects of executing the same transaction t in two different states are equal, i.e., $(\sigma_a \neq \sigma_b) \wedge (\text{DIFF}(t, \sigma_a) = \text{DIFF}(t, \sigma_b))$.

We consider a transaction t to be a sequence of operations (computations) that lead to a state transition. A transaction is *semantically malleable*, if the available operations, which are used to define the semantics of the transaction, allow the control flow of the execution to branch conditionally based on the particular input state σ . The following two properties we define are necessary, but not sufficient, for a ledger to be robust against semantic malleability. We refer to these properties as *replay equivalence* and *eventual replay validity*. Replaying the same ordered set of transactions on some initial state σ_0 should always yield the same state transitions and final state, and the validity of transactions should not be affected by the environment, e.g., the number and size of already mined blocks.

Definition 2 (replay equivalence). *Assuming that no transaction equivocation happens: A transaction t satisfies replay equivalence, if executing t in all candidate states where t is executable (valid) leads to the same changes in the respective (sub)states:*

$$\begin{aligned} \forall \sigma_a, \sigma_b \in \mathcal{S}, \\ (\text{VALID}(t, \sigma_a) \wedge \text{VALID}(t, \sigma_b)) \implies (\text{DIFF}(t, \sigma_a) = \text{DIFF}(t, \sigma_b)). \end{aligned} \quad (1)$$

Definition 3 (eventual replay validity). *Assuming that no transaction equivocation happens: If a transaction t is found executable (valid) in some state σ_a , then it either remains executable (valid) or has already been executed in predecessor states of σ_a :*

$$\begin{aligned} \forall \sigma_a, \sigma_b \in \mathcal{S}, \\ (\text{VALID}(t, \sigma_a) \wedge \sigma_a \prec \sigma_b) \implies (t \in \text{PAST}(\sigma_b) \vee \text{VALID}(t, \sigma_b)). \end{aligned} \quad (2)$$

Definition 4 (semantic malleability). *A transaction t is semantically malleable if it violates replay equivalence and/or eventual replay validity.*³

5 Semantic Malleability of Bitcoin and Cardano

For the following investigation, we set aside the orthogonal topic of *how* to create blockchain forks of sufficient depth to facilitate double-spending attacks. Instead, we are interested in identifying if, in principle, the designs are vulnerable to semantic malleability by evaluating whether the aforementioned necessary properties are violated. We first consider Bitcoin and Cardano within this Section, and then cover Ethereum separately in Section 6. Each ledger represents an instantiation of a Nakamoto-style blockchain with distinct design differences. Bitcoin [63] is UTXO based and facilitates a highly limited, non-Turing complete scripting language for transaction semantics.[4] Cardano [15] adopts the EUTXO model [11], which leverages advantages of a stateless UTXO design with the expressiveness of Turing-complete smart contracts that can carry state.

Bitcoin: In Bitcoin, transactions are based on the so-called *unspent transaction outputs* (UTXO) model [18] and contain simple (deterministic) Boolean functions, called *Scripts*, that determine the transaction semantics [4]. Bitcoin’s UTXO model is stateless and non-Turing complete. A key aspect of the UTXO model is that transactions are deterministic and bound to a single execution by committing to the exact input (sub)states, i.e., UTXOs, that a transaction consumes, and a precise set of output UTXOs, that the transaction produces. We refer the interested reader to Atzei et al. for a formal model of Bitcoin transactions and Script.[4]

Furthermore, within Bitcoin transactions the access to external ledger state is not made explicitly by including it as an input in the transaction, but implicitly through Scripts or when defining the validity of the transaction in terms of the block height or current time at the protocol level. There currently exist only a limited number of primitives that can be used to constrain the validity of a

³ An equivalent standalone definition of this property is given in Appendix B.

transaction to some external context. Specifically, it is possible to define some relative or absolute time, in relation to that of the ledger context, from which point onward a transaction may become *valid* [77]. However, *it is not possible to permanently invalidate a previously valid transaction that depends on ledger context*, i.e., in a live blockchain, there is a future point in time where this dependency is satisfied. Therefore, in principle, the Bitcoin UTXO model could satisfy eventual replay validity. However, we show that in case of deep forks, eventual replay validity can be violated by coinbase transactions, making Bitcoin-like UTXO cryptocurrencies theoretically vulnerable to semantic malleability.

Theorem 2 (Semantic malleability of Bitcoin-like UTXO cryptocurrencies with coinbase transaction). *A Bitcoin-like UTXO based cryptocurrency is affected by semantic malleability if it programmatically allows the issuance of special per-block transactions as payout, i.e., coinbase transactions, transferring collected fees and/or rewards for block creation.*

Proof. We show that A Bitcoin-like UTXO cryptocurrency is affected by semantic malleability by constructing a counterexample violating the *eventual replay validity* property: Let σ_a be some blockchain state and $t_c \neq t'_c$ two different coinbase transactions (e.g., rewarding different miners) that are valid in this state if included by a newly mined block, i.e., $\text{VALID}(t_c, \sigma_a) \wedge \text{VALID}(t'_c, \sigma_a)$. Let there be a blockchain with a new block containing t_c st. $\sigma_a \xrightarrow{t_c} \sigma_b$ and thus $\sigma_a \prec \sigma_b$. In Bitcoin-like UTXO cryptocurrencies, the coinbase transaction can only be issued at the beginning of each block and is tied to the respective block height⁴. Therefore, the other coinbase transaction t'_c cannot be included anymore in state σ_b . The reason for this is that executing the block containing t_c (and potentially other transactions) necessarily leads to a state σ_b with increased block height. Therefore, there exists a σ_b st. $t'_c \notin \text{PAST}(\sigma_b) \wedge \neg \text{VALID}(t'_c, \sigma_b)$ which violates *eventual replay validity*. \square

In practice, the potential consequences of the semantic malleability of coinbase transactions are mitigated by the maturation period of 100 blocks, after which transactions can be included that spend coinbase UTXOs. As an attack example, consider a transaction with an output from a recent coinbase transaction that is spendable (i.e., has matured for 100 Blocks) as one of its input UTXOs. If a sufficiently deep blockchain fork, of say 144 blocks, occurs and this coinbase transaction does not exist in the chain anymore, the depending transaction using its UTXO as input can not be replayed within a fork and becomes invalid. Therefore, coinbase transactions could facilitate algorithmic double-spending.

The design of forkable Nakamoto-style cryptocurrencies, which provide payouts in terms of fees/block rewards to incentivize participation, necessarily require payments depending on the state of the blockchain, i.e., context, which inherently violates eventual replay validity. Thus, in Section 8.1, we raise the question whether characterizing Nakamoto-style ledgers as replicated state machines (RSM) is accurate in light of algorithmic double-spending.

⁴ cf. <https://github.com/bitcoin/bips/blob/master/bip-0034.mediawiki>

Cardano: Cardano [15] is based on a line of research on provably secure proof-of-stake Nakamoto-style blockchains [47,17,5,46], which we subsume under the term *Ouroboros*. Ouroboros is a provably secure proof-of-stake Nakamoto-style blockchain protocol. Ouroboros, as it is currently realized in Cardano, offers probabilistic finality guarantees and the existence of temporary blockchain forks is possible. Cardano adopts the *Extended UTXO* (EUTXO) model [14,11], that was conceived to leverage desirable properties of Bitcoin’s UTXO design for more expressive transaction semantics [14]. Conceptually, to support stateful Turing-complete smart contracts in EUTXO, the UTXO model is extended in the following (from Chakravarty et al. [14]): i) outputs can contain arbitrary contract-specific data; ii) Scripts, which are referred to as *validators* in the EUTXO model, receive the entire transaction information, including its outputs, as *context* next to the contract specific data, and can impose arbitrary validity constraints on the transaction; iii) a *validity interval* is added for transactions, which is specified as an interval of “ticks”⁵, whereby any Scripts which run during validation can assume that the current time is within that interval, but do not know the precise value of the current tick;

A key property the EUTXO model inherits from the UTXO model is that the execution of a transaction during validation is entirely deterministic and solely determined by its inputs. Equivocation is hence required to achieve a different semantic result. In terms of our necessary properties to achieve robustness against semantic malleability, *replay equivalence* follows analogous to Bitcoin.

However, as Brünjes and Gabbay [11] crucially point out, the EUTXO model allows restricting the validity of transactions to time intervals, which renders the result of transaction processing dependent on the ledger context. Unlike Bitcoin, in Cardano transactions can be permanently invalidated based on ledger context. Hence, *eventual replay validity* is not satisfied and semantic malleability possible.

Corollary 1 (Semantic malleability of Cardano-like EUTXO cryptocurrencies that support limited validity transactions). *A EUTXO based cryptocurrency is affected by semantic malleability, if it programmatically allows the issuance of limited validity transactions which are valid at some point in the chain, but become invalid after a certain block height or time interval.*

Proof. We show that Cardano-like EUTXO cryptocurrencies that support limited validity transactions are semantically malleable, by pointing out that the desired properties of such transactions directly negate and thus violate *eventual replay validity*. Let t_v be a limited validity transaction and σ_a be some blockchain state where $\text{VALID}(t_v, \sigma_a)$, which is true when the specified criteria (block height or time) is satisfied. By definition of a limited validity transaction, there must exist a state $\sigma_a \prec \sigma_b$ st. $\neg \text{VALID}(t_v, \sigma_b)$. Due to forks, or congestion, it might be the case, that t_v is not included until σ_b is reached, thus $t_v \notin \text{PAST}(\sigma_b)$. Therefore, t_v is invalidated after this point and cannot be included in any other

⁵ Chakravarty et al. [14] assume that in practice a tick will correspond to a block number or block height.

subsequent block. Hence, by the construction of limited validity transactions,

$$\exists \sigma_a, \sigma_b (\text{VALID}(t_v, \sigma_a) \wedge \sigma_a \prec \sigma_b \wedge t_v \notin \text{PAST}(\sigma_b) \wedge \neg \text{VALID}(t_v, \sigma_b)),$$

which is exactly the negation of our definition of *eventual replay validity*. \square

As an example, consider a payment transaction to a merchant where the validity is constrained to a specific block height. Thus, an *OpAl* attack is triggered if the transaction does not make it into a block in time during a fork.

6 Semantic Malleability in Ethereum

Ethereum [82] adopts an account-based model and offers expressive transaction semantics that can draw upon stateful Turing-complete smart contract functionality. Due to the various ways in which replay equivalence and eventual replay validity can be violated in Ethereum, we omit a formal analysis and directly discuss a proof-of-concept (PoC) *OpAl* attack and its practical implications.

Our attack design is inspired, on the one hand, by *hardfork oracles*, which McCorry et al. [59] discusses in the context of atomic-trade protocols during hardforks, and, on the other hand, by the notion of *context sensitive transactions* Gazi et al. [30] describes as a replay protection mechanism in *stake-bleeding attacks*. An informal statement that encapsulates the intended transaction semantics for our PoC *OpAl* attack is the following:

“IF *this transaction is included in a blockchain that contains a block with hash 0xa79d* THEN *pay the merchant*, ELSE *don’t pay the merchant*.”

Essentially, our attack is based on the insight that a transaction can act as its own *fork oracle* for conditionally branching its execution. In the following, we first outline the construction of such a fork oracle in more detail and then present a PoC attack that allows transactions with the above semantics to be created.

6.1 How to Construct an *OpAl* Fork Oracle in Ethereum

The concept of employing a fork oracle to distinguish between branches of (hard)forks within a transaction or smart contract execution was proposed in cryptocurrency communities [57,24], as well as research [59,60,38]. Hereby, a frequent goal is achieving *replay protection* [59]. McCorry et al. [59] outlines how fork oracles can be leveraged to realize atomic trades across hardforks. Constructing a smart contract based fork oracle if the underlying forks do not offer replay protection can be challenging [59]. McCorry et al. [60] demonstrate through *history revision bribery* how (equivocation-based) double-spending can be leveraged to realize a fork oracle for a smart contract based bribing scheme for incentivizing forks. Hereby, the fork oracle is not used to facilitate (algorithmic) double-spending. Rather, the mutually exclusive outcomes of the double-spend in different forks are relied upon to actually implement the oracle. Surprisingly, to the best of our knowledge, the idea of using fork oracles to *algorithmically* trigger double-spending was not yet considered.

Block-Hash Based Fork Oracle The fork oracle we propose is inspired by a simple and elegant technique to achieve replay protection considered in the proof-of-stake (PoS) setting [52,30]. Hereby, the hash of a recent block is included in a transaction, and the transaction is only considered valid for blockchains that contain this block in their prefix. Gaži et al. [30] refer to this mechanism as *context sensitive transactions*. Essentially, context sensitive transactions already implicitly realize the attack semantics described above.⁶ In case a fork of sufficient depth occurs, this replay protection mechanism ensures that transactions become invalid at the protocol level, and the double-spending “attack” is realized algorithmically through the underlying protocol rules. Ethereum does not natively support context sensitive transactions, however, this functionality can be emulated with smart contract code using EVM primitives that expose ledger context, such as the BLOCKHASH opcode [82]. It is hence possible to programmatically act upon the existence of a particular block, or other ledger context, as part of an Ethereum transaction.

Fork Oracle Discussion A downside of hash-based fork oracles is the reliance on a commitment to *previous* ledger state, thereby requiring a fork of at least depth-2 to trigger the attack. However, it is also possible to construct oracles for forks of depth-1. The key difference between a depth-1 fork oracle and a hash-based fork oracle is that the latter is based on ledger state which is *known*, whereas the former is based on some *prediction* of the future state at the time the transaction is processed. Hence, depth-1 fork oracles generally offer weaker probabilistic guarantees for identifying forks. For example, consider the EVM COINBASE opcode that returns the *current* block’s beneficiary address [82]. Instead of specifying the highest known block hash as the branching condition, an adversary could use the beneficiary address of a large mining pool⁷ in a depth-1 *OpAl* attack. Hereby the transaction semantics depend on whether the transaction is included in a block from the targeted mining pool or some other miner. Generally speaking, in Nakamoto-style proof-of-work ledgers the next block producer is not known in advance. However, we note that in some PoS protocols this can be different [73], thereby allowing for more reliable depth-1 fork oracles.

Another limitation of the hash-based fork oracle specific to the EVM is the restriction that the BLOCKHASH opcode only returns hashes within a 256 block lookback window, and 0 otherwise [82]. Hence, if a transaction is processed in a block that exceeds 257 blocks after the height of the blockhash commitment, the oracle will falsely report a fork and trigger the attack branch. We argue that in the case of our intended *OpAl* semantics this limitation is unproblematic, as the transaction would simply transfer the funds back to the attacker.

⁶ Thereby introducing the possibility of *unintentional OpAl* attacks (see Section 3).

⁷ We note that in Ethereum address reuse in the coinbase by miners is prevalent.

```

1  pragma solidity 0.8.4;
2  // This contract acts as an OpAl forwarding proxy for transactions.
3  contract Opal {
4      address public owner;
5
6      modifier onlyOwner() {
7          require(isOwner(msg.sender));
8      }
9
10     constructor() {
11         owner = msg.sender;
12     }
13
14     fallback() external payable {}
15     receive() external payable {}
16
17     function isOwner(address addr) public view returns(bool) {
18         return addr == owner;
19     }
20
21     function cashOut(address payable _to) public onlyOwner {
22         _to.transfer(address(this).balance);
23     }
24
25     // forwarding function implementing opportunistic double-spending (OpAl)
26     function forward(address payable destination, bytes32 commitblockHash,
27         uint commitblockNumber, bytes memory data)
28         onlyOwner public payable returns(bool success) {
29         if (blockhash(commitblockNumber) == commitblockHash)
30             assembly { success := call(gas(), destination, callvalue(),
31                 add(data, 0x20), mload(data), 0, 0)
32         }
33     }
34 }

```

Listing 1.1. Solidity *OpAl* contract that implements a basic fork oracle by only forwarding transactions if the provided commitment to a block hash can be resolved.

6.2 Proof of Concept *OpAl* Attack Contract

Di Angelo and Salzer present a comprehensive empirical analysis of wallet contracts on Ethereum [19,20]. Of the identified properties, in particular, designs that support *flexible transactions*, i.e., forwarding of arbitrary calls, appear suitable for augmentation to support the creation of *OpAl* transactions. Their empirical data shows that at least *tens of thousands* of contracts supporting flexible transactions are currently deployed in Ethereum, suggesting practical use-cases for such contract patterns, even without an *OpAl* augmentation. Our attack requires minimal modifications, and the interaction pattern is almost identical.

In the following, we present a minimal *fully viable PoC OpAl attack smart contract* written in Solidity [81], that relies on the aforementioned hash-based fork oracle. Our contract code (Listing 1.1) is loosely based on the Executor contract from the Gnosis-Safe Wallet [61], which allows the forwarding of arbitrary function calls. Instead of forwarding a call directly, the contract first evaluates if the block hash at a particular height of the current ledger matches the commitment hash that is provided as an additional parameter in the transaction data. This is realized through the *blockhash()* function [25]. If the blockhash matches the

commitment, the function call is forwarded. Else, no action is performed, i.e., the action is reversed whenever the transaction is replayed in a fork.

Outline of the Attack An adversary wishing to engage in *OpAl* first needs to deploy the attack contract. Once the contract is successfully deployed, whenever they wish to perform a transaction with *OpAl* functionality, instead of calling a function $f()$ in the target contract or sending funds directly, they simply forward this call to the *forward()* function (Line 15 in Listing 1.1) of the deployed attack contract, together with the appropriate parameters. Specifically, the adversary generates transaction t that calls *forward* in the attack contract with the following parameters: i) the target address; ii) the block hash and height h of the current chain tip; iii) the encoded function name to be called at the target $f()$ together with its parameters; iv) any Ether that shall be sent; and broadcasts t to the network. Ideally, the transaction fee is high enough for t to be immediately included in the next block $h + 1$. Otherwise, the required fork depth increases in the number of blocks the chain grows between the creation and inclusion of t .

To the recipient of t , the interaction pattern will appear as if the user employed a regular wallet contract. Unless they perform an analysis of the execution trace, the malicious behavior only becomes apparent once the attack conditions are triggered, i.e., during a fork. In case the adversary is lucky and a fork at, or before, height h occurs, and their transaction is replayed within this fork, the alternative attack branch of the contract is executed automatically.

6.3 Cost Overhead of PoC Attack in Ethereum

We quantify the additional costs incurred when augmenting a transaction with *OpAl* capabilities by deploying our attack contract in a private Ethereum test-net and measuring the gas utilization for basic interactions, such as ERC-20 token [79] transfers. Our PoC *OpAl* attack adds a constant overhead of gas that depends on the number of parameters supplied to the target function $f()$. The deployment transaction for the contract in Listing 1.1 required 393 175 gas. As it is not essential for the contract to be deployed in a recent block, and can be done well in advance of any attacks, we assume a gas price of 50 GWei, which translates to deployment costs of ≈ 0.02 Ether or, at an exchange rate of 2 000 USD, approximately 40 USD. Note that this contract needs to be deployed once, after which the only overhead derives from using the forwarding function. For ERC-20 token interactions (*approve*, *transfer*, *transferFrom*), using *OpAl* adds $\approx 3\,000$ gas, which equates to $\approx 8\%$ overhead. At the time of writing, assuming a gas price of 100 GWei for timely inclusion⁸ of the transaction, this overhead translates to ≈ 0.6 USD higher fees if a transaction is augmented to support *OpAl* attacks, rendering our attack a viable default strategy for most cases.

⁸ For simplicity we consider legacy transactions and omit pricing based on EIP-1559.

7 Empirical Analysis of Ethereum Transaction Traces

We empirically analyze the execution traces of 922 562 transactions from 5 000 Ethereum blocks in order to identify and characterize transactions where ledger context is accessed. Hereby, block selection for the analysis was performed in batches of 100 consecutive blocks every 1000 blocks, starting from block height 14 010 000 up to block 14 059 099 to obtain a sample spread over a wider time window. The selection of blocks for our analysis was necessitated due to the steep storage and processing requirements for analyzing full execution traces. For every considered block, we parse the execution trace of all included transactions and record whether the trace contains EVM opcodes that are characteristic for accessing the ledger context. The specific opcodes⁹ that we considered are highlighted in Table 1. Our analysis reveals that 231 271 transactions, or $\approx 25\%$, include at least one of these opcodes, whereby roughly every 5th transaction uses `TIMESTAMP`, while the other opcodes are encountered considerably less often.

Opcode (OP)	TIMESTAMP	SELFBALANCE	NUMBER	BALANCE	CHAINID	BASEFEE	BLOCKHASH	COINBASE	DIFFICULTY	GASLIMIT
TX containing OP	199731	63594	36859	4324	8253	777	3425	3882	1251	906
% of TX with OP	21.65%	6.893%	3.995%	0.469%	0.895%	0.084%	0.371%	0.421%	0.136%	0.098%
Blocks cont. OP	4886	4767	4529	2265	3071	641	1830	1897	812	545
% of Blocks with OP	97.72%	95.34%	90.58%	45.3%	61.42%	12.82%	36.6%	37.94%	16.24%	10.9%

Table 1. EVM Opcode occurrence within the analyzed block range.

Of particular interest are transactions that include *both* `BLOCKHASH` and `NUMBER` opcodes in their traces, as this combination is also present in our PoC *OpAl* attack. We are able to identify 3 338 transactions with such an *OpAl*-like opcode signature within 1 823 ($\approx 36\%$) of the analyzed blocks. Table 2 in the Appendix shows the top 10 contract addresses that these transactions were directed at, as well as a generalized categorization of their purpose based on publicly available information. Analyzing the decompiled¹⁰ bytecode of the contract with the second most *OpAl*-like transaction interactions, we indeed discover an *OpAl*-like pattern. Listing 1.2 highlights the relevant code section, which, in plaintext, evaluates whether the first 4 Bytes of the previous block hash match those stored as part of the transaction data and reverts the execution otherwise. We further confirm this behavior by observing transactions to the aforementioned contract that were reverted due to an incorrect commitment¹¹. While we believe that this pattern is likely intended to render the transaction *context sensitive* to prevent execution in an undesirable state, it could nevertheless be used for *OpAl* attacks simply by subsequently using the transferred/traded funds for payments to a victim.

⁹ Cf. the Ethereum Yellow paper [82] for details on EVM opcodes and their behavior.

¹⁰ Cf. <https://ethervm.io/decompile/0x000000000035B5e5ad9019092C665357240f594e>

¹¹ Cf. txn: 0x2368617cf02cf083eed2d8691004c1ad0176976b6fa83873bc6b0fd7de4cc7fc

```

1 function func_060C() {
2   if (msg.data[0x04:0x24] >> 0xe0 ==
3       block.blockHash(block.number + ~0x00) >> 0xe0) { return; }
4   // ... code omitted for brevity
5   revert(memory[0x60:0xc4]); }

```

Listing 1.2. Code snippet from decompiled contract (tagged as MEV bot) showing *OpAl*-like pattern. Notice that `~0x00` corresponds to -1 in Two’s complement representation.

8 Mitigation Strategies against *OpAl*

Having outlined the principles behind algorithmic double-spending, we now discuss possible prevention or mitigation strategies. Hereby, we broadly distinguish between two categories: i) Approaches that address instability in consensus, i.e., a *lack of finality*. ii) Approaches that seek to limit the effects of *semantic malleability*. Finally, we discuss if the characterization of blockchains as replicated state machines is accurate in light of semantic malleability.

Mitigating *OpAl* through stronger consensus guarantees: Essentially, the majority of distributed ledgers rely on *consensus*¹² to agree upon the order of transactions among participants in order to prevent double-spending [33]. Thus, one possible defensive approach against *OpAl* attacks is to prevent players from concurrently interacting with malleable ledger state until it has sufficiently stabilized. However, such a pattern may not be desirable for users, as it can lead to long waiting periods. In this regard, it appears advantageous to achieve fast and guaranteed *consensus finality*, which remains an active research topic for decentralized ledger designs [12,68,66]. Our Definition (Def. 1) of double-spending highlights the requirement of some stale or invalid system state in order to fool a victim. The existence of hasty players who are willing to act upon such state renders double-spending attacks feasible in practice, even if the consensus protocol in principle could provide stronger guarantees. In this regard, effective mitigation strategies to combat double-spending may also entail the stricter enforcement of safe interaction patterns in client software and cryptocurrency wallets, and a better understanding of the behavior and mental models of cryptocurrency users [22,50,58].

However, if the security assumptions of the underlying system are compromised, in particular, Nakamoto-style distributed ledgers can suffer from deep forks where previously assumed stable ledger state is reverted. Aside from the potential of targeted attacks against the protocol [3,78], technical failures¹³ can also lead to such a violation of the security assumptions [59,54,62]. Table 3 in the Appendix highlights several documented unintentional forks in Bitcoin and

¹² An interesting recent result in this regard is presented in Guerraoui et al. [33], which proves that for simple *asset transfer* consensus is, in principle, not necessary for double-spending prevention.

¹³ We note that scheduled protocol updates carry a risk of unintentional forks, and an adversary may try to leverage this by performing *OpAl* transactions at that time.

Ethereum due to technical failures. Notice that in this regard there is a crucial difference between *OpAl* and equivocation-based double-spending. In the latter, an adversary has to actively monitor the network for forks and disseminate conflicting double-spending transactions that are at risk of being easily detected and prevented at the peer-to-peer layer [42,31]. *OpAl* attacks and algorithmic double-spending, on the other hand, may prove particularly severe. Any transaction that was included in a blockchain that is replayed on a fork faces the risk of triggering a hidden *OpAl* attack without requiring any active participation from the initial creator of the transaction. If a fork in excess of k blocks occurs, triggered *OpAl* attacks have a high probability of success, as the corresponding economic transaction is likely to have already happened. A possible mitigation strategy to limit the effects of *OpAl* in deep forks is the utilization of checkpointing [41]. Another line of research seeks to strengthen the guarantees of Nakamoto consensus by achieving consensus finality [12,69,21,66]. It may also be preferable to sacrifice *liveness* by halting execution rather than risking systemic risk through *OpAl* attacks.

Mitigating Semantic Malleability: As we have shown in Sections 3 and 4, semantic malleability lies at the core of enabling algorithmic double-spending. Semantically malleable transactions allow for different state transitions, depending on the input state and execution environment at the time of processing – a property that is generally observed within smart contract platforms. In this regard, we believe that the expressive transaction semantics associated with smart contract functionality poses a fundamental challenge when trying to combat algorithmic double-spending. Drawing upon the concept of *guard functions* from Luu et al. [55] and *context sensitive transactions* Gaži et al. [30] and Botta et al. [10] rely on, transaction validity should more explicitly be constrained to input states that only lead to desirable outcomes for the sender. While such patterns do not prevent the possibility of algorithmic double-spending, they can avert that a user’s transaction executes in a state that leads to an undesirable outcome. In light of recent research in regard to *order-fairness* in consensus [45,87,51,44], the aforementioned pattern could also help to mitigate the potential negative impact of malicious orderings. Similar to the concept of the *Let’s Go Shopping Defense* [37], a highly questionable mitigation strategy might be to oneself proactively engage in *OpAl* (counter-) attacks in order to reduce counterparty risk and try to hedge against the potentially detrimental effects of any deep blockchain fork, should it ever occur.

Another mitigation strategy by which to address semantic malleability and algorithmic double-spending is through the analysis and classification of transaction semantics, in order to try and identify potential threats and malicious behavior. Hereby, the challenges lie on the one side, in finding efficient techniques for static and dynamic code analysis that can be applied, in real-time, to identify potentially malicious transactions before they are processed, and on the other side, in how to define what is considered malicious behavior and also enforce any transaction rejection policies within decentralized systems [70,80,27,83,85,35,28,9]

For platforms that do not support expressive transaction semantics, it may appear that the solution to this problem is to enforce only a single valid state transition for transactions, such as the EUTXO model [14] employed by Cardano. However, in this case the possibility of algorithmic double-spending still arises if the *validity* of a transaction can be tied to particular ledger states, which is generally the case. In the UTXO model of Bitcoin [4], transaction expressiveness and access to ledger state are sufficiently constrained to prevent practicable *OpAt* attacks, apart from the possibility of using recent coinbase transactions to limit replay validity in case of deep-forks. However, since the mechanism design of most cryptocurrencies relies on the issuance of rewards to incentivize participation [13], it is unclear if the underlying issue could be completely avoided in practice.

8.1 Can Blockchains be Characterized as State Machines?

In his seminal work on the state machine approach, F. B. Schneider provides the following semantic characterization of a replicated state machine (RSM): “*Outputs of a state machine are completely determined by the sequence of requests it processes, independent of time and any other activity in a system.*” [72] Interestingly, while blockchains are often considered to realize RSMs, e.g., in the model we adopt from Luu et al., we observe (Section 5) that *in practice*, ledger designs appear to actually deviate from this characterization.

First, consider the herein discussed property of semantic malleability in transactions. The existence of Semantic malleability in itself does not violate the above characterization, as a mere reordering of transactions, i.e., requests, may lead to semantic malleability without requiring any access to time or activity within the system. However, in practice, ledger designs often allow transaction semantics to depend on external ledger context that is not solely defined by such requests, i.e., *time* or other *external data* (See Section 6). In essence, being able to define functions that can act upon such context within transaction semantics, such as previous block hashes, the block height, coinbase transactions, or block time, can cause a violation of replay equivalence or eventual replay validity, both of which can be directly derived as required properties of a RSM from the above characterization.

Second, blockchain designs generally offer *rewards* as an incentive mechanism for block producers to participate in the consensus protocol. Under the assumption that a block merely represents an ordered set of transactions, i.e., requests, and transactions can not access any external state defined within blocks, this model would appear to realize a RSM. However, if we include the fact that block rewards represent transactions or state transitions that depend on a particular external state, namely the block itself that justifies the reward, the model is no longer independent of the system state.

We note that one possibility to amend this issue is to either include the creation of blocks as requests, or model state updates entirely from the perspective of blocks and not at the transaction level. The latter approach is, for instance, taken by formal models that analyze Nakamoto consensus [29,67,6]. Nevertheless, even if one considers state machine replication only from the perspective of

blocks and not individual transactions, there can still exist external dependencies on the environment, in particular on time. Consider that receiving late or early blocks may render them (temporarily) invalid by the protocol rules, leading to different possible interpretations of the same sequence of requests and resulting final state depending on the current system time.

9 Conclusion

We have described and analyzed a novel class of double-spending attacks, called (opportunistic) algorithmic double-spending (*OpAl*), and shown that *OpAl* can readily be realized in stateful smart contract platforms, by presenting a proof-of-concept implementation for EVM-based designs. *OpAl* itself does not increase the likelihood or severity of blockchain forks, which are a prerequisite for most double-spending attacks. Instead, *OpAl* allows regular transactions performed by *anyone* to opportunistically leverage forking events for double-spending attacks while evading common detection strategies and offering a degree of plausible deniability. A particularly worrying property of *OpAl* is the ability for already processed transactions to trigger hidden double-spending attacks whenever they are replayed in a fork. Hereby, our empirical analysis of 922 562 transaction traces in Ethereum reveals that transactions with *OpAl*-like semantics already exist in practice. While these transactions are likely intended for a different use case, the effect in case of a fork could still lead to unintentional double-spending. Attacks or technical failures that lead to deep forks may hence pose an even greater systemic risk than previously assumed. It would appear that the most promising mitigation strategy against *OpAl* is achieving fast consensus finality, combined with avoiding semantic malleability in transactions.

We believe that the introduction of algorithmic double-spending as a novel attack category opens up new research directions and highlights the interconnectedness of many important insights in the domain of distributed ledgers. The advent of expressive smart contract systems has created a vast new range of exciting use-cases, but with them also come novel security challenges [56,40,16,?,89] that need to be thoroughly addressed.

References

1. Adams, H., Zinsmeister, N., Salem, M., Keefer, R., Robinson, D.: Uniswap v3 core. Tech. rep., Tech. rep., Uniswap (2021)
2. Androulaki, E., Capkun, S., Karame, G.O.: Two Bitcoins at the Price of One? Double-Spending Attacks on Fast Payments in Bitcoin. In: CCS (2012), <http://eprint.iacr.org/2012/248.pdf>
3. Apostolaki, M., Zohar, A., Vanbever, L.: Hijacking bitcoin: Routing attacks on cryptocurrencies. In: 2017 IEEE Symposium on Security and Privacy (SP). pp. 375–392. IEEE (2017)
4. Atzei, N., Bartoletti, M., Lande, S., Zunino, R.: A formal model of Bitcoin transactions. In: Proceedings of the 22nd International Conference on Financial

- Cryptography and Data Security (FC). Springer (2018), <http://fc18.ifca.ai/preproceedings/92.pdf>
5. Badertscher, C., Gazi, P., Kiayias, A., Russell, A., Zikas, V.: Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In: Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security. pp. 913–930 (2018)
 6. Badertscher, C., Maurer, U., Tschudi, D., Zikas, V.: Bitcoin as a Transaction Ledger: A Composable Treatment (2017), <https://eprint.iacr.org/2017/149.pdf>, published: Cryptology ePrint Archive, Report 2017/149
 7. Birman, K., Malkhi, D., Van Renesse, R.: Virtually synchronous methodology for dynamic service replication. Appears as Appendix A in [4] (2010)
 8. Bonneau, J.: Why buy when you can rent? Bribery attacks on Bitcoin consensus. In: BITCOIN '16: Proceedings of the 3rd Workshop on Bitcoin and Blockchain Research (Feb 2016), <http://fc16.ifca.ai/bitcoin/papers/Bon16b.pdf>
 9. Bose, P., Das, D., Chen, Y., Feng, Y., Kruegel, C., Vigna, G.: Sailfish: Vetting smart contract state-inconsistency bugs in seconds. arXiv preprint arXiv:2104.08638 (2021)
 10. Botta, V., Friolo, D., Venturi, D., Visconti, I.: Shielded computations in smart contracts overcoming forks. In: Financial Cryptography and Data Security-25th International Conference, FC. pp. 1–5 (2021)
 11. Brünjes, L., Gabbay, M.J.: Utxo-vs account-based smart contract blockchain programming paradigms. In: International Symposium on Leveraging Applications of Formal Methods. pp. 73–88. Springer (2020)
 12. Buterin, V., Griffith, V.: Casper the Friendly Finality Gadget (2017), <https://arxiv.org/pdf/1710.09437.pdf>, published: arXiv:1710.09437
 13. Carlsten, M., Kalodner, H., Weinberg, S.M., Narayanan, A.: On the instability of bitcoin without the block reward. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 154–167. ACM (2016), <https://www.cs.princeton.edu/~smattw/CKWN-CCS16.pdf>
 14. Chakravarty, M.M., Chapman, J., MacKenzie, K., Melkonian, O., Jones, M.P., Wadler, P.: The extended utxo model. In: International Conference on Financial Cryptography and Data Security. pp. 525–539. Springer (2020)
 15. Corduan, J., Vinogradova, P., Gudemann, M.: A formal specification of the cardano ledger (2019)
 16. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., Juels, A.: Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In: 2020 IEEE Symposium on Security and Privacy (SP). pp. 910–927. IEEE (2020)
 17. David, B., Gazi, P., Kiayias, A., Russell, A.: Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 66–98. Springer (2018)
 18. Delgado-Segura, S., Pérez-Sola, C., Navarro-Arribas, G., Herrera-Joancomartí, J.: Analysis of the bitcoin utxo set. In: International Conference on Financial Cryptography and Data Security. pp. 78–91. Springer (2018)
 19. Di Angelo, M., Salzer, G.: Wallet contracts on ethereum. In: 2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–2. IEEE (2020)
 20. Di Angelo, M., Salzer, G.: Wallet contracts on ethereum—identification, types, usage, and profiles. arXiv preprint arXiv:2001.06909 (2020)

21. Dinsdale-Young, T., Magri, B., Matt, C., Nielsen, J.B., Tschudi, D.: Afgjort: A partially synchronous finality layer for blockchains. In: International Conference on Security and Cryptography for Networks. pp. 24–44. Springer (2020)
22. Eskandari, S., Barrera, D., Stobert, E., Clark, J.: A first look at the usability of bitcoin key management. In: Workshop on Usable Security (USEC) (2015), http://users.encs.concordia.ca/clark/papers/2015_usec_full.pdf
23. Eskandari, S., Moosavi, S., Clark, J.: SoK: Transparent Dishonesty: front-running attacks on Blockchain. In: arXiv preprint arXiv:1902.05164 (2019), <https://arxiv.org/pdf/1902.05164.pdf>
24. Ethereum Community: Replay attack protection: Include blocklimit and blockhash in each transaction issue#134 ethereum/eips (Jul 2016), <https://github.com/ethereum/EIPs/issues/134>
25. Ethereum Community: Units and globally available variables (Aug 2021), <https://github.com/ethereum/solidity/blob/develop/docs/units-and-global-variables.rst>
26. Eyal, I., Sirer, E.G.: Majority is not enough: Bitcoin mining is vulnerable. In: Financial Cryptography and Data Security. pp. 436–454. Springer (2014), <http://arxiv.org/pdf/1311.0243>
27. Ferreira Torres, C., Baden, M., Norvill, R., Jonker, H.: Ægis: Smart shielding of smart contracts. In: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security. pp. 2589–2591 (2019)
28. Ferreira Torres, C., Iannillo, A.K., Gervais, A., et al.: The eye of horus: Spotting and analyzing attacks on ethereum smart contracts. In: International Conference on Financial Cryptography and Data Security, Grenada 1-5 March 2021 (2021)
29. Garay, J., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: Analysis and applications. In: Advances in Cryptology-EUROCRYPT 2015. pp. 281–310. Springer (2015), <http://courses.cs.washington.edu/courses/cse454/15wi/papers/bitcoin-765.pdf>
30. Gazi, P., Kiayias, A., Russell, A.: Stake-Bleeding Attacks on Proof-of-Stake Blockchains (2018), <https://eprint.iacr.org/2018/248.pdf>, published: Cryptology ePrint Archive, Report 2018/248
31. Grundmann, M., Neudecker, T., Hartenstein, H.: Exploiting Transaction Accumulation and Double Spends for Topology Inference in Bitcoin. In: 5th Workshop on Bitcoin and Blockchain Research, Financial Cryptography and Data Security 18 (FC). Springer (2018), <http://fc18.ifca.ai/bitcoin/papers/bitcoin18-final10.pdf>
32. Gudgeon, L., Moreno-Sanchez, P., Roos, S., McCorry, P., Gervais, A.: Sok: Layer-two blockchain protocols. In: International Conference on Financial Cryptography and Data Security. pp. 201–226. Springer (2020)
33. Guerraoui, R., Kuznetsov, P., Monti, M., Pavlovič, M., Seredinschi, D.A.: The consensus number of a cryptocurrency. In: Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing. pp. 307–316 (2019)
34. Heilman, E., Kendler, A., Zohar, A., Goldberg, S.: Eclipse Attacks on Bitcoin’s Peer-to-Peer Network. In: 24th USENIX Security Symposium (USENIX Security 15). pp. 129–144 (2015), <https://www.usenix.org/system/files/conference/usenixsecurity15/sec15-paper-heilman.pdf>
35. Hu, T., Liu, X., Chen, T., Zhang, X., Huang, X., Niu, W., Lu, J., Zhou, K., Liu, Y.: Transaction-based classification and detection approach for ethereum smart contract. Information Processing & Management **58**(2), 102462 (2021)
36. Iqbal, M., Matulevičius, R.: Exploring sybil and double-spending risks in blockchain systems. IEEE Access **9**, 76153–76177 (2021)

37. Judmayer, A., Stifter, N., Schindler, P., Weippl, E.: Estimating (miner) extractable value is hard, let's go shopping! (2021), <https://ia.cr/2021/1231>
38. Judmayer, A., Stifter, N., Zamyatin, A., Tsabary, I., Eyal, I., Gazi, P., Meiklejohn, S., Weippl, E.: Pay to win: Cheap, crowdfundable, cross-chain algorithmic incentive manipulation attacks on pow cryptocurrencies (2019), <https://ia.cr/2019/775>
39. Judmayer, A., Stifter, N., Zamyatin, A., Tsabary, I., Eyal, I., Gazi, P., Meiklejohn, S., Weippl, E.: Sok: Algorithmic incentive manipulation attacks on permissionless pow cryptocurrencies. In: International Conference on Financial Cryptography and Data Security. pp. 507–532. Springer (2021)
40. Juels, A., Kosba, A., Shi, E.: The ring of Gyges: Investigating the future of criminal smart contracts. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 283–295. ACM (2016), <http://www.arjueels.com/wp-content/uploads/2013/09/Gyges.pdf>
41. Karakostas, D., Kiayias, A.: Securing proof-of-work ledgers via checkpointing. In: 2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–5. IEEE (2021)
42. Karame, G.O., Androulaki, E., Capkun, S.: Double-spending fast payments in bitcoin. In: Proceedings of the 2012 ACM conference on Computer and communications security. pp. 906–917 (2012)
43. Karame, G.O., Androulaki, E., Roeschlin, M., Gervais, A., Čapkun, S.: Misbehavior in Bitcoin: A Study of Double-Spending and Accountability. In: ACM Transactions on Information and System Security (TISSEC). vol. 18, p. 2. ACM (2015), http://www.syssec.ethz.ch/content/dam/ethz/special-interest/infk/inst-infsec/system-security-group-dam/research/publications/pub2015/tissec15_karame.pdf, issue: 1
44. Kelkar, M., Deb, S., Kannan, S.: Order-fair consensus in the permissionless setting (2021)
45. Kelkar, M., Zhang, F., Goldfeder, S., Juels, A.: Order-fairness for byzantine consensus. In: Annual International Cryptology Conference. pp. 451–480. Springer (2020)
46. Kerber, T., Kiayias, A., Kohlweiss, M., Zikas, V.: Ouroboros cryptsinous: Privacy-preserving proof-of-stake. In: 2019 IEEE Symposium on Security and Privacy (SP). pp. 157–174. IEEE (2019)
47. Kiayias, A., Russell, A., David, B., Oliynykov, R.: Ouroboros: A provably secure proof-of-stake blockchain protocol. In: Annual International Cryptology Conference. pp. 357–388. Springer (2017)
48. Kolluri, A., Nikolic, I., Sergey, I., Hobor, A., Saxena, P.: Exploiting the laws of order in smart contracts. In: Proceedings of the 28th ACM SIGSOFT international symposium on software testing and analysis. pp. 363–373 (2019)
49. Kosba, A., Miller, A., Shi, E., Wen, Z., Papamanthou, C.: Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In: Symposium on Security & Privacy. IEEE (2016), <http://eprint.iacr.org/2015/675.pdf>
50. Krombholz, K., Judmayer, A., Gusenbauer, M., Weippl, E.R.: The Other Side of the Coin: User Experiences with Bitcoin Security and Privacy. In: International Conference on Financial Cryptography and Data Security (FC) (2016), https://www.sba-research.org/wp-content/uploads/publications/TheOtherSideOfTheCoin_FC16preConf.pdf
51. Kursawe, K.: Wendy, the Good Little Fairness Widget. arXiv preprint arXiv:2007.08303 (2020)

52. Larimer, D.: Transactions as proof-of-stake (November 2013), <https://github.com/super3/invictus.io/blob/master/assets/pdf/TransactionsAsProofOfStake10.pdf>
53. Liao, K., Katz, J.: Incentivizing blockchain forks via whale transactions. In: International Conference on Financial Cryptography and Data Security. pp. 264–279. Springer (2017), <http://www.cs.umd.edu/jkatz/papers/whale-txs.pdf>
54. Lovejoy, J.P.T.: An empirical analysis of chain reorganizations and double-spend attacks on proof-of-work cryptocurrencies. Master’s thesis, Massachusetts Institute of Technology (2020)
55. Luu, L., Chu, D.H., Olickel, H., Saxena, P., Hobor, A.: Making Smart Contracts Smarter. In: 23rd ACM Conference on Computer and Communications Security (ACM CCS 2016) (Oct 2016), <https://eprint.iacr.org/2016/633.pdf>
56. Luu, L., Teutsch, J., Kulkarni, R., Saxena, P.: Demystifying incentives in the consensus computer. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. pp. 706–719. ACM (2015), <http://www.comp.nus.edu.sg/prateeks/papers/VeriEther.pdf>
57. Maersk, N.: Thedaohardforkoracle (Jul 2016), <https://github.com/veox/solidity-contracts/blob/TheDAOHardForkOracle-v0.1/TheDAOHardForkOracle/TheDAOHardForkOracle.sol>
58. Mai, A., Pfeffer, K., Gusenbauer, M., Weippl, E., Krombholz, K.: User mental models of cryptocurrency systems—a grounded theory approach. In: Sixteenth Symposium on Usable Privacy and Security ({SOUPS} 2020). pp. 341–358 (2020)
59. McCorry, P., Heilman, E., Miller, A.: Atomically Trading with Roger: Gambling on the success of a hardfork. In: CBT’17: Proceedings of the International Workshop on Cryptocurrencies and Blockchain Technology (Sep 2017), <http://homepages.cs.ncl.ac.uk/patrick.mc-corry/atomically-trading-roger.pdf>
60. McCorry, P., Hicks, A., Meiklejohn, S.: Smart Contracts for Bribing Miners. In: 5th Workshop on Bitcoin and Blockchain Research, Financial Cryptography and Data Security 18 (FC). Springer (2018), <http://fc18.ifca.ai/bitcoin/papers/bitcoin18-final14.pdf>
61. Meissner, R., Gnosis community: Gnosis safe contracts - executor. <https://github.com/gnosis/safe-contracts/blob/34c87b783dfd04ff09ef7c358c3182c3c151e086/contracts/base/Executor.sol>, accessed: 2021-09-07
62. Moroz, D.J., Aronoff, D.J., Narula, N., Parkes, D.C.: Double-spend counterattacks: Threat of retaliation in proof-of-work systems. arXiv preprint arXiv:2002.10736 (2020), <https://arxiv.org/pdf/2002.10736.pdf>
63. Nakamoto, S.: Bitcoin: A Peer-to-Peer Electronic Cash System (Dec 2008), <https://bitcoin.org/bitcoin.pdf>
64. Natoli, C., Gramoli, V.: The blockchain anomaly. In: Network Computing and Applications (NCA), 2016 IEEE 15th International Symposium on. pp. 310–317. IEEE (2016), <https://arxiv.org/pdf/1605.05438.pdf>
65. Nayak, K., Kumar, S., Miller, A., Shi, E.: Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. In: 1st IEEE European Symposium on Security and Privacy, 2016. IEEE (2016), <http://eprint.iacr.org/2015/796.pdf>
66. Neu, J., Tas, E.N., Tse, D.: Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 446–465. IEEE (2021)

67. Pass, R., Seeman, L., Shelat, A.: Analysis of the blockchain protocol in asynchronous networks. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 643–673. Springer (2017), <https://eprint.iacr.org/2016/454.pdf>
68. Pass, R., Shi, E.: Thunderella: Blockchains with Optimistic Instant Confirmation (2017), <http://eprint.iacr.org/2017/913.pdf>, published: Cryptology ePrint Archive, Report 2017/913
69. Pass, R., Shi, E.: Thunderella: Blockchains with optimistic instant confirmation. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 3–33. Springer (2018)
70. Rodler, M., Li, W., Karame, G.O., Davi, L.: Sereum: Protecting Existing Smart Contracts Against Re-Entrancy Attacks (2018), <https://arxiv.org/pdf/1812.05934.pdf>, published: arXiv:1812.05934
71. Rosenfeld, M.: Analysis of Hashrate-Based Double Spending, vol. abs/1402.2009 (2014), <https://arxiv.org/pdf/1402.2009.pdf>, publication Title: CoRR
72. Schneider, F.B.: Implementing fault-tolerant services using the state machine approach: A tutorial. In: ACM Computing Surveys (CSUR). vol. 22, pp. 299–319. ACM (1990), <http://www-users.cselabs.umn.edu/classes/Spring-2014/csci8980-sds/Papers/ProcessReplication/p299-schneider.pdf>, issue: 4
73. Schwarz-Schilling, C., Neu, J., Monnot, B., Asgaonkar, A., Tas, E.N., Tse, D.: Three attacks on proof-of-stake ethereum. In: International Conference on Financial Cryptography and Data Security (2022)
74. Sergey, I., Hobor, A.: A Concurrent Perspective on Smart Contracts (2017), <https://arxiv.org/pdf/1702.05511.pdf>, publication Title: arXiv preprint arXiv:1702.05511
75. Sompolinsky, Y., Zohar, A.: Bitcoin’s Security Model Revisited. arXiv preprint arXiv:1605.09193 (2016), <http://arxiv.org/pdf/1605.09193.pdf>
76. Sonnino, A., Bano, S., Al-Bassam, M., Danezis, G.: Replay attacks and defenses against cross-shard consensus in sharded distributed ledgers. In: 2020 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 294–308. IEEE (2020)
77. Todd, P.: Op_checkclocktimeverify (Oct 2014), <https://github.com/bitcoin/bips/blob/master/bip-0065.mediawiki>
78. Tran, M., Choi, I., Moon, G.J., Vu, A.V., Kang, M.S.: A Stealthier Partitioning Attack against Bitcoin Peer-to-Peer Network. In: To appear in Proceedings of IEEE Symposium on Security and Privacy (IEEE S&P) (2020), <https://erebus-attack.comp.nus.edu.sg/erebus-attack.pdf>
79. Victor, F., Lüders, B.K.: Measuring ethereum-based ERC20 token networks. In: International Conference on Financial Cryptography and Data Security. pp. 113–129. Springer (2019)
80. Wang, X., He, J., Xie, Z., Zhao, G., Cheung, S.C.: Contractguard: Defend ethereum smart contracts with embedded intrusion detection. *IEEE Transactions on Services Computing* **13**(2), 314–328 (2019)
81. Wohrer, M., Zdun, U.: Smart contracts: security patterns in the ethereum ecosystem and solidity. In: 2018 International Workshop on Blockchain Oriented Software Engineering (IWBOSE). pp. 2–8. IEEE (2018)
82. Wood, G., et al.: Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper* **151**(2014), 1–32 (2014)
83. Wu, L., Wu, S., Zhou, Y., Li, R., Wang, Z., Luo, X., Wang, C., Ren, K.: EthScope: A Transaction-centric Security Analytics Framework to Detect Malicious Smart Contracts on Ethereum. arXiv:2005.08278 [cs] (May 2020), <http://arxiv.org/abs/2005.08278>, arXiv: 2005.08278

84. Zamyatin, A., Al-Bassam, M., Zindros, D., Kokoris-Kogias, E., Moreno-Sanchez, P., Kiayias, A., Knottenbelt, W.J.: SoK: Communication Across Distributed Ledgers. IACR Cryptology ePrint Archive, 2019: 1128 (2019), <https://eprint.iacr.org/2019/1128.pdf>
85. Zhang, M., Zhang, X., Zhang, Y., Lin, Z.: {TXSPECTOR}: Uncovering attacks in ethereum from transactions. In: 29th {USENIX} Security Symposium ({USENIX} Security 20). pp. 2775–2792 (2020)
86. Zhang, R., Preneel, B.: Lay down the common metrics: Evaluating proof-of-work consensus protocols’ security. In: 2019 IEEE Symposium on Security and Privacy (SP). IEEE (2019), <https://www.esat.kuleuven.be/cosic/publications/article-3005.pdf>
87. Zhang, Y., Setty, S., Chen, Q., Zhou, L., Alvisi, L.: Byzantine ordered consensus without byzantine oligarchy. In: 14th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 20). pp. 633–649 (2020)
88. Zhou, L., Qin, K., Gervais, A.: A2mm: Mitigating frontrunning, transaction re-ordering and consensus instability in decentralized exchanges. arXiv preprint arXiv:2106.07371 (2021)
89. Zhou, L., Qin, K., Torres, C.F., Le, D.V., Gervais, A.: High-frequency trading on decentralized on-chain exchanges. In: 2021 IEEE Symposium on Security and Privacy (SP). pp. 428–445. IEEE (2021)

A Responsible Disclosure and Ethical Considerations

After careful consideration and review of the potential impact of this work, we believe that the presentation of opportunistic algorithmic double-spending and *OpAl* attacks does not require prior notification or responsible disclosure. As we outline within this paper, the presented double-spending techniques do not fundamentally increase the success probability of attacks aimed at causing blockchain forks, which are a necessary prerequisite for performing most double-spending attacks. Hence, exchanges, merchants, and users that adhere to best practices and wait for sufficiently many confirmations on transactions do not appear to face substantially higher risk compared to other forms of double-spending.

However, we do see a potential for systemic risks in case of severe technical failures or attacks that cause deep forks, because of blockchain interlinking [84,76], layer-two protocols [32] and the existence of large centralized cryptocurrency exchanges, which all face the risk of becoming victims of algorithmic double-spending during such events. The principal mechanisms that can lead to *unintentional* algorithmic double-spending are already present in various smart contracts today. Consider the *slippage tolerance* users can specify as part of a trade when interacting with DEXs such as Uniswap [1]. In case of a deep fork, the necessary preconditions for a user’s trade to be executed within a transaction may change, e.g., due to transaction reordering. The effect can be similar or even equivalent to intentional *OpAl* attacks.

We are hence convinced that raising awareness of these potential issues through openly publishing our findings is the best course of action.

B Definition of Semantic Malleability

Definition 5 (semantic malleability). *A transaction t is semantically malleable if at least one of the following conditions hold:*

1. *Executing t on two states σ_a, σ_b results in different changes to their respective substates, i.e., $\text{DIFF}(t, \sigma_a) \neq \text{DIFF}(t, \sigma_b)$.*

$$\exists \sigma_a, \sigma_b (\text{VALID}(t, \sigma_a) \wedge \text{VALID}(t, \sigma_b) \wedge \text{DIFF}(t, \sigma_a) \neq \text{DIFF}(t, \sigma_b)) \quad (3)$$

2. *After a transaction t becomes valid for execution in state σ_a , t is invalidated at some future state σ_b without being executed.*

$$\exists \sigma_a, \sigma_b (\text{VALID}(t, \sigma_a) \wedge \sigma_a \prec \sigma_b \wedge t \notin \text{PAST}(\sigma_b) \wedge \neg \text{VALID}(t, \sigma_b)) \quad (4)$$

C Additional Data on *OpAl*-like Transactions and Forks

Contract Address	TX int.	Purpose	Name	Source	Opcode Purpose
0xc5F85281d4402850ff436b959a925a0e811D78d3	557	Game/Token	CnMGame	yes	randomness?
0x000000000035B5e5ad9019092c665357240f594e	411	MEV Bot?	?	no	context sensitivity?
0xEef86c2E49E11345F1a693675dF9a38f7d880CBF	313	MEV Bot?	?	no	context sensitivity?
0x5E4e65926BA27467555EB562121fac00D24E9dD2	264	Layer 2 rollup	optimism.io	yes	caching/processing
0x56a76bc92361f6DF8D7547f6eD8843EdC70e1C9	227	Layer 2 rollup	metis.io	yes	caching/processing
0xB6eD7644C69416d67B522e20bc294A9a9B405B31	222	Token	0xbitcoin.org	yes	context sensitivity
0xd6e382aa7A09fc4A09C2fb99Cfce6A429985E65d	221	Game/Token	Doomsday NFT (BUNKER)	yes	randomness
0x75E9Abc7E69fc46177d2F3538C0B92d89054eC91	130	Token/NFT	EnterDAO Sharded Minds	yes	randomness
0x563bDabAa8846ec445b25Bfbcd88d160890a02ED	115	MEV Bot?	?	no	context sensitivity?
0xa10FcA31A2Cb432C9Ac976779DC947CfDb003EF0	111	MEV Bot?	?	no	context sensitivity?

Table 2. Contracts with the highest number of transaction interactions with EVM opcodes that are also characteristic of *OpAl*. (?) denotes uncertain categorizations.

Currency	Date	Start blockheight	Length	Duration (hours)
BTC	2010-08-15	74638	51	8.5
BTC	2013-03-11	225430	24	4
BTC	2015-07-04	363731	6	1
ETH	2016-11-24	2686351	≈ 165	1.25
ETH	2020-11-11	11234873	> 30	?
ETH 2.0	2022-05-25	3887075	7	0.02

Table 3. Long range forks in Bitcoin and Ethereum due to technical failures.

D Alternative OpAl Attack Designs in Ethereum

We hereby illustrate alternative constructions to the *OpAl* attack that is outlined in Section 6. Figure 2 shows two different attack scenarios that rely on depth-1 fork oracles. Hereby the first part of the illustration (left) captures an attack that relies on querying the *coinbase address* of the current block and conditionally branches if a particular mining pool address is returned. The second part of the illustration (right) shows a scenario where the block height is used to trigger the attack. Figure 3 serves to illustrate how *OpAl* attacks can also be facilitated through *invalidation* of transactions, for example if the underlying protocol supports *context sensitive transactions*. Hereby, the transaction may only be valid for specific blockchain states, or during specific time intervals, at the protocol level, preventing the transaction, such as a payment, to be replayed in a fork.

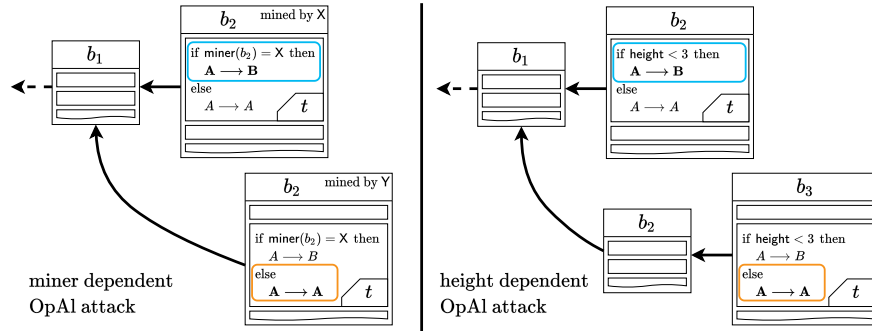


Fig. 2. Depth-1 *OpAl* attacks




Fig. 3. *OpAl* attack based on transaction invalidation