

# Smart Contract Synthesis Modulo Hyperproperties

Norine Coenen

CISPA Helmholtz Center for Information Security  
Saarbrücken, Germany  
norine.coenen@cispa.de

Bernd Finkbeiner

CISPA Helmholtz Center for Information Security  
Saarbrücken, Germany  
finkbeiner@cispa.de

Jana Hofmann

CISPA Helmholtz Center for Information Security  
Saarbrücken, Germany  
jana.hofmann@cispa.de

Julia Tillman

CISPA Helmholtz Center for Information Security  
Saarbrücken, Germany  
julia.tillman@cispa.de

**Abstract**—Smart contracts are small but highly security-critical programs that implement wallets, token systems, auctions, crowd funding systems, elections, and other multi-party transactions on the blockchain. A broad range of methods has been developed to ensure that a smart contract is functionally correct. However, smart contracts often additionally need to satisfy certain hyperproperties, such as symmetry, determinism, or an information flow policy. In this paper, we show how a synthesis method for smart contracts can ensure that the contract satisfies its desired hyperproperties. We build on top of a recently developed synthesis approach from specifications in the temporal logic TSL. We present HyperTSL, an extension of TSL for the specification of hyperproperties of infinite-state software. As a preprocessing step, we show how to detect if a hyperproperty has an equivalent formulation as a (simpler) trace property. Finally, we describe how to refine a synthesized contract to adhere to its HyperTSL specification.

**Index Terms**—hyperproperties, smart contracts, reactive synthesis, temporal logics

## I. INTRODUCTION

Smart contracts are small programs which describe transactions between multiple parties. Common smart contracts implement wallets, token systems, auctions, crowd funding systems, and elections. Smart contracts run on the blockchain and thus remove the need for a trusted third party. Bugs in smart contracts have led to huge monetary losses in the last decade, making the need for rigorous formal foundations for smart contracts obvious. In the absence of a trusted third party, it is especially important to ensure that a contract implements the agreed order of transactions. The large body of existing work on the correctness of a contract’s control flow mainly consists of verification methods to ensure that a contract exhibits a correct behavior with respect to *trace properties*.

Smart contracts often additionally need to satisfy *hyperproperties*, however. Hyperproperties describe the relation between

several execution traces of the system. Prominent hyperproperties are information flow policies like noninterference and observational determinism. Existing work on hyperproperties in smart contracts focuses on verifying concrete information flow policies such as integrity, which are enforced using language-based methods [1], [2]. Hyperproperties are not limited to information flow policies, though. To establish the users’ trust in a contract, one has to additionally show that a contract satisfies hyperproperties such as robustness and symmetry. In a voting contract, for example, all candidates should be treated symmetrically, and a vote from the owner of the contract should not count more than any other vote.

In this work, we develop a general-purpose temporal logic, HyperTSL, to express the wide range of hyperproperties relevant in the context of smart contracts. We then show how to automatically construct the underlying state machine of a smart contract directly from the logical specification of the hyperproperties. We approach the task from two angles. First, we investigate the synthesis problem of the logic and show how to approximate the in general undecidable problem to obtain correct-by-design implementations. Second, we propose a two-step approach to include HyperTSL properties in a recently proposed smart contract synthesis approach [3].

**HyperTSL.** We develop two variants of the logic, HyperTSL and HyperTSL<sub>rel</sub>, which are both based on temporal stream logic (TSL) [4]. TSL extends linear temporal logic (LTL) with the concept of cells and uninterpreted functions and predicates. These mechanisms separate the temporal control flow from the concrete data in the system, which enables reasoning about infinite-state systems. TSL has been applied successfully to synthesize hardware [5], functional programs [6], and also to describe the underlying state machines of smart contracts [3]. As an example, we can state that the winner of an election remains the same after the method `close` has been called.

$$\Box(\text{close} \rightarrow \bigcirc \Box[\text{winner} \leftarrow \text{winner}])$$

The contract’s method `close` is modeled as a boolean and the field `winner` is modeled as a cell. We define HyperTSL as a conservative extension of TSL. It adds quantification over

This work was partially supported by the German Research Foundation (DFG) as part of the Collaborative Research Center “Foundations of Perspicuous Software Systems” (TRR 248, 389792660), and by the European Research Council (ERC) Grant OSARES (No. 683300). N. Coenen, J. Hofmann, and J. Tillman carried out this work as members of the Saarbrücken Graduate School of Computer Science.

multiple executions; predicates, functions, and updates then refer to one of the quantified executions. In an election with two candidates A and B, we can state that two traces have the same winner A if they agree on the votes for A.

$$\forall\pi\forall\pi'.\Box(\text{vote}_{A,\pi} \leftrightarrow \text{vote}_{A,\pi'}) \\ \rightarrow \Box(\llbracket \text{winner} \leftarrow A() \rrbracket_\pi \leftrightarrow \llbracket \text{winner} \leftarrow A() \rrbracket_{\pi'})$$

The second hyperlogic  $\text{HyperTSL}_{\text{rel}}$  allows to relate executions within a predicate. Here, we can state that the winners on two executions are always the same as long as the votes are always the same.

$$\forall\pi\forall\pi'.\Box(\text{vote}_\pi = \text{vote}_{\pi'}) \rightarrow \Box(\text{winner}_\pi = \text{winner}_{\pi'})$$

Note that the  $=$  predicate ranges over two different execution variables  $\pi$  and  $\pi'$ . In TSL, all functions and predicates are uninterpreted, but  $=$  would probably be implemented as actual equality.

**Smart contract synthesis from HyperTSL.** Synthesis constructs correct-by-design systems from a formal specification. It therefore shifts the development from software writing to specification writing. This idea fits especially well to smart contracts. Most contracts belong to one of the major classes like token systems or election protocols, and their behavior can be characterized with similar specifications. This also extends to hyperproperties: the set of actually relevant hyperproperties is often the same across different systems.

We therefore investigate the synthesis problem of HyperTSL. Inherited from TSL, the problem is undecidable already for a single quantifier. We show, however, that the  $\forall^*$  fragment of HyperTSL can be approximated in  $\forall^*$  HyperLTL, for which there exists a bounded synthesis approach implemented in HyperBoSy [7]. For the  $\exists^*$  fragment of HyperTSL, we present an approximation based on LTL satisfiability checking.

As a step towards actual smart contract synthesis, we build on a recent workflow that implements the synthesis of a smart contract control flow from TSL specifications [3]. We present a two-step approach to synthesize a contract that adheres to HyperTSL specifications. First, we check if the combination of a given TSL specification with a  $\forall^*$  HyperTSL specification lets the hyperproperty “collapse” to a simple trace property. We show that the check is undecidable in general, but can to some extent be approximated in a decidable fragment of with HyperLTL satisfiability checking. If this is not the case, we synthesize the most general solution (the winning region) for the TSL specification using the tool presented in [3]. We then prune the system to find a strategy that implements a  $\forall^*$  HyperTSL property. Our implementation of the approach shows that we can automatically construct a voting contract which satisfies several hyperproperties.

**Contributions.** We present the first solution to automatically construct smart contracts with respect to hyperproperties.

- We define HyperTSL and  $\text{HyperTSL}_{\text{rel}}$ , two hyperlogics based on TSL, and show how they can specify hyperproperties occurring in the context of smart contracts.

- We show that the synthesis problem of  $\forall^*$  HyperTSL can be over-approximated by  $\forall^*$  HyperLTL synthesis and that the synthesis problem of  $\exists^*$  HyperTSL can be under-approximated by LTL satisfiability checking.
- We describe a preprocessing step to detect if a HyperTSL formula has an equivalent TSL formulation.
- We describe and implement a repair-like pruning algorithm that refines a smart contract synthesized from a TSL specification to satisfy a  $\forall^*$  HyperTSL formula.

**Structure of the paper.** We first give the necessary preliminaries, including the definitions of HyperLTL and TSL. In Section III, we present HyperTSL and  $\text{HyperTSL}_{\text{rel}}$ , and show how to specify hyperproperties in a voting contract. We proceed by investigating the realizability problem of HyperTSL and argue why HyperTSL is more suited for synthesis than  $\text{HyperTSL}_{\text{rel}}$ . In Section V, we recap the synthesis workflow based on TSL [3] and discuss the combination of the TSL specification of the voting contract with the HyperTSL specifications. Section VI discusses the detection of pseudo hyperproperties. Section VII presents the pruning of a contract with respect to HyperTSL, including a prototype implementation.

## II. PRELIMINARIES

We recap the definitions of HyperLTL and TSL.

### A. HyperLTL

Let  $AP$  be an infinite set of atomic propositions. A trace  $t$  is an element of  $(2^{AP})^\omega$ . We write  $t[i]$  for the  $i$ th element of  $t$ . A trace property  $P$  is a set of traces, a hyperproperty  $H$  is a set of sets of traces. HyperLTL [8] describes hyperproperties by extending LTL with explicit trace quantification. It is defined with respect to an infinite set of trace variables  $\mathbb{V}$ .

$$\varphi ::= \forall\pi. \varphi \mid \exists\pi. \varphi \mid \psi \\ \psi ::= a_\pi \mid \neg\psi \mid \psi \wedge \psi \mid \bigcirc\psi \mid \psi \mathcal{U} \psi,$$

where  $a \in AP$  and  $\pi \in \mathbb{V}$ . In HyperLTL, atomic propositions are indexed with the trace variable they refer to. The semantics of HyperLTL is defined with respect to a set of traces  $T$ . Let  $\Pi : \mathbb{V} \rightarrow T$  be a trace assignment that maps trace variables to traces in  $T$ . To update  $\Pi$ , we write  $\Pi[\pi \mapsto t]$ , where  $\pi$  maps to  $t$  and all other trace variables are as in  $\Pi$ . The satisfaction relation of HyperLTL is defined with respect to a trace assignment  $\Pi$ , a set of traces  $T$ , and a point in time  $i$ .

$$\begin{aligned} \Pi, T, i \models a_\pi & \quad \text{iff } a \in \Pi(\pi)[i] \\ \Pi, T, i \models \neg\psi & \quad \text{iff } \Pi, T, i \not\models \psi \\ \Pi, T, i \models \psi_1 \wedge \psi_2 & \quad \text{iff } \Pi, T, i \models \psi_1 \text{ and } \Pi, T, i \models \psi_2 \\ \Pi, T, i \models \bigcirc\psi & \quad \text{iff } \Pi, T, i+1 \models \psi \\ \Pi, T, i \models \psi_1 \mathcal{U} \psi_2 & \quad \text{iff } \exists j \geq i. \Pi, T, j \models \psi_2 \\ & \quad \text{and } \forall i \leq k < j. \Pi, T, k \models \psi_1 \\ \Pi, T, i \models \exists\pi. \varphi & \quad \text{iff } \exists t \in T. \Pi[\pi \mapsto t], T, i \models \varphi \\ \Pi, T, i \models \forall\pi. \varphi & \quad \text{iff } \forall t \in T. \Pi[\pi \mapsto t], T, i \models \varphi \end{aligned}$$

As for LTL, we can derive the usual boolean logical connectives and temporal operators “eventually”  $\Diamond \psi = \top \mathcal{U} \psi$ , “globally”  $\Box \psi = \neg \Diamond \neg \psi$ , and “weak until”  $\psi_1 \mathcal{W} \psi_2 = (\psi_1 \mathcal{U} \psi_2) \vee \Box \psi_1$ . When specifying smart contracts, we assume specifications to belong to the *syntactic safety* fragment, i.e., that they only contain  $\bigcirc, \Box$ , and  $\mathcal{W}$  as temporal operators. Formulas from the syntactic safety fragment can be equivalently expressed using *past-time temporal operators*, which reason about the past instead of the future of a trace. “Yesterday”  $\ominus$  is the past-time equivalent of  $\bigcirc$ , and “since”  $\mathcal{S}$  corresponds to  $\mathcal{U}$ . We say that a trace set  $T$  satisfies a HyperLTL formula  $\varphi$ , written as  $T \models \varphi$ , if  $\emptyset, T, 0 \models \varphi$ , where  $\emptyset$  denotes the empty trace assignment.

We define the HyperLTL realizability problem as in [9]. We assume  $AP$  to be a disjoint union of a set of input propositions and a set of output propositions, i.e.,  $AP = AP_{in} \dot{\cup} AP_{out}$ . A *strategy*  $\sigma : (2^{AP_{in}})^+ \rightarrow 2^{AP_{out}}$  maps finite sequences of input valuations to an output valuation. For an infinite word  $w = w_0 w_1 w_2 \dots \in (2^{AP_{in}})^\omega$ , the trace generated by a strategy  $\sigma$  is defined as  $\sigma_w = (w_0 \cup \sigma(w_0))(w_1 \cup \sigma(w_0 w_1)) \dots \in (2^{AP})^\omega$ . We write  $traces(\sigma)$  for the set  $\{\sigma_w \mid w \in (2^{AP_{in}})^\omega\}$ . We say that a strategy  $\sigma$  satisfies a HyperLTL formula  $\varphi$  over  $AP = AP_{in} \dot{\cup} AP_{out}$  if  $traces(\sigma) \models \varphi$ .

**Definition II.1** (HyperLTL realizability). A HyperLTL formula  $\varphi$  over  $AP = AP_{in} \dot{\cup} AP_{out}$  is realizable if there is a strategy  $\sigma : (2^{AP_{in}})^+ \rightarrow 2^{AP_{out}}$  that satisfies  $\varphi$ .

## B. TSL

We introduce the syntax and semantics of TSL following [3], [4]. TSL extends LTL with the concept of cells, which hold data from an arbitrary domain. To abstract from concrete data points, TSL employs uninterpreted functions and predicates. A TSL formula describes a system that receives a stream of inputs, abstracts from the concrete values using predicates, and produces a stream of cell updates using function applications. We formally define values, functions, and predicates. The set of all values is denoted by  $\mathcal{V}$ , the Boolean values by  $\mathbb{B} \subseteq \mathcal{V}$ . An  $n$ -ary function  $f : \mathcal{V}^n \rightarrow \mathcal{V}$  computes a value from  $n$  values. An  $n$ -ary predicate  $p : \mathcal{V}^n \rightarrow \mathbb{B}$  assigns a boolean value to  $n$  values. The sets of all functions and predicates are denoted by  $\mathcal{F}$  and  $\mathcal{P} \subseteq \mathcal{F}$ , respectively. Constants are both 0-ary functions and values.

Let  $\mathbb{I}$  and  $\mathbb{C}$  be the set of inputs and cells, and let  $\Sigma_F$  and  $\Sigma_P \subseteq \Sigma_F$  be the set of function and predicate symbols. *Function terms*  $\tau^f$  are recursively defined by

$$\tau^f ::= s \mid \mathfrak{f} \tau^f \dots \tau^f$$

where  $s \in \mathbb{I} \cup \mathbb{C}$  and  $\mathfrak{f}$  is an  $n$ -ary function symbol. *Predicate terms*  $\tau^p$  are obtained by applying a predicate to function terms. The sets of all function and predicate terms are denoted by  $\mathcal{T}_F$  and  $\mathcal{T}_P \subseteq \mathcal{T}_F$ , respectively. TSL formulas are built according to the following grammar:

$$\psi ::= \neg \psi \mid \psi \wedge \psi \mid \bigcirc \psi \mid \psi \mathcal{U} \psi \mid \tau^p \mid \llbracket c \leftarrow \tau^f \rrbracket$$

where  $c \in \mathbb{C}$ ,  $\tau^p \in \mathcal{T}_P$ , and  $\tau^f \in \mathcal{T}_F$ . An *update term*  $\llbracket c \leftarrow \tau^f \rrbracket$  denotes that the value of function term  $\tau^f$  is assigned to cell  $c$ . We denote the set of update terms with  $\mathcal{T}_U$ .

Function and predicate terms are syntactic objects, i.e. the term  $p(\mathfrak{f} c)$  only becomes meaningful when we assign, for example,  $c$  to the value 3,  $\mathfrak{f}$  to the function  $f(x) = x + 1$  and  $p$  to  $p(x) = x > 0$ . To assign a semantic interpretation to function and predicate symbols, we use an *interpretation*  $\langle \cdot \rangle : \Sigma_F \rightarrow \mathcal{F}$ . A cell assignment  $\mathbb{C} \rightarrow \mathcal{T}_F$  is a total function assigning exactly one function term to each cell. The set of all assignments  $\mathbb{C} \rightarrow \mathcal{T}_F$  is denoted by  $\mathcal{C}$ . A *computation*  $\varsigma \in \mathcal{C}^\omega$  describes the control flow of the cells, i.e., which function term is associated with a cell at which point in time. For every cell  $c \in \mathbb{C}$ , let  $init_c$  be a designated value assigned initially to  $c$ . Input streams  $\mathcal{I}^\omega$  are infinite sequences of assignments of inputs to values. Given a computation, an input stream, and a point in time  $i$ , we can *evaluate* a function term. The evaluation function  $\eta_{\langle \cdot \rangle} : \mathcal{C}^\omega \times \mathcal{I}^\omega \times \mathbb{N} \times \mathcal{T}_F \rightarrow \mathcal{V}$  is defined as

$$\eta_{\langle \cdot \rangle}(\varsigma, \iota, i, s) := \begin{cases} \iota[i] s & s \in \mathbb{I} \\ init_s & s \in \mathbb{C} \wedge i = 0 \\ \eta_{\langle \cdot \rangle}(\varsigma, \iota, i-1, \varsigma[i-1] s) & s \in \mathbb{C} \wedge i > 0 \end{cases}$$

$$\eta_{\langle \cdot \rangle}(\varsigma, \iota, i, \mathfrak{f} \tau_0 \dots \tau_n) := \langle \mathfrak{f} \rangle \eta_{\langle \cdot \rangle}(\varsigma, \iota, i, \tau_0) \dots \eta_{\langle \cdot \rangle}(\varsigma, \iota, i, \tau_n)$$

Note that  $\iota[i] s$  denotes the value that input stream  $\iota$  assigns to input  $s$  at position  $i$ . Likewise,  $\varsigma[i] s$  is the function term that  $\varsigma$  assigns to cell  $s$  at point in time  $i$ . As an example, to compute the value of cell  $x$  in step  $i$ , we might obtain from the computation that  $x$  is updated to  $\mathfrak{f} x$  in step  $i$ , so we recursively evaluate  $x$  in step  $i-1$  and apply the function assigned to  $\mathfrak{f}$  to the result. We evaluate a TSL formula  $\psi$  with respect to an assignment function  $\langle \cdot \rangle$ , an input stream  $\iota \in \mathcal{I}^\omega$ , a computation  $\varsigma \in \mathcal{C}^\omega$ , and a time step  $i \in \mathbb{N}$ .

$$\begin{aligned} \varsigma, \iota, i \models_{\langle \cdot \rangle} \neg \psi & \quad \text{iff} \quad \varsigma, \iota, i \not\models_{\langle \cdot \rangle} \psi \\ \varsigma, \iota, i \models_{\langle \cdot \rangle} \psi_1 \wedge \psi_2 & \quad \text{iff} \quad \varsigma, \iota, i \models_{\langle \cdot \rangle} \psi_1 \text{ and } \varsigma, \iota, i \models_{\langle \cdot \rangle} \psi_2 \\ \varsigma, \iota, i \models_{\langle \cdot \rangle} \bigcirc \psi & \quad \text{iff} \quad \varsigma, \iota, i+1 \models_{\langle \cdot \rangle} \psi \\ \varsigma, \iota, i \models_{\langle \cdot \rangle} \psi_1 \mathcal{U} \psi_2 & \quad \text{iff} \quad \exists j \geq i. \varsigma, \iota, j \models_{\langle \cdot \rangle} \psi_2 \\ & \quad \text{and } \forall i \leq k < j. \varsigma, \iota, k \models_{\langle \cdot \rangle} \psi_1 \\ \varsigma, \iota, i \models_{\langle \cdot \rangle} \llbracket c \leftarrow \tau \rrbracket & \quad \text{iff} \quad \varsigma[i] c \equiv \tau \\ \varsigma, \iota, i \models_{\langle \cdot \rangle} p \tau_1 \dots \tau_n & \quad \text{iff} \quad \eta_{\langle \cdot \rangle}(\varsigma, \iota, i, p \tau_1 \dots \tau_n) \end{aligned}$$

We use  $\equiv$  to syntactically compare two terms. An *execution*  $(\varsigma, \iota)$  satisfies a TSL formula  $\psi$  if  $\varsigma, \iota, 0 \models_{\langle \cdot \rangle} \psi$  holds. In this case, we write  $\varsigma, \iota \models_{\langle \cdot \rangle} \psi$ . A TSL formula  $\psi$  is *satisfiable* iff there exists an interpretation  $\langle \cdot \rangle$  and an execution  $(\varsigma, \iota)$  such that  $\varsigma, \iota \models_{\langle \cdot \rangle} \psi$  [10].

**TSL realizability.** The realizability problem of a TSL formula  $\psi$  asks whether there exists a strategy which reacts to predicate evaluations with cell updates according to  $\psi$ . Formally, a strategy is a function  $\sigma : (2^{\mathcal{T}_P})^+ \rightarrow \mathcal{C}$ . Given a strategy  $\sigma$  and an input stream  $\iota$ , we can define the resulting computation

$\sigma(\iota)$ . To compute the cell updates at point in time  $i$ , we require the current inputs as well as the history of cell updates:

$$\begin{aligned} \sigma(\iota)[i] := & \sigma(\{\tau^p \in \mathcal{T}_P \mid \eta_{\langle \cdot \rangle}(\sigma(\iota), \iota, 0, \tau^p)\} \cdot \dots \\ & \dots \cdot \{\tau^p \in \mathcal{T}_P \mid \eta_{\langle \cdot \rangle}(\sigma(\iota), \iota, i, \tau^p)\}) \end{aligned}$$

Note that in order to define  $\sigma(\iota)[i]$ , the definition uses  $\sigma(\iota)$ . This is well-defined since the evaluation function  $\eta_{\langle \cdot \rangle}(\varsigma, \iota, i, \tau)$  only uses  $\varsigma[0] \dots \varsigma[i-1]$ .

**Definition II.2** (TSL realizability [4]). A TSL formula  $\psi$  is *realizable* iff there exists a strategy  $\sigma : (2^{\mathcal{T}_P})^+ \rightarrow \mathcal{C}$  such that for every input stream  $\iota \in \mathcal{I}^\omega$  and every assignment function  $\langle \cdot \rangle : \Sigma_P \rightarrow \mathcal{F}$ , it holds that  $\sigma(\iota), \iota \models_{\langle \cdot \rangle} \psi$ .

### III. HYPERTSL

We extend TSL with execution quantifiers to allow reasoning about hyperproperties for infinite-state systems. There is not a single obvious way to lift TSL to a hyperlogic, so we motivate in the following our choices for HyperTSL.

**Function and predicate interpretations.** There are two options to deal with function and predicate interpretations. The first is to interpret functions and predicates differently across multiple executions. This would make sense, for example, to model multiple components of a system. Most of the hyperproperties, however, concern the behavior of the same system when executed several times. Functions and predicates like equality and addition mostly stay the same across several executions. We therefore decide to quantify the function and predicate interpretation *before* quantifying executions.

**Domain of updates and predicates.** Another question is whether predicates, functions, and update terms may contain terms evaluated on different executions. In the case of update terms the answer should probably be no: an update of the form  $\llbracket c_\pi \leftarrow f(c_{\pi'}) \rrbracket$  would mean that the system could access the value of a cell at the same point in time but on a different execution. This does not seem realistic for a real system, which would have to store all possible values a cell  $c$  could hold at some point in time. Additionally, the execution chosen for  $\pi'$  differs with every quantifier evaluation. Predicates ranging over multiple executions, on the other hand, would be useful. One could state a stronger version of observational determinism, namely “if two traces agree on the value of input  $x$ , then also cell  $c$  is always updated to the same term”, written as  $\Box(x_\pi = x_{\pi'}) \rightarrow \Box(\llbracket c \leftarrow f(x) \rrbracket_\pi \leftrightarrow \llbracket c \leftarrow f(x) \rrbracket_{\pi'})$ . We therefore define two logics: HyperTSL, in which we allow predicates to range only over the *same* execution, and HyperTSL<sub>rel</sub>, which relates multiple executions using predicates.

#### A. HyperTSL

We define HyperTSL as an extension of TSL. The definition of function terms  $\tau^f$  and predicate terms  $\tau^p$  are the same as in TSL. The syntax of HyperTSL is given as follows.

$$\begin{aligned} \varphi ::= & \forall \pi. \varphi \mid \exists \pi. \varphi \mid \psi \\ \psi ::= & \neg \psi \mid \psi \wedge \psi \mid \psi \circ \psi \mid \psi \mathcal{U} \psi \mid \tau_\pi^p \mid \llbracket c \leftarrow \tau^f \rrbracket_\pi \end{aligned}$$

where  $\pi \in \mathbb{V}$ ,  $c \in \mathbb{C}$ ,  $\tau^p \in \mathcal{T}_P$ , and  $\tau^f \in \mathcal{T}_F$ . We evaluate a HyperTSL formula  $\varphi$  with respect to a set of executions  $E \subseteq \mathcal{C}^\omega \times \mathcal{I}^\omega$ , an execution assignment function  $\Pi$ , and a point in time  $i$ . We omit the cases for  $\neg$ ,  $\wedge$ ,  $\circ$ , and  $\mathcal{U}$ , which follow closely the respective cases in TSL.

$$\begin{aligned} \Pi, E, i \models_{\langle \cdot \rangle} \llbracket c \leftarrow \tau^f \rrbracket_\pi & \text{ iff } \#_\varsigma(\Pi(\pi)) \ i \ c \equiv \tau^f \\ \Pi, E, i \models_{\langle \cdot \rangle} (\mathfrak{p} \ \tau_1 \dots \tau_n)_\pi & \text{ iff } \eta_{\langle \cdot \rangle}(\varsigma, \iota, i, \mathfrak{p} \ \tau_1 \dots \tau_n) \\ & \text{ for } \varsigma = \#_\varsigma(\Pi(\pi)) \text{ and} \\ & \iota = \#_\iota(\Pi(\pi)) \end{aligned}$$

$$\begin{aligned} \Pi, E, i \models_{\langle \cdot \rangle} \exists \pi. \varphi & \text{ iff } \exists e \in E. \Pi[\pi \mapsto e], E, i \models_{\langle \cdot \rangle} \varphi \\ \Pi, E, i \models_{\langle \cdot \rangle} \forall \pi. \varphi & \text{ iff } \forall e \in E. \Pi[\pi \mapsto e], E, i \models_{\langle \cdot \rangle} \varphi \end{aligned}$$

Here,  $\#_\iota(\cdot)$  and  $\#_\varsigma(\cdot)$  are used to project an execution on its input stream and its computation. Given an assignment function  $\langle \cdot \rangle : \Sigma_F \rightarrow \mathcal{F}$ , a set of executions  $E$  satisfies a HyperTSL formula  $\varphi$ , written  $E \models_{\langle \cdot \rangle} \varphi$ , if  $\emptyset, E, 0 \models_{\langle \cdot \rangle} \varphi$ .

#### B. HyperTSL<sub>rel</sub>

In HyperTSL<sub>rel</sub>, predicate terms are defined with respect to the set of executions variables  $\mathbb{V}$ . If  $\tau^1, \dots, \tau^n$  are function terms,  $\pi_1, \dots, \pi_n \in \mathbb{V}$ , and  $\mathfrak{p}$  is an  $n$ -ary predicate symbol, then  $\mathfrak{p} \ \tau_{\pi_1}^1 \dots \tau_{\pi_n}^n$  is a predicate term. The syntax of HyperTSL<sub>rel</sub> is that of HyperTSL with the exception that we denote predicate terms as  $\tau_{\pi_1, \dots, \pi_n}^p$  instead of  $\tau_\pi^p$ . The semantics of HyperTSL<sub>rel</sub> is that of HyperTSL except for the semantics of the predicate, which we define as follows.

$$\begin{aligned} \Pi, E, i \models_{\langle \cdot \rangle} \mathfrak{p} \ \tau_{\pi_1}^1 \dots \tau_{\pi_n}^n & \text{ iff } \langle \mathfrak{p} \rangle \ v_1 \dots v_n \\ \text{where } v_j = & \eta_{\langle \cdot \rangle}(\#_\varsigma(\Pi(\pi_j)), \#_\iota(\Pi(\pi_j)), i, \tau^j) \text{ for } 1 \leq j \leq n \end{aligned}$$

Satisfaction of a HyperTSL<sub>rel</sub> formula is defined as for HyperTSL.

#### C. HyperTSL for Smart Contracts

As for TSL, the strength of HyperTSL and HyperTSL<sub>rel</sub> is in their use of uninterpreted predicates and functions together with a cell mechanism. That way, the logics can describe control flow properties like “if predicate  $\mathfrak{p}(c)$  evaluates to *true*, cell  $c$  should be updated by applying function  $f$  to it”. Since the concrete implementation of  $\mathfrak{p}$  and  $f$  is not relevant for the property, the synthesized system can be later completed to a concrete system by implementing  $\mathfrak{p}$  and  $f$  as needed.

In this section, we showcase how the logics can express temporal hyperproperties of smart contracts. As a running example, we use a simple voting contract. We informally describe the contract and specify hyperproperties that should hold in such a contract. This showcases the expressiveness of the logics and makes the differences in expressiveness become clear. We postpone the formal specification of the trace properties of the contract to Section V.

For simplicity, we describe the core features of an exemplary voting contract, which are enough to demonstrate typical hyperproperties occurring in a voting scenario. The contract describes an election where users can vote for either candidate A or candidate B by calling methods `voteA` and

`voteB`. The contract can be closed by the owner, who can call method `close`. `voteA`, `voteB`, and `close` are modeled as booleans. We assume that in each step, exactly one of `voteA`, `voteB`, and `close` holds. The contract has two fields, `votesA` and `votesB`, to store the number of votes recorded for the candidates. Both fields are modeled as cells in TSL. In the smart contract setting, we use the term field and cell interchangeably. Furthermore, `winner` holds the current winner chosen by the contract. The winner is updated in every step where the contract receives a vote, so either  $\llbracket \text{winner} \leftarrow A() \rrbracket$  or  $\llbracket \text{winner} \leftarrow B() \rrbracket$  holds globally.  $A()$  and  $B()$  are constants.

The hyperproperties relevant in such a setting are typical hyperproperties such as determinism and symmetry. Consider, for example, the question which candidate is chosen as winner if there is a tie. One would expect a form of symmetry, e.g., that the winner is not simply always candidate A. Similarly, one typically requires some sort of determinism, i.e., that in similar situations the same winner is chosen. We use the following syntactic sugar.

$$\begin{aligned} \text{sameWinner}(\pi, \pi') := & \\ & (\llbracket \text{winner} \leftarrow A() \rrbracket_\pi \leftrightarrow \llbracket \text{winner} \leftarrow A() \rrbracket_{\pi'}) \\ & \wedge (\llbracket \text{winner} \leftarrow B() \rrbracket_\pi \leftrightarrow \llbracket \text{winner} \leftarrow B() \rrbracket_{\pi'}) \end{aligned}$$

The formula states that two executions  $\pi$  and  $\pi'$  choose the same winner.

**Determinism.** We state that the winner of the election is determined by the sequence of votes received.

$$\forall \pi. \forall \pi'. \text{sameWinner}(\pi, \pi') \mathcal{W} \left[ (\text{voteA}_\pi \leftrightarrow \text{voteA}_{\pi'}) \vee (\text{voteB}_\pi \leftrightarrow \text{voteB}_{\pi'}) \right]$$

The formula states that for any two executions, as long as they receive the exact same votes, they should agree on who is the winner. Alternatively, we could have stated a local version of determinism using a greater predicate. Instead of writing greater `votesA votesB`, we use the better readable infix notation with  $>$ . Note that the predicate is uninterpreted at this stage although we intend it to be implemented as expected. We express that the choice of the winner is strictly determined by the current evaluation of the  $>$  predicate on the votes and the last vote.

$$\begin{aligned} \forall \pi. \forall \pi'. \square [ & \\ & ((\text{votesA} > \text{votesB})_\pi \leftrightarrow (\text{votesA} > \text{votesB})_{\pi'}) \\ & \wedge ((\text{votesB} > \text{votesA})_\pi \leftrightarrow (\text{votesB} > \text{votesA})_{\pi'}) \\ & \wedge (\text{voteA}_\pi \leftrightarrow \text{voteA}_{\pi'}) \wedge (\text{voteB}_\pi \leftrightarrow \text{voteB}_{\pi'}) \\ & \rightarrow \text{sameWinner}(\pi, \pi') ] \end{aligned} \quad (1)$$

This formula states that if A and B received the same number of votes, the winner must be the same on both executions. Note that the formula implicitly entails that any other predicate/input does not influence the winner. For example, if there is a predicate term `sender = owner()` testing whether the caller of the method is the contract's owner, this should not

influence the `winner` field. This formulation of determinism additionally entails that the past of an execution cannot influence the choice of the winner.

In  $\text{HyperTSL}_{\text{rel}}$ , we can express determinism by relating executions with predicates. Instead of abstracting from the concrete number of votes using the  $>$  predicate, we could state local determinism as follows. As before, the  $=$  predicate used in the formula is uninterpreted.

$$\begin{aligned} \forall \pi. \forall \pi'. & \\ \square [ & \text{votesA}_\pi = \text{votesA}_{\pi'} \wedge \text{votesB}_\pi = \text{votesB}_{\pi'} \\ & \rightarrow \text{sameWinner}(\pi, \pi') ] \end{aligned}$$

**Symmetry.** A prominent fairness condition in systems with multiple agents is *symmetry*, which requires that agents are treated symmetrically. In the voting case, this means, e.g., that if two traces swap the votes for A and B, then the winner must also be swapped.

$$\begin{aligned} \forall \pi. \forall \pi'. [ & (\llbracket \text{winner} \leftarrow A() \rrbracket_\pi \leftrightarrow \llbracket \text{winner} \leftarrow B() \rrbracket_{\pi'}) \\ & \wedge (\llbracket \text{winner} \leftarrow B() \rrbracket_\pi \leftrightarrow \llbracket \text{winner} \leftarrow A() \rrbracket_{\pi'}) \\ & \mathcal{W} [(\text{voteA}_\pi \leftrightarrow \text{voteB}_{\pi'}) \vee (\text{voteB}_\pi \leftrightarrow \text{voteA}_{\pi'})] ] \end{aligned} \quad (2)$$

The fairness property implies that in case of a tie, the winner cannot be brute-forced to always be candidate A or candidate B.

**No harm.** A typical monotonicity criterion in elections states that a vote for the currently leading candidate should not harm the candidate. Translated to our context this means that a vote for candidate A does not lead to B being the winner instead of A. We formulate this condition as follows in  $\text{HyperTSL}$ .

$$\begin{aligned} \forall \pi. \forall \pi'. ( \bigwedge_{x \in \mathcal{T}_P} x_\pi \leftrightarrow x_{\pi'} ) \mathcal{U} \left[ \square \square ( \bigwedge_{x \in \mathcal{T}_P} x_\pi \leftrightarrow x_{\pi'} ) \right. \\ \left. \wedge \text{voteA}_\pi \wedge \text{voteB}_{\pi'} \wedge ( \bigwedge_{x \in (\mathcal{T}_P) \setminus \{\text{voteA}, \text{voteB}\}} x_\pi \leftrightarrow x_{\pi'} ) \right] \\ \rightarrow \square (\llbracket \text{winner} \leftarrow A() \rrbracket_{\pi'} \rightarrow \llbracket \text{winner} \leftarrow A() \rrbracket_\pi) \end{aligned} \quad (3)$$

The formula describes two executions on which all predicate terms evaluate to the same value except for one point in time at which the first execution receives a vote for A while the second receives a vote for B. Then, A is the winner in the first trace whenever A is in the second one.

**Elections with multiple candidates.**  $\text{HyperTSL}_{\text{rel}}$  can specify hyperproperties in contracts with an unknown number of candidates. Assume that instead of `voteA` and `voteB`, the contract has a single method `vote` with a parameter `cand` indicating which candidate the caller is voting for. We model `vote` and `cand` as inputs. Now, the `winner` field might range over a domain of unknown size. We express determinism as follows.

$$\forall \pi. \forall \pi'. \square [ (\text{winner}_\pi = \text{winner}_{\pi'}) \mathcal{W} (\neg(\text{cand}_\pi = \text{cand}_{\pi'})) ]$$






Fig. 1: The strategy tree for a HyperTSL strategy  $\sigma$  branching on a single predicate term  $\mathcal{T}_P = \{p(x)\}$ .

#### IV. HYPERTSL SYNTHESIS VIA REDUCTIONS

In this section, we define the synthesis problem of HyperTSL and argue why HyperTSL is better suited for synthesis than HyperTSL<sub>rel</sub>. HyperTSL inherits the undecidable realizability problem from TSL. We therefore show that  $\forall^*$  HyperTSL synthesis can be soundly approximated by HyperLTL, i.e., a strategy for the HyperLTL approximation can be translated to a strategy in HyperTSL. The  $\exists^*$  fragment of HyperTSL synthesis can be approximated by LTL satisfiability to prove unrealizability.

**Definition IV.1.** A HyperTSL formula  $\varphi$  is *realizable* iff there exists a strategy  $\sigma : (2^{\mathcal{T}^P})^+ \rightarrow \mathcal{C}$  such that for every interpretation  $\langle \cdot \rangle : \Sigma_P \rightarrow \mathcal{F}$  the set constructed from all input streams satisfies  $\varphi$ , i.e.,  $\{(\sigma(\iota), \iota) \mid \iota \in \mathcal{I}^\omega\} \models_{\langle \cdot \rangle} \varphi$ .

The above definition defines the realizability problem of HyperTSL by generalizing the TSL definition, similarly to how HyperLTL generalizes LTL. The strategy has the same type as a TSL strategy, i.e., it generates an execution by reacting to the predicate evaluations for the current input stream. The resulting set of executions must satisfy the HyperTSL formula. A visualization of the definition can be found in Figure 1.

In HyperTSL<sub>rel</sub>, formulas contain predicates that relate multiple executions. However, a system can still only react to the inputs it receives, i.e., it cannot include the evaluation of such predicates in its decision making. A definition of the realizability problem could therefore only contain predicates not ranging over multiple executions. This makes many formulas unrealizable or only realizable by trivial strategies most likely not intended by the developer. Consider the two HyperTSL<sub>rel</sub> formulas from Section III-C again. In the first formula, the evaluation of the terms  $\text{votesA}_\pi = \text{votesA}_{\pi'}$  and  $\text{votesB}_\pi = \text{votesB}_{\pi'}$  depends on the chosen executions  $\pi$  and  $\pi'$  and can therefore not be known by the system. The predicate symbol  $=$  is uninterpreted and could be implemented with any binary predicate. A strategy must be winning for all interpretations. The only winning strategy is therefore the trivial strategy to set `winner` to either always  $A()$  or always  $B()$ . We observe a similar phenomenon for the second formula. Since  $=$  is an uninterpreted predicate, the only winning strategy in this case is to always copy the value of `cand` to `winner`.

The above observations make HyperTSL the better candidate to extend TSL synthesis to hyperproperties. This is only

the case for synthesis, though, and as long as predicates and functions remain uninterpreted. If we considered HyperTSL<sub>rel</sub> synthesis with the theory of equality, much more meaningful strategies would be possible. HyperTSL<sub>rel</sub> might also be useful when performing model checking, where formulas can contain predicates that are not part of the system. TSL synthesis with theories has only been studied in one work so far [11] and is therefore out of the scope of this paper. TSL model checking has not been studied at all so far. Both directions will be exciting future work.

#### A. Realizability of $\forall^*$ HyperTSL

We show that the fact that TSL synthesis can be approximated by LTL synthesis [4] carries over to  $\forall^*$  HyperTSL. To do so we recap the approximation of TSL in LTL [4]. The main idea of the approximation is to forget the fact that predicates must evaluate to the same value if evaluated on cells that hold the same value. Then, every predicate  $\tau^P$  and update term  $\llbracket x \leftarrow \tau^f \rrbracket$  can be replaced with an atomic proposition  $a_{\tau^P}$  and  $a_{x\_to\_tau^f}$ , respectively, which is denoted by *syntacticConversion*( $\psi$ ) for a given TSL formula  $\psi$ . For example, *syntacticConversion*( $\Box(p \ x \rightarrow \llbracket x \leftarrow f \ x \rrbracket)$ ) translates to  $\Box(a_{p\_x} \rightarrow a_{x\_to\_f\_x})$ . We call the set of atomic propositions obtained in that translation  $AP_\psi$ . Additionally, we ensure that in each step, exactly one update proposition holds for each cell:

$$\text{cellProps}(\varphi) := \Box \left[ \bigwedge_{c \in \mathcal{C}} \bigvee_{\tau \in \mathcal{T}_c^U} (a_\tau \wedge \bigwedge_{\tau' \in \mathcal{T}_c^U \setminus \{\tau\}} \neg a_{\tau'}) \right]$$

where  $\mathcal{T}_c^U$  is the set of all update terms of  $c$  and  $a_\tau$  is the atomic proposition associated with term  $\tau$ . Now we define

$$[\psi]_{\text{atomic}} := \text{syntacticConversion}(\psi) \wedge \text{cellProps}(\psi)$$

Given an execution  $e = (\varsigma, \iota) \in \mathcal{C}^\omega \times \mathcal{I}^\omega$  and an interpretation  $\langle \cdot \rangle$ , we define the corresponding LTL trace  $[e]_{\text{atomic}}^{\langle \cdot \rangle}$  over  $AP_\psi$ .

$$[e]_{\text{atomic}}^{\langle \cdot \rangle}[i] := \{a_{\tau^P} \mid \eta_{\langle \cdot \rangle}(\varsigma, \iota, i, \tau^P)\} \cup \{a_{x\_to\_tau^f} \mid \varsigma \ i \ x \equiv \tau^f\}$$

Theorem 1 in [4] shows that for a TSL formula  $\psi$ , if  $[\psi]_{\text{atomic}}$  is realizable, then  $\psi$  is realizable. Actually, their proof shows the following stronger result.

**Proposition IV.1** (Proof of Theorem 1 in [4]). *Let  $\psi$  be a TSL formula. For any execution  $e$ , any interpretation  $\langle \cdot \rangle$ , and any point in time  $i$ ,  $e, i \models_{\langle \cdot \rangle} \psi$  iff  $[e]_{\text{atomic}}^{\langle \cdot \rangle}, i \models [\psi]_{\text{atomic}}$ .*

This result entails that realizability of  $[\psi]_{\text{atomic}}$  implies realizability of  $\psi$ ; the opposite direction does not hold. This is because the semantics of TSL forbids *spurious* behavior, which is not captured by the approximation in LTL. For example, if we have  $\llbracket x \leftarrow y \rrbracket$  at some point in time,  $p(x)$  will be equivalent to  $p(y)$  but  $a_{p\_x}$  and  $a_{p\_y}$  are different atomic propositions with possibly different truth values assigned to them. Therefore, every execution  $e$  can be mapped to a trace  $[e]_{\text{atomic}}^{\langle \cdot \rangle}$ , but there are traces  $t$  over  $AP_\psi$  such that there is no execution  $e$  with  $[e]_{\text{atomic}}^{\langle \cdot \rangle} = t$ . The approximation is still very

valuable, because strategies found for the LTL approximations carry over to strategies for the TSL specification. We show that this result can be lifted to  $\forall^*$  HyperTSL and  $\forall^*$  HyperLTL. Given a set of executions  $E$ , we set  $[E]_{\text{atomic}}^{\langle \cdot \rangle} = \{[e]_{\text{atomic}}^{\langle \cdot \rangle} \mid e \in E\}$  and also lift  $[\cdot]_{\text{atomic}}$  to HyperTSL formulas by setting  $[\tau\pi]_{\text{atomic}} = ([\tau]_{\text{atomic}})_{\pi}$  with  $\tau$  being either a predicate term or an update term.

**Lemma IV.2.** *Let  $\varphi$  be a HyperTSL formula. For any set of executions  $E$  and any interpretation  $\langle \cdot \rangle$ ,  $E \models_{\langle \cdot \rangle} \varphi$  iff  $[E]_{\text{atomic}}^{\langle \cdot \rangle} \models [\varphi]_{\text{atomic}}$ .*

*Proof.* Let  $\varphi = (\exists^* \forall^*)^k \cdot \psi$  be a HyperTSL formula consisting of  $k$  blocks of  $\exists^* \forall^*$  quantifiers. For the first direction, let  $E \models_{\langle \cdot \rangle} \varphi$ . We show  $[E]_{\text{atomic}}^{\langle \cdot \rangle} \models [\varphi]_{\text{atomic}}$ . Let  $\varphi[i]$  be the subformula of  $\varphi$  starting from the  $i$ th quantifier block with  $1 \leq i \leq k$ . We keep the invariant to only construct trace assignments  $\Pi$  such that there exists an execution assignment  $\hat{\Pi}$  with  $\Pi(\pi) = [\hat{\Pi}(\pi)]_{\text{atomic}}^{\langle \cdot \rangle}$  and  $\hat{\Pi}, E, 0 \models_{\langle \cdot \rangle} \varphi[i]$ . When choosing the witness traces for the existential variables  $\pi_1^{i+1}, \dots, \pi_n^{i+1}$  of the  $(i+1)$ th quantifier block, we choose  $\Pi(\pi_j^{i+1}) = [e_j^{i+1}]_{\text{atomic}}^{\langle \cdot \rangle}$  where  $e_j^{i+1}$  is assigned to  $\pi_j^{i+1}$  by the proof of satisfaction of  $\varphi[i]$  with respect to the current execution assignment  $\hat{\Pi}$ . Now, by Proposition IV.1 and a simple induction over the structure of  $\psi$ , we get  $\Pi, [E]_{\text{atomic}}^{\langle \cdot \rangle}, 0 \models \psi$ . The argument is similar for the other direction since for every  $t \in [E]_{\text{atomic}}^{\langle \cdot \rangle}$ , there exists an  $e \in E$  such that  $[e]_{\text{atomic}}^{\langle \cdot \rangle} = t$ .  $\square$

Note that the above lemma does not extend to  $\text{HyperTSL}_{\text{rel}}$  as predicates over multiple executions cannot be mapped to HyperLTL. From the lemma we get the following theorem. We will also reuse the lemma in Section V.

**Theorem IV.3.** *Let  $\varphi$  be a  $\forall^*$  HyperTSL formula. If  $[\varphi]_{\text{atomic}}$  is realizable, then  $\varphi$  is realizable.*

*Proof.* Let  $AP_{\varphi} = AP_{\text{in}} \dot{\cup} AP_{\text{out}}$  be the atomic propositions obtained in the translation  $[\varphi]_{\text{atomic}}$ , where  $AP_{\text{in}}$  are the propositions generated for predicate terms and  $AP_{\text{out}}$  the propositions for update terms, respectively. Let  $\sigma_{\text{HyperLTL}} : (2^{AP_{\text{in}}})^+ \rightarrow 2^{AP_{\text{out}}}$  be the realizing strategy for  $[\varphi]_{\text{atomic}}$ . We claim that the following strategy  $\sigma$  realizes  $\varphi$ .

$$\begin{aligned} \sigma(P_1 \dots P_i) \text{c} &:= \tau^f \\ \text{s.t. } [[\text{c} \leftarrow \tau^f]]_{\text{atomic}} &\in \sigma_{\text{HyperLTL}}([P_1]_{\text{atomic}} \dots [P_i]_{\text{atomic}}) \end{aligned}$$

Here, we lift  $[\cdot]_{\text{atomic}}$  to sets by  $[P_j]_{\text{atomic}} = \{[\tau^p]_{\text{atomic}} \mid \tau^p \in P_j\}$ . By definition of *cellProps*, the function term  $\tau^f$  is unique for every  $i$  and every cell  $\text{c}$ . We show that  $\sigma$  realizes  $\varphi$ . Let  $\langle \cdot \rangle$  be any interpretation. Let  $E = \{(\sigma(\iota), \iota) \mid \iota \in \mathcal{I}^{\omega}\}$  be the set of executions obtained from  $\sigma$  and  $T = \text{traces}(\sigma)$  be the set of traces obtained from  $\sigma_{\text{HyperLTL}}$ . Notice that  $[E]_{\text{atomic}}^{\langle \cdot \rangle} \subseteq T$ , but  $T$  may not be a subset of  $[E]_{\text{atomic}}^{\langle \cdot \rangle}$ . Since  $T \models [\varphi]_{\text{atomic}}$  and since universal properties are downwards-closed, also  $[E]_{\text{atomic}}^{\langle \cdot \rangle} \models [\varphi]_{\text{atomic}}$ . Therefore, by Lemma IV.2 we have  $E \models_{\langle \cdot \rangle} \varphi$ .  $\square$

This theorem approximates  $\forall^*$  HyperTSL synthesis by  $\forall^*$  HyperLTL synthesis. Even for HyperLTL, this problem

is undecidable [7]. However, there exists an implementation of bounded synthesis for  $\forall^*$  HyperLTL in the tool BoSyHyper [7]. This approach searches for smallest systems implementing the formula. The efficient reduction from  $\forall^*$  HyperTSL to HyperLTL opens the door to apply BoSyHyper to HyperTSL.

## B. Realizability of $\exists^*$ HyperTSL

Theorem IV.3 uses the downwards closedness of  $\forall^*$  HyperLTL properties. This is necessary since a strategy for HyperTSL produces “fewer” traces than a strategy for the translation of the formula to HyperLTL does. Consequently, the proof does not extend to HyperTSL formulas with existential quantifiers. We show how to translate the synthesis problem of  $\exists^*$  HyperTSL formulas to a TSL satisfiability problem, which we can approximate by LTL satisfiability checking using the translation from TSL to LTL from [4]. In contrast to the case of  $\forall^*$  HyperTSL formulas, the approximation of existential properties in LTL can be used to show unrealizability instead of realizability.

**Theorem IV.4.** *Given an  $\exists^*$  HyperTSL formula  $\varphi$ , there exists a TSL formula  $\psi_{\text{TSL}}$  such that  $\varphi$  is realizable iff  $\psi$  is satisfiable.*

*Proof.* Let  $\varphi = \exists \pi_1 \dots \exists \pi_n \cdot \psi$ . We encode the fact that all traces in a model of  $\varphi$  are producible by a strategy:

$$\varphi_{\text{strat}} := \forall \pi. \forall \pi'. \left[ \bigwedge_{\tau^u \in \mathcal{T}_U} \tau_{\pi}^u \leftrightarrow \tau_{\pi'}^u \right] \mathcal{W} \left[ \bigvee_{\tau^p \in \mathcal{T}_P} \tau_{\pi}^p \leftrightarrow \tau_{\pi'}^p \right]$$

The formula states that while two executions have the same predicate evaluations, they have to perform the same updates. Formula  $\varphi' = \exists \pi_1 \dots \exists \pi_n \cdot \psi \wedge \varphi_{\text{strat}}$  ensures that the executions chosen as witnesses for  $\pi_1, \dots, \pi_n$  can be arranged in a strategy tree. Thus,  $\varphi$  and  $\varphi'$  are equi-realizable. The additional conjunct also ensures that  $\varphi'$  is satisfiable iff it is realizable. Formula  $\varphi'$  is satisfiable iff it is satisfiable by  $n$  executions. The  $\forall$  quantifiers in  $\varphi'$  quantify from these  $n$  executions. We can therefore create an equi-satisfiable  $\exists^n$  formula by “unrolling” the  $\forall$  quantifiers. Let  $\varphi_{\text{strat}} = \forall \pi. \forall \pi'. \psi_{\text{strat}}$ . We define:

$$\varphi_{\text{new}} := \exists \pi_1 \dots \exists \pi_n \cdot \psi \wedge \bigwedge_{1 \leq i, j \leq n} \psi_{\text{strat}}[\pi \mapsto \pi_i, \pi' \mapsto \pi_j]$$

Now we have that  $\varphi$  is realizable iff  $\varphi_{\text{new}}$  is satisfiable. The above formula is satisfiable iff a TSL formula  $\psi_{\text{TSL}}$  is satisfiable, where we construct  $\psi_{\text{TSL}}$  by creating  $n$  copies of each input and each cell.  $\square$

It is an open question whether the two approximations for  $\forall^*$  HyperTSL and  $\exists^*$  HyperTSL can be combined to approximate formulas with quantifier alternations. The challenge here is that universal properties are best over-approximated to obtain realizability results, whereas existential properties would need an under-approximation.

## V. HYPERPROPERTIES FOR SMART CONTRACT SYNTHESIS

We have shown that universal and existential HyperTSL formulas can be approximated by HyperLTL synthesis and LTL satisfiability. For the case of universal properties, this result is, as of now, mostly of theoretic interest. The realizability problem of  $\forall^*$  HyperLTL is undecidable [7]. Bounded hyperproperty synthesis is also still in its infancy.

In the following three sections, we therefore present a more feasible approach to synthesize smart contract control flows. We exploit the fact that we do not synthesize a system from hyperproperties only, but in combination with trace properties describing the temporal behavior of the contract. We build on a workflow recently developed for smart contract control flows from TSL specifications [3] and augment it with methods for HyperTSL. In this section, we recap the TSL approach by showing how to specify some of the trace properties expected of a voting protocol as described in Section III-C. We also show how to include the hyperproperties discussed in Section III-C. In the next section, we analyze whether the combination of hyperproperties with trace properties leads to “pseudo” hyperproperties. In Section VII, we synthesize the most general system from the trace properties and resolve free choices with a repair-like algorithm.

### A. Recap: Smart Contract Synthesis from TSL Specifications

It was recently demonstrated that the temporal behavior of a smart contract can be extensively specified in the syntactic safety fragment of TSL [3]. The idea of the approach is to synthesize the underlying state machine of a smart contract from a TSL specification. The state machine is then translated to Solidity code, which can then be augmented with additional functionality to obtain the desired contract. TSL is a convenient logic to not only describe the correct order of method calls (method “`vote` should not be called after `close` has been successfully called”) but also how fields need to be updated to ensure a correct control flow. This can be achieved through TSL’s cells and uninterpreted functions and predicates. The TSL specification abstracts from the concrete implementation of functions and predicates and instead states how fields have to be updated if certain predicates hold (“if method `vote` is called, update the `winner` field, depending on whether `votesA > votesB`”).

While TSL synthesis is in general undecidable [4], correct solutions can be found through approximation in LTL, which works very well in practice [4]–[6]. By restricting the TSL specification to the past-time fragment, synthesis can be solved comparatively fast using a BDD-based implementation in a tool called SCSYNT [3].

The synthesis is integrated in a feedback loop as shown in Figure 2. The user provides a past-time TSL specification, which is translated to LTL. If the LTL formula is unrealizable, this might be either a specification error or is due to the LTL approximation (in most cases it’s the former). If the specification is realizable, SCSYNT returns the full *winning region* – a Mealy machine that implements the specification and includes all realizing strategies. If there is more than one




Fig. 2: Workflow for the synthesis of a smart contract control flow as developed in [3].

admissible strategy, the specification might be incomplete and should be refined. If not, or if either strategy is fine, a strategy is extracted from the winning region. In the last step, the developer provides an implementation of the functions and predicates (such that the implementation respects the assumptions stated in the specification), and SCSYNT generates a Solidity smart contract.

SCSYNT specializes on a correct temporal control flow and does not check if the provided function and predicate implementations satisfy the stated assumptions. This can be checked by specialized tools based on SMT solving. For recent approaches on how to combine TSL synthesis with reasoning over theories and syntax-guided synthesis, see [11], [12].

### B. Voting Protocol

The part of the TSL voting specification that is relevant for this paper is depicted in Specification 1. Specifications are divided into three parts. *Assumptions* describe properties we can assume, e.g., that in each step only one method is called. We also formulate assumptions on the predicates which must be met by the user providing the implementation of the functions and predicates. *Requirements* state in which order and under which conditions methods can be called. *Obligations* describe how the contract’s fields must be updated.

The specification has four inputs. There is one input for each method indicating the called method. The `sender` input contains the caller of the method (`msg.sender` in Solidity). Cells describe the fields of the contract. We use a function `addOne`, which will be implemented as an increment. There are two predicates `>` and `=`, which will be implemented as the actual “greater” and “equals” predicate. Lastly, we use three constants, one for each candidate A and B, and one for the owner of the contract. Constants will be translated to immutable Solidity fields which are set in the constructor.

Method calls are assumed to be atomic, i.e., exactly one method is called at a time. This follows other work on modeling the control flow of smart contracts, e.g., [13], [14]. Initially, neither candidate has more votes than the other, what can be achieved by setting `votesA` and `votesB` to 0. Lastly, keeping in mind that the `>` predicate will be implemented as the actual “greater” predicate, we assume that `votesA > votesB` and `votesB > votesA` cannot hold true at the same time. The requirements ensure that `close` can only be invoked by the owner of the contract and that no vote can be cast after the voting is closed. The obligations describe the updates of the fields. When `voteA` or `voteB` is invoked,



```

1 Inputs: voteA, voteB, close, reveal, sender
2 Cells: votesA, votesB
3 Functions: addOne
4 Predicates: >, =
5 Constants: owner(), A(), B()
6
7 --- Assumptions ---
8 // Mutex on methods
9 □(voteA → ¬ voteB ∧ ¬ close);
10 □(voteB → ¬ voteA ∧ ¬ close);
11 □(close → ¬ voteA ∧ ¬ voteB);
12 □(voteA ∨ voteB ∨ close);
13
14 //Arithmetic assumptions on > predicate
15 □ ¬(votesA > votesB ∧ votesB > votesA);
16 //Initial assumption on number of votes
17 ¬(votesA > votesB) ∧ ¬(votesB > votesA);
18
19 --- Requirements ---
20 □(close → sender = owner());
21 □(close → ○□ ¬voteA ∧ ¬voteB);
22
23 --- Obligations ---
24 □(voteA → [[votesA ← addOne votesA]]);
25 □(voteB → [[votesB ← addOne votesB]]);
26
27 □(¬ voteA → [[votesA ← votesA]]);
28 □(¬ voteB → [[votesB ← votesB]]);
29
30 □((voteA ∨ voteB)
31   → [[winner ← A()]] ∨ [[winner ← B()]]);
32 □((voteA ∨ voteB) ∧ votesA > votesB
33   → [[winner ← A()]]);
34 □((voteA ∨ voteB) ∧ votesB > votesA
35   → [[winner ← B()]]);
36
37 □(close → □[[winner ← winner]]);

```

Specification 1: Functional specification of the voting contract in TSL.

the respective field counting the votes must be increased. Otherwise, the field must self-update, i.e., stay unchanged. If a vote is cast, then the contract must update the winner to either `A()` or `B()`. If `votesA` holds more votes than `votesB`, the `winner` field is set to `A()` and correspondingly to `B()` if there are more votes in `votesB`. Once `close` has been called, the winner cannot change anymore.

In this paper, we include in Specification 1 only the trace properties which interact with the hyperproperties we are interested in. To obtain a satisfactory voting protocol, one would, for example, also record the voters who already cast their vote and require that every person only votes once by using a predicate. This could be achieved by using a field `voters` and guard calls to `voteA` and `voteB` with an uninterpreted predicate `inVoters voters sender`, which is later implemented as expected.

The winning region synthesized from the above specification is depicted in Figure 3. We write `w` instead of `winner`, and `vA` and `vB` instead of `votesA` and `votesB`, respectively. We also use `+1` instead of `addOne`. We simplify the system

by summarizing multiple transitions between the same states with boolean formulas. For example,  $(\neg(vA > vB) \wedge [w \leftarrow B()]) \vee (\neg(vB > vA) \wedge [w \leftarrow A()])$  expresses that `A` can be chosen as winner as long as `B` does not have strictly more votes.

From the winning region we can see that the TSL specification leaves the free choice which candidate is the winner when neither one has the majority of votes. Strategies permitted by the winning region are, for example, to always choose `A` as winner, or go for the candidate who got the last vote. Which strategies are considered to be suitable is described by the HyperTSL formulations of local determinism (Formula 1 in Section III-C), symmetry (Formula 2 in Section III-C), and the no harm property (Formula 3 in Section III-C). Our goal is to combine the TSL specification with the HyperTSL specifications to obtain a satisfactory strategy. Since the specification of the voting contract abstracts from the concrete number of votes cast with the `>` predicate, we need to add the following assumption to make the combination of the specifications realizable.

$$\begin{aligned}
& \forall \pi. \forall \pi'. \\
& [((votesA > votesB)_\pi \leftrightarrow (votesA > votesB)_{\pi'}) \\
& \quad \wedge ((votesB > votesA)_\pi \leftrightarrow (votesB > votesA)_{\pi'})] \\
& \mathcal{W}(voteA_\pi \leftrightarrow voteA_{\pi'} \vee voteB_\pi \leftrightarrow voteB_{\pi'})
\end{aligned}$$

The property states that the evaluation of the `>` predicate is the same for every two executions as long as they receive the same sequence of votes. The property is satisfied when implementing `addOne` as increment and `>` as “greater than”.

When combining the TSL specification of the voting contract with local determinism, symmetry, and the no harm property, we observe that there is only one valid strategy left to resolve a tie. Local determinism states that the winner may only depend on the greater predicate on the votes and the vote cast in the current step. This forbids to take into account the past of the trace, e.g., by choosing the winner differently in the first and in the second step. Symmetry excludes to brute force and always let the same candidate win in case of a tie. Together with local determinism, the winner must therefore truly depend on the current vote. In combination with local determinism and symmetry, the no harm property forbids to let `A` win if the current vote is for `B` (and vice versa). This leaves as only option to resolve a tie to let the candidate win who got the current vote. In this case, we could therefore replace the hyperproperties with a trace property describing exactly this strategy, which leads to a much easier synthesis problem.

## VI. PSEUDO HYPERPROPERTIES

The previous section has shown that it can be hard to see the implications of combining several hyperproperties with trace properties. It might be the case that the conjunction leaves only one strategy to resolve free choices. To recognize such situations is not easy for a human. We therefore propose to




Fig. 3: Winning region for Specification 1 as synthesized by SCSYNT [3].

preprocess the specification by checking whether the hyperproperties in conjunction with the trace properties effectively define a trace property. To do so, we introduce the notion of pseudo hyperproperties. We first investigate the problem for the general definition based on equivalence checking, and subsequently consider the special case of synthesis.

**Definition VI.1.** A hyperproperty  $H$  is a pseudo hyperproperty iff there is a trace property  $P$  such that

$$\forall T \subseteq (2^{AP})^\omega. T \in H \text{ iff } \forall t \in T. t \in P$$

If  $H$  is a pseudo hyperproperty,  $H$  describes the trace property  $P = \bigcup_{T \in H} T$ . For proofs we will often use the following fact.

**Proposition VI.1.** If  $H$  is a pseudo hyperproperty, then  $H$  is closed under union and subsets.

The downwards closure property in the above proposition implies that only hyperproperties expressible with  $\forall^*$  formulas can be pseudo hyperproperties. The following proposition establishes the convenient fact that if a  $\forall^*$  HyperTSL formula  $\varphi$  is a pseudo hyperproperty, then the formula obtained by using only a single execution variable in the body of the formula is equivalent to  $\varphi$ . The proposition is close to the corresponding result for HyperLTL, which has been observed in context of using HyperLTL synthesis for the synthesis of linear distributed architectures [7]. Given a HyperTSL formula  $\varphi = \forall \pi_1 \dots \forall \pi_n. \psi$ , we define its  $\forall^1$  counterpart as  $\varphi[\pi] := \forall \pi. \psi[\pi_1 \mapsto \pi, \dots, \pi_n \mapsto \pi]$ .

**Proposition VI.2.** A  $\forall^*$  HyperTSL formula  $\varphi$  describes a pseudo hyperproperty iff  $\varphi \leftrightarrow \varphi[\pi]$ .

*Proof.* We observe that a hyperproperty  $H$  is a pseudo hyperproperty iff

$$\forall T \subseteq (2^{AP})^\omega. T \in H \text{ iff } \forall t \in T. \{t\} \in H$$

Thus, for any set of executions  $E$  and any interpretation  $\langle \cdot \rangle$ , we have  $E \models_{\langle \cdot \rangle} \varphi$  iff  $\forall e \in E. \{e\} \models_{\langle \cdot \rangle} \varphi$  iff  $E \models_{\langle \cdot \rangle} \varphi[\pi]$ .  $\square$

The proposition describes how to check if a HyperTSL property is a pseudo hyperproperty, namely by checking if  $\varphi$  is equivalent to  $\varphi[\pi]$ . If it is, the TSL formula  $\varphi[\pi]$  can be used for synthesis. The check is undecidable for HyperTSL, however.

**Theorem VI.3.** It is undecidable whether a HyperTSL formula  $\varphi$  describes a pseudo hyperproperty.

*Proof.* The proof follows from undecidability of the satisfiability problem of TSL [10]. The formula  $\varphi = \forall \pi. \forall \pi'. (\text{p } a)_\pi \leftrightarrow (\text{p } a)_{\pi'}$  is not a pseudo hyperproperty as there exists an interpretation of  $\text{p}$  such that  $E$  consists of executions for which  $\text{p } a$  always evaluates to *true*, and  $E'$  consists of executions for which  $\text{p } a$  always evaluates to *false*.  $E$  and  $E'$  are not closed under union for  $\varphi$ . Now, let a TSL formula  $\psi$  be given for which we assume w.l.o.g. that  $\text{p}$  is a fresh predicate and  $a$  a fresh input. Since  $\forall \pi. \forall \pi'. \text{false}$  is a pseudo hyperproperty, the formula  $\forall \pi. \forall \pi'. ((\text{p } a)_\pi \leftrightarrow (\text{p } a)_{\pi'}) \wedge \psi_\pi$  is a pseudo hyperproperty iff  $\psi$  is unsatisfiable. Here,  $\psi_\pi$  denotes the formula obtained by lifting  $\psi$  to a HyperTSL formula by annotating predicate and update terms with the trace variables  $\pi$ .  $\square$

Fortunately, the above negative result comes with the remedy that we can approximate the problem in HyperLTL, which is a result of Lemma IV.2.

**Theorem VI.4.** Let  $\varphi$  be a HyperTSL formula. If  $[\varphi]_{\text{atomic}}$  is a pseudo hyperproperty, then so is  $\varphi$ .

*Proof.* Closedness under union and subsets follows from Lemma IV.2 and since for every set of executions  $E$  and interpretation  $\langle \cdot \rangle$ ,  $[E]_{\text{atomic}}^{\langle \cdot \rangle} = \bigcup_{e \in E} [e]_{\text{atomic}}^{\langle \cdot \rangle}$ , i.e.,  $[E \cup E']_{\text{atomic}}^{\langle \cdot \rangle} = [E]_{\text{atomic}}^{\langle \cdot \rangle} \cup [E']_{\text{atomic}}^{\langle \cdot \rangle}$ , and  $E \subseteq E'$  iff  $[E]_{\text{atomic}}^{\langle \cdot \rangle} \subseteq [E']_{\text{atomic}}^{\langle \cdot \rangle}$ .  $\square$

Using the above result, we can translate a  $\forall^*$  HyperTSL formula to HyperLTL and check its equivalence with the corresponding  $\forall^1$  formula. For HyperLTL, this check is decidable, because of the decidability of the  $\exists^* \forall^*$  fragment of HyperLTL [7]: a HyperLTL formula  $\varphi = \forall \pi_1 \dots \forall \pi_n. \psi$  is equivalent to its  $\forall^1$  counterpart  $\varphi[\pi]$  iff  $\varphi[\pi] \rightarrow \varphi$  (because of the semantics of the  $\forall$ ). By merging the quantifiers, this implication is valid iff the formula

$$\exists \pi_1 \dots \exists \pi_n. \forall \pi. \neg \psi \wedge \psi[\pi_1 \mapsto \pi, \dots, \pi_n \mapsto \pi]$$

is unsatisfiable. Thus, if the check returns UNSAT, the HyperTSL formula is guaranteed to describe a pseudo hyperproperty, so we can replace the hyperproperty with its TSL version.

So far, we described the check for pseudo hyperproperties in terms of satisfiability checking. This can be useful to detect superfluous specifications or mistakes. For example, the

developer could have already specified in TSL that in case of a tie, the current vote determines the winner.

$$\begin{aligned} & \Box(\neg(\text{votesA} > \text{votesB}) \wedge \neg(\text{votesB} > \text{votesA})) \\ & \rightarrow \llbracket \text{winner} \leftarrow A() \rrbracket \leftrightarrow \text{voteA} \end{aligned}$$

With this specification, the hyperproperties stated in Section III-C are entailed by the specification. In the case of the “no harm” property, for example, the check with EAHyper reveals correctly within 0.003 seconds that the conjunction of the trace properties with the “no harm” property is a pseudo hyperproperty.

Without the above TSL specification however, the check for the conjunction of local determinism, symmetry, and the no harm property is labeled as a hyperproperty by EAHyper, even though there is only one possible strategy. The reason for that is the difference between realizability and satisfiability. Consider the following two traces.

- (1)  $\{\text{voteA}, \llbracket \text{winner} \leftarrow A() \rrbracket\}^\omega$
- (2)  $\{\text{voteB}, \llbracket \text{winner} \leftarrow A() \rrbracket\}^\omega$

In the first trace, neither  $\text{votesA} > \text{votesB}$  nor  $\text{votesB} > \text{votesA}$ , and there is always a vote for A, who is always the winner. The second trace is similar but there is always a vote for B. As single sets, both traces satisfy all three hyperproperties. Together, they do not (because of the symmetry requirement). Therefore, EAHyper returns that the conjunction of the properties is not a pseudo hyperproperty (even in combination with the full contract specification). When solving the realizability problem, however, we do not need to find an *equivalent* trace property, it is enough to find a trace property which evaluates the same on all sets generated by strategies. We therefore adapt the notion of pseudo hyperproperties to realizability. We give the definition for hyperproperties of traces over  $AP_{in} \dot{\cup} AP_{out}$ , but the definition applies to traces of any type, particularly also to hyperproperties over TSL-like executions.

**Definition VI.2.** Let  $H$  be a hyperproperty over atomic propositions  $AP = AP_{in} \dot{\cup} AP_{out}$ .  $H$  is a pseudo hyperproperty in the context of realizability if there is a trace property  $P$  such that

$$\begin{aligned} \forall \sigma : (2^{AP_{in}})^+ & \rightarrow 2^{AP_{out}}. \\ \text{traces}(\sigma) \in H & \text{ iff } \forall t \in \text{traces}(\sigma). t \in P \end{aligned}$$

The definition implies that a HyperTSL formula  $\varphi$  is a pseudo hyperproperty in the context of realizability iff there is a TSL formula  $\psi$  such that a strategy  $\sigma$  realizes  $\varphi$  iff  $\sigma$  realizes  $\psi$ . In the context of realizability, however,  $\psi$  may not be equivalent to  $\varphi[\pi]$ . To show why, we give an example in HyperLTL but the reasoning carries over to HyperTSL. Consider the example of the following HyperLTL formulas over input  $i$  and output  $o$ , which are abstract, simplified

versions of local determinism, symmetry, and the no harm property.

$$\begin{aligned} & \forall \pi. \forall \pi'. (i_\pi \leftrightarrow i_{\pi'}) \rightarrow (o_\pi \leftrightarrow o_{\pi'}) \\ & \forall \pi. \forall \pi'. (i_\pi \leftrightarrow i_{\pi'}) \rightarrow (o_\pi \leftrightarrow o_{\pi'}) \\ & \forall \pi. \forall \pi'. i_\pi \wedge \neg i_{\pi'} \wedge o_{\pi'} \rightarrow o_\pi \end{aligned}$$

The properties are not temporal, so they are realizable iff there are values  $o_1, o_2$  such that  $\{\{i, o_1\}^\omega, \{\neg i, o_2\}^\omega\}$  satisfies the three formulas. Let  $\varphi$  be their conjunction. There are only four possible assignments of  $o_1, o_2$  to boolean values, each of which corresponds to a possible strategy. Indeed, for every possible strategy,  $\varphi$  is satisfied iff the trace property  $i \leftrightarrow o$  is satisfied. However, the same does not hold for the trace property *true*, which is equivalent to  $\varphi[\pi]$ . Unfortunately, the general problem is undecidable already for HyperLTL, due to the undecidability of the realizability problem of the  $\forall^*$  fragment of HyperLTL.

**Theorem VI.5.** *Given a HyperLTL formula  $\varphi$ , it is in general undecidable if there exists an LTL formula  $\psi$  such that for all strategies  $\sigma$ ,  $\text{traces}(\sigma) \models \varphi$  iff  $\forall t \in \text{traces}(\sigma). t \models \psi$ .*

*Proof.* We show undecidability by a reduction from the  $\forall^*$  HyperLTL realizability problem, which is undecidable [7]. Let  $\rho$  be a  $\forall^*$  HyperLTL formula over  $AP = AP_{in} \dot{\cup} AP_{out}$ . We assume w.l.o.g. that  $AP_{in}$  is non-empty. We define  $\varphi$  as follows, where  $o$  is an output proposition that does not occur in  $\rho$ .

$$\varphi := \rho \wedge \forall \pi. \forall \pi'. o_\pi \leftrightarrow o_{\pi'}$$

We claim that  $\rho$  is unrealizable iff there exists an LTL formula  $\psi$  such that for all strategies  $\sigma$ ,  $\text{traces}(\sigma) \models \varphi$  iff  $\forall t \in \text{traces}(\sigma). t \models \psi$ . If  $\rho$  is unrealizable, then we choose  $\psi := \text{false}$ . Since  $\rho$  is unrealizable, it holds for all  $\sigma$  that  $\text{traces}(\sigma) \not\models \varphi$  iff  $\forall t \in \text{traces}(\sigma). t \models \psi$ . For the other direction, assume that  $\rho$  is realizable and that there exists a suitable LTL formula  $\psi$ . Let  $\sigma$  be the realizing strategy of  $\rho$ . Since  $o$  is fresh for  $\rho$ , we can extend  $\sigma$  to  $\sigma_1$  and  $\sigma_2$ , where  $\sigma_1$  adds  $o$  to the first output (for any input), and  $\sigma_2$  does not add  $o$  (also for any input). Both strategies realize  $\varphi$ , therefore, by our assumption, both strategies realize  $\psi$ . Now, let  $i \in AP_{in}$  be any input. We define  $\sigma_3$  as the strategy that adds  $o$  to the first output exactly if  $i$  holds in the first position of the input sequence. Since every trace generated by  $\sigma_3$  is either a trace of  $\sigma_1$  or  $\sigma_2$ ,  $\sigma_3$  realizes  $\psi$ . However, it does not realize  $\varphi$ , which contradicts our assumption.  $\square$

Lastly, we observe that we can decide if a HyperTSL formula is a pseudo hyperproperty if the formula contains as conjunct a local determinism formula like the one in our running example. We define local determinism in the general case as follows.

*localDeterminism* :=

$$\forall \pi. \forall \pi'. \Box \left( \left( \bigwedge_{\tau^p \in \mathcal{T}_P} \tau_\pi^p \leftrightarrow \tau_{\pi'}^p \right) \rightarrow \left( \bigwedge_{\tau^c \in \mathcal{T}_U} \tau_\pi^c \leftrightarrow \tau_{\pi'}^c \right) \right)$$

**Proposition VI.6.** For every  $\forall^*$  HyperTSL formula  $\varphi$  over predicate terms  $\mathcal{T}_P$  and update terms  $\mathcal{T}_U$ , if  $\varphi = \text{localDeterminism} \wedge \varphi'$ , then deciding whether  $\varphi$  is a pseudo hyperproperty in the context of realizability is equivalent to a Boolean SAT problem.

The above proposition follows from the observation that  $\text{localDeterminism} \wedge \varphi$  is realizable iff there is a positional strategy that always assigns the same cell updates for the same predicate evaluations, independently of the trace's past. In this case, there are only finitely many combinations of predicate evaluations and the problem becomes finite.

## VII. RESOLVING CHOICES WITH REPAIR

If a hyperproperty is not a pseudo hyperproperty, we synthesize the winning region of the trace property and then check if free choices can be resolved such that the hyperproperty is satisfied. This is a repair algorithm additionally respecting the distinction between inputs and outputs. We first give a formal definition of the problem we are interested in for HyperTSL formulas. As observed several times, the formal problem is undecidable for HyperTSL, but a sound approximation can be achieved through HyperLTL. We then discuss how to simplify the problem for  $\forall^*$  formulas and present a prototype implementation, which successfully repairs the synthesized voting contract with respect to determinism, symmetry, and the no harm property.

The result of synthesizing the extended control flow of a smart contract from TSL specifications is a Mealy machine as depicted in Figure 3. A *Mealy machine* is a tuple  $(Q, \Sigma, \delta, q_0)$ , where  $Q$  is a set of states,  $\Sigma = 2^{AP}$  with  $AP = AP_{in} \dot{\cup} AP_{out}$  is a labeling alphabet,  $\delta : Q \times \Sigma \rightarrow Q$  is a transition function, and  $q_0$  is the initial state.

**Definition VII.1.** A *free choice* in a Mealy machine  $(Q, \Sigma, \delta, q_0)$  consists of a state  $q \in Q$  and an input  $i \in 2^{AP_{in}}$  such that there are at least two outputs  $o_1, o_2 \in 2^{AP_{out}}$  with  $(q, i \cup o_1, q_1) \in \delta$  and  $(q, i \cup o_2, q_2) \in \delta$  for some  $q_1, q_2$ .

**Definition VII.2.** A Mealy machine  $\hat{S}$  is a *refinement* of a Mealy machine  $S$  if  $\hat{Q} = Q$ ,  $\hat{\Sigma} = \Sigma$ ,  $\hat{q}_0 = q_0$  and  $\hat{\delta} \subseteq \delta$  such that for every  $q \in Q, i \in 2^{AP_{in}}$ , if there are  $o \in 2^{AP_{out}}$  and  $q' \in Q$  such that  $\delta(q, i \cup o, q')$ , then there are  $\hat{o} \in 2^{AP_{out}}$  and  $\hat{q}' \in \hat{Q}$  such that  $\hat{\delta}(q, i \cup \hat{o}, \hat{q}')$  and  $\delta(q, i \cup \hat{o}, \hat{q}')$ .

The above definition ensures that if  $S$  defines a set of strategies for a TSL property  $\psi$ , then a refinement  $\hat{S}$  still describes at least one strategy for  $\psi$ . Our goal is to refine  $S$  such that  $\hat{S}$  models a HyperTSL property  $\varphi$ .

The system  $S$  is produced from the LTL approximation, it might therefore contain spurious traces which, for any interpretation, cannot be produced by a combination of an input stream and a computation. We define when such a system satisfies a HyperTSL formula.

**Definition VII.3.** Let  $\varphi$  be a HyperTSL formula over  $\mathcal{T}_P, \mathcal{T}_U$  and let  $S$  be a Mealy machine over  $AP = [\mathcal{T}_P]_{\text{atomic}} \cup$

$[\mathcal{T}_U]_{\text{atomic}}$ .  $S$  models  $\varphi$  if for every interpretation  $\langle \cdot \rangle$ :

$$\{(\varsigma, \iota) \mid \exists t \in \text{traces}(S). t = [(\varsigma, \iota)]_{\text{atomic}}^{\langle \cdot \rangle}\} \models_{\langle \cdot \rangle} \varphi$$

The definition creates for each interpretation the set of executions which have a trace through  $S$  and then checks if  $\varphi$  is satisfied on these sets. We show that we cannot check the satisfaction of a HyperTSL formula on such a system directly.

**Lemma VII.1.** Given a HyperTSL formula  $\varphi$ , it is undecidable whether a Mealy machine  $S$  over  $AP = [\mathcal{T}_P]_{\text{atomic}} \cup [\mathcal{T}_U]_{\text{atomic}}$  models  $\varphi$ .

*Proof.* We show that the problem is already undecidable for TSL formulas expressed in HyperTSL as  $\forall \pi. \psi_\pi$ , where  $\psi_\pi$  is the  $\pi$ -indexed version of a TSL formula  $\psi$ . We proceed by reduction from the unsatisfiability problem of TSL. Let  $\psi$  be a TSL property.  $\psi$  is unsatisfiable if for every  $\langle \cdot \rangle$  and every  $\varsigma, \iota$ , it holds that  $\varsigma, \iota \models_{\langle \cdot \rangle} \neg \psi$ . Let  $S$  be the system that produces every trace over  $AP = [\mathcal{T}_P]_{\text{atomic}} \cup [\mathcal{T}_U]_{\text{atomic}}$ . Then,  $S$  models  $\forall \pi. \neg \psi_\pi$  iff  $\psi$  is unsatisfiable.  $\square$

However, owing to Lemma IV.2, for every  $\forall^*$  HyperTSL formula  $\varphi$ , if a Mealy machine  $S$  models  $[\varphi]_{\text{atomic}}$ , then it also models  $\varphi$ .

**Theorem VII.2.** Let a  $\forall^*$  HyperTSL formula and a Mealy machine  $S$  over  $AP = [\mathcal{T}_P]_{\text{atomic}} \cup [\mathcal{T}_U]_{\text{atomic}}$  be given. If  $S$  models  $[\varphi]_{\text{atomic}}$ , then  $S$  models  $\varphi$ .

*Proof.* For every  $\langle \cdot \rangle$ , let  $S_{\langle \cdot \rangle} = \{(\varsigma, \iota) \mid \exists t \in \text{traces}(S). t = [(\varsigma, \iota)]_{\text{atomic}}^{\langle \cdot \rangle}\}$ . By construction,  $[S_{\langle \cdot \rangle}]_{\text{atomic}}^{\langle \cdot \rangle} \subseteq \text{traces}(S)$ . Since  $S \models [\varphi]_{\text{atomic}}$  and since universal properties are downwards closed,  $[S_{\langle \cdot \rangle}]_{\text{atomic}}^{\langle \cdot \rangle} \models [\varphi]_{\text{atomic}}$  and by Lemma IV.2 we have  $S_{\langle \cdot \rangle} \models_{\langle \cdot \rangle} \varphi$ .  $\square$

Based on this theorem, we repair the synthesized transition system with respect to the corresponding HyperLTL formula. To do so, we enumerate all refinements which leave exactly one transition for each choice. Thereby, the refinement still implements a strategy. Since the HyperLTL is a universal formula, this approach detects a correct refinement iff there is one. Note that also with respect to HyperLTL, this cannot be a complete method, as  $\forall^*$  HyperLTL synthesis is undecidable [7]. The general problem of repairing HyperLTL formulas has been discussed in [15], the problem is NP-complete for  $\forall^*$  HyperLTL. An approach similar to our has been described for controller synthesis in [16], which distinguishes between controllable and uncontrollable inputs. Implementations have not been developed yet.

## VIII. IMPLEMENTATION AND EXPERIMENTS

We implemented a prototype of the repair algorithm as a Python script. Given a synthesized smart contract in form of a Mealy machine and a universal HyperLTL formula, the script checks if the complete system satisfies the property. If not, it searches for free choices, performs a self-composition, and checks if one of the possibilities to resolve the choices satisfies the LTL body of the HyperLTL formula. For LTL

TABLE I: Results of the prototype implementation of the repair algorithm. Times (t) are given in seconds. # refers to the number of nuXmv calls that were needed to find a repair.

Property	only <code>vote</code>		+ <code>close</code>		+ <code>owner</code>		full	
	t	#	t	#	t	#	t	#
Local determinism	0.170	1	0.475	1	2.577	1	7.049	1
Local symmetry	0.229	2	1.251	6	64.90	86	308.1	120
Global no harm	0.163	1	0.473	1	2.786	1	219.9	86
Determinism	0.147	1	0.612	1	27.63	35	6.825	1
Symmetry	0.254	2	1.630	6	29.03	35	6.571	1
No harm	0.170	1	0.704	1	2.993	1	90.81	35
Determinism + Symmetry + No harm	0.274	3	2.105	6	217.7	256	760.4	256

model checking we use the state-of-the-art model checker nuXmv [17]. Using our implementation, we conducted experiments on different version of the voting contract and on a blinded auction contract.

### A. Voting Contract

We repaired four variants of the voting protocol with respect to the hyperproperties described in this paper. The results are depicted in Table I. None of the three systems initially satisfies either of the properties, the number of calls given in the table therefore refers to the number of refinements tested with nuXmv.

The first variant (“only `vote`”) reduces the contract to the core specification needed for the hyperproperties to make sense. It has only a `vote` method and can be implemented as a single-state system. It has two state-input combinations with a free choice, in each case there are two possible transitions. The second variant (“+ `close`”) adds the `close` method and initial assumptions, resulting in a three-state system with four choices, again each with two options. The third variant (“+ `owner`”) is the contract as described in Specification 1. It has eight choices. This results in  $2^8$  different combinations that might need to be checked.

Lastly, we extended the voting contract with additional features following the voting example of the Solidity Documentation [18]. The extended version additionally records the registered voters in a field `voters` and which addresses have voted in a field `voted`. The contract also has two additional methods: `giveRightToVote` may be called by the owner of the contract and adds addresses to `voters`; `getWinner` may be called after the voting has been closed to learn which candidate won. The synthesized state machine has again eight choices but is naturally larger than the state machine depicted in Figure 3.

Determinism, local determinism, symmetry, and the no harm property indicate the formulas given in Section III-C. For symmetry, we included the necessary assumption described in Section V-B. The “global no harm” property formulates the no harm property with a single  $\square$  instead of  $\mathcal{W}$ . Local symmetry

states symmetry with respect to the greater predicate as well as the `voteA` and `voteB` inputs, similar to local determinism.

### B. Blind Auction

As a second contract, we specified a blind auction in TSL, following [19]. Similar to the voting protocol, we had to restrict ourselves to a finite number of bidders, otherwise the hyperproperties would not have been expressible in HyperTSL. The idea of a blind auction is that bidders send a hash of their actual bid and deposit a value which might be higher or lower than the actual bid. After the bidding is closed, bidders reveal their bids. If the hash fits the revealed bid and the deposited value is higher than the actual bid, the bid is valid. The winner of the auction is the bidder with the highest valid bid. Bidders can also withdraw their deposits once a higher bid was revealed.

We specified the contract with two bidders. We use methods `bidA`, `bidB`, `closeBidding`, `revealA`, `revealB`, `closeRevealing`, and `withdraw`. One of the temporal requirements is that bids can only be placed as long as the bidding has not been closed.

$$\square(\text{closeBidding} \rightarrow \bigcirc \square \neg(\text{bidA} \vee \text{bidB}))$$

The specification also reasons about fields `bidsA`, `bidsB`, `highestBidder`, `highestBid`. One of the obligations on the fields is to update `highestBid` if bidder A reveals a correct bid which is higher than all bids revealed so far.

$$\begin{aligned} &\square(\text{revealA} \wedge \text{valid bid secret} \\ &\quad \wedge \text{bid} > \text{highestBid} \\ &\quad \rightarrow \llbracket \text{highestBid} \leftarrow \text{bid} \rrbracket) \end{aligned}$$

In the above specification, `bid` and `secret` are parameters of the method `revealA`. As for the voting specification, the TSL specification describes the control flow of the contract, i.e., in which order methods have to be called and how fields need to be updated. Other functionality like checking if a bid is valid needs to be implemented in the predicate `validBid` to obtain a working contract.

Similar to the voting case, we left the free choice which bid is stored as highest bid if the currently revealed bid is the same as the current highest bid. The synthesized state machine leaves free choices at two nodes. Using our implementation, we obtained strategies to resolve a tie in a way that local determinism or local symmetry are satisfied. We report the running times and number of calls to nuXmv in Table II.

TABLE II: Repairing the blind auction state machine. Times (t) are given in seconds. # refers to the number of nuXmv calls.

Property	t	#
Local determinism	11.13s	1
Local symmetry	17.42s	2



### C. Evaluation

Our experiments show that the hyperproperties discussed in this paper are realizable on top of variants of a voting contract and a blinded auction. The evaluation shows that the runtime mainly depends on the number of choice options that had to be tested with NUXMV. The runtime also increases with larger formulas and larger state machines.

This implementation is only a prototype to see if the idea to resolve free choices via repair works for relevant hyperproperties. There are many options to improve performance. For example, one could switch to a model checker that is based on Büchi or Parity automata (e.g., spot [20]). Like this, the LTL formula (which stays the same for every call) would have to be translated to an automaton only once.

## IX. RELATED WORK

Our work brings together research on formal methods for smart contract control flows, reactive synthesis using TSL, and logics for hyperproperties.

**Formal methods for smart contract control flows.** Most closely related work is the recent approach to synthesize the control and data flow of a smart contract from TSL specifications [3]. We build on the fact that they synthesize the full winning region of the specification, which we refine to find a strategy that satisfies a HyperTSL specification. Another approach to synthesize smart contract control flows is [21], which synthesizes the order of method calls from an LTL specification. Formal guarantees have been also obtained by modeling the control flow of a contract directly with a finite state machine and transform the state machine to Solidity code. In the FSOLIDM framework [13] the state machine can be modeled in a graphical editor. A lot of work has been invested to verify a contract against correctness properties instead of generating it automatically. The VERISOL [14], [22] framework checks a smart contract against properties defined with finite state machines. VERX [23] and SMARTPULSE [24] verify that a contract satisfies specifications given in LTL-like logics, where SMARTPULSE can handle not only safety but also liveness properties.

**Temporal Stream Logic.** TSL is a temporal logic that has been developed to simplify reasoning about infinite-state systems by separating control and data [4]. Due to its expressiveness, TSL synthesis is undecidable in general but sound CEGAR loops based on LTL synthesis work well in practice [4]. TSL has been successfully applied to specify and synthesize an arcade shooter game running on an FPGA [5] and functional reactive programs [6]. The TSL synthesis approach has been extended with decidable theories like linear integer arithmetic [11]. Here, counterexamples provided by the CEGAR loop are checked for consistency with the theory using an SMT solver. TSL modulo first-order theories has also been studied with respect to its satisfiability problem [10].

**Hyperproperties.** The interest in hyperproperties, or relational properties, stems from information flow policies like non-interference and observational determinism. Since the term has been coined for the general class of properties [25], a

range of logics have been developed for the specification of hyperproperties. Examples are temporal logics like HyperLTL, HyperCTL\* [8], and HyperPDL [26], first-order and second-order logics [27], [28], or asynchronous extensions mostly based on HyperLTL [29]–[31]. All of the above logics are targeted to specify properties in the context of hardware and therefore assume a finite state space. Recently, the temporal logic of actions (TLA) has also been used to verify hyperproperties via self-composition [32]. The synthesis from hyperproperties is known to be challenging. HyperLTL can express distributed architectures, which makes the synthesis from  $\forall^*$  properties already undecidable [7]. A simpler variant of synthesis is controller synthesis, where the state space is given. Related to our approach, the resulting problem is similar to the repair problem and is therefore decidable for HyperLTL [16]. The classic program repair problem does not distinguish between inputs and outputs and has been studied for HyperLTL as well [15]. The analysis of smart contract with respect to hyperproperties has not received enough attention yet. A recent approach uses type-checking to ensure information flow policies such as integrity in smart contracts in the context of reentrancy attacks [1]. Other hyperproperties identified in the context of smart contracts are integrity properties to prevent reentrancy attacks [33], [34]

## X. CONCLUSION AND FUTURE WORK

We have presented two logics, HyperTSL and HyperTSL<sub>rel</sub>, which extend the temporal logic TSL to express hyperproperties. Inherited from TSL, both logics express the control flow of infinite-state systems by abstracting from concrete data with cells and uninterpreted functions and predicates. HyperTSL<sub>rel</sub> predicates relate several executions, which makes the logic more expressive than HyperTSL. HyperTSL, on the other hand, is more suitable for synthesis tasks. We have demonstrated that both logics can express typical hyperproperties that occur in the context of smart contracts. The realizability problem of the universal fragment of HyperTSL can be approximated with HyperLTL realizability, the unrealizability problem of the existential fragment can be approximated with LTL satisfiability. As a step towards making synthesis from hyperproperties reality, we have described two approaches to integrate universal HyperTSL specifications in the synthesis workflow of smart contracts specified with TSL. The first approach checks for overly complicated specifications which have equivalent formulations in TSL. The second approach refines the winning region obtained from the TSL synthesis to resolve choices in a way that satisfy the hyperproperty.

**Future work.** This work has opened many interesting paths for future work. To profit from the expressiveness of HyperTSL<sub>rel</sub>, one would have to include theories (e.g., the theory of equality) into the synthesis from HyperTSL<sub>rel</sub> specifications. Prior work has made steps in that direction for TSL [10], [11], one could extend these ideas to HyperTSL<sub>rel</sub>. Using theories, HyperTSL<sub>rel</sub> would also be interesting in the context of model checking smart contracts or other infinite-state systems. Another open question is how to approximate

the HyperTSL realizability problem for formulas with quantifier alternations.

## REFERENCES

- [1] E. Cecchetti, S. Yao, H. Ni, and A. C. Myers, “Compositional security for reentrant applications,” in *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 2021, pp. 1249–1267. [Online]. Available: <https://doi.org/10.1109/SP40001.2021.00084>
- [2] —, “Securing smart contracts with information flow,” in *International Symposium on Foundations and Applications of Blockchain*, 2020.
- [3] B. Finkbeiner, J. Hofmann, F. Kohn, and N. Passing, “Reactive synthesis of smart contract control flows,” 2022. [Online]. Available: <https://arxiv.org/abs/2205.06039>
- [4] B. Finkbeiner, F. Klein, R. Piskac, and M. Santolucito, “Temporal Stream Logic: Synthesis Beyond the Booleans,” in *Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part I*, ser. Lecture Notes in Computer Science, I. Dillig and S. Tasiran, Eds., vol. 11561. Springer, 2019, pp. 609–629.
- [5] G. Geier, P. Heim, F. Klein, and B. Finkbeiner, “Syntriods: Synthesizing a game for fpgas using temporal logic specifications,” *CoRR*, vol. abs/2101.07232, 2021. [Online]. Available: <https://arxiv.org/abs/2101.07232>
- [6] B. Finkbeiner, F. Klein, R. Piskac, and M. Santolucito, “Synthesizing functional reactive programs,” in *Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2019, Berlin, Germany, August 18-23, 2019*, R. A. Eisenberg, Ed. ACM, 2019, pp. 162–175. [Online]. Available: <https://doi.org/10.1145/3331545.3342601>
- [7] B. Finkbeiner, C. Hahn, P. Lukert, M. Stenger, and L. Tentrup, “Synthesis from hyperproperties,” *Acta Informatica*, vol. 57, no. 1-2, pp. 137–163, 2020. [Online]. Available: <https://doi.org/10.1007/s00236-019-00358-2>
- [8] M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez, “Temporal logics for hyperproperties,” in *Principles of Security and Trust - Third International Conference, POST 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, ser. Lecture Notes in Computer Science, M. Abadi and S. Kremer, Eds., vol. 8414. Springer, 2014, pp. 265–284. [Online]. Available: [https://doi.org/10.1007/978-3-642-54792-8\\_15](https://doi.org/10.1007/978-3-642-54792-8_15)
- [9] B. Finkbeiner, C. Hahn, P. Lukert, M. Stenger, and L. Tentrup, “Synthesizing reactive systems from hyperproperties,” in *Proceedings of CAV*, ser. LNCS, vol. 10981. Springer, 2018, pp. 289–306.
- [10] B. Finkbeiner, P. Heim, and N. Passing, “Temporal stream logic modulo theories,” in *Foundations of Software Science and Computation Structures - 25th International Conference, FOSSACS 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings*, ser. Lecture Notes in Computer Science, P. Bouyer and L. Schröder, Eds., vol. 13242. Springer, 2022, pp. 325–346. [Online]. Available: [https://doi.org/10.1007/978-3-030-99253-8\\_17](https://doi.org/10.1007/978-3-030-99253-8_17)
- [11] B. Maderbacher and R. Bloem, “Reactive synthesis modulo theories using abstraction refinement,” *CoRR*, vol. abs/2108.00090, 2021. [Online]. Available: <https://arxiv.org/abs/2108.00090>
- [12] W. Choi, B. Finkbeiner, R. Piskac, and M. Santolucito, “Can reactive synthesis and syntax-guided synthesis be friends?” in *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, R. Jhala and I. Dillig, Eds. ACM, 2022, pp. 229–243. [Online]. Available: <https://doi.org/10.1145/3519939.3523429>
- [13] A. Mavridou and A. Laszka, “Designing secure ethereum smart contracts: A finite state machine based approach,” in *Financial Cryptography and Data Security - 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26 - March 2, 2018, Revised Selected Papers*, ser. Lecture Notes in Computer Science, S. Meiklejohn and K. Sako, Eds., vol. 10957. Springer, 2018, pp. 523–540. [Online]. Available: [https://doi.org/10.1007/978-3-662-58387-6\\_28](https://doi.org/10.1007/978-3-662-58387-6_28)
- [14] Y. Wang, S. K. Lahiri, S. Chen, R. Pan, I. Dillig, C. Born, I. Naseer, and K. Ferles, “Formal verification of workflow policies for smart contracts in azure blockchain,” in *Verified Software. Theories, Tools, and Experiments - 11th International Conference, VSTTE 2019, New York City, NY, USA, July 13-14, 2019, Revised Selected Papers*, ser. Lecture Notes in Computer Science, S. Chakraborty and J. A. Navas, Eds., vol. 12031. Springer, 2019, pp. 87–106. [Online]. Available: [https://doi.org/10.1007/978-3-030-41600-3\\_7](https://doi.org/10.1007/978-3-030-41600-3_7)
- [15] B. Bonakdarpour and B. Finkbeiner, “Program repair for hyperproperties,” in *Automated Technology for Verification and Analysis - 17th International Symposium, ATVA 2019, Taipei, Taiwan, October 28-31, 2019, Proceedings*, ser. Lecture Notes in Computer Science, Y. Chen, C. Cheng, and J. Esparza, Eds., vol. 11781. Springer, 2019, pp. 423–441. [Online]. Available: [https://doi.org/10.1007/978-3-030-31784-3\\_25](https://doi.org/10.1007/978-3-030-31784-3_25)
- [16] —, “Controller synthesis for hyperproperties,” in *33rd IEEE Computer Security Foundations Symposium, CSF 2020, Boston, MA, USA, June 22-26, 2020*. IEEE, 2020, pp. 366–379. [Online]. Available: <https://doi.org/10.1109/CSF49147.2020.00033>
- [17] R. Cavada, A. Cimatti, M. Dorigatti, A. Griggio, A. Mariotti, A. Micheli, S. Mover, M. Roveri, and S. Tonetta, “The nuxmv symbolic model checker,” in *CAV*, 2014, pp. 334–342.
- [18] “Voting example from the solidity documentation,” <https://docs.soliditylang.org/en/v0.8.13/solidity-by-example.html#voting>, 2022, accessed: 2022-05-12.
- [19] “Blind auction example from the solidity documentation,” <https://docs.soliditylang.org/en/v0.8.13/solidity-by-example.html#blind-auction>, 2022, accessed: 2022-05-12.
- [20] A. Duret-Lutz, A. Lewkowicz, A. Fauchille, T. Michaud, E. Renault, and L. Xu, “Spot 2.0 — a framework for LTL and  $\omega$ -automata manipulation,” in *Proceedings of the 14th International Symposium on Automated Technology for Verification and Analysis (ATVA'16)*, ser. Lecture Notes in Computer Science, vol. 9938. Springer, Oct. 2016, pp. 122–129.
- [21] D. Suvorov and V. Ulyantsev, “Smart contract design meets state machine synthesis: Case studies,” *CoRR*, vol. abs/1906.02906, 2019. [Online]. Available: <http://arxiv.org/abs/1906.02906>
- [22] S. K. Lahiri, S. Chen, Y. Wang, and I. Dillig, “Formal specification and verification of smart contracts for azure blockchain,” *CoRR*, vol. abs/1812.08829, 2018. [Online]. Available: <http://arxiv.org/abs/1812.08829>
- [23] A. Permenev, D. K. Dimitrov, P. Tsankov, D. Drachslers-Cohen, and M. T. Vechev, “Verx: Safety verification of smart contracts,” in *2020 IEEE Symposium on Security and Privacy, S&P 2020, San Francisco, CA, USA, May 18-21, 2020*. IEEE, 2020, pp. 1661–1677. [Online]. Available: <https://doi.org/10.1109/SP40000.2020.00024>
- [24] J. Stephens, K. Ferles, B. Mariano, S. K. Lahiri, and I. Dillig, “Smartpulse: Automated checking of temporal properties in smart contracts,” in *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 2021, pp. 555–571. [Online]. Available: <https://doi.org/10.1109/SP40001.2021.00085>
- [25] M. R. Clarkson and F. B. Schneider, “Hyperproperties,” in *Proceedings of the 21st IEEE Computer Security Foundations Symposium, CSF 2008, Pittsburgh, Pennsylvania, USA, 23-25 June 2008*. IEEE Computer Society, 2008, pp. 51–65. [Online]. Available: <https://doi.org/10.1109/CSF.2008.7>
- [26] J. O. Gutsfeld, M. Müller-Olm, and C. Ohrem, “Propositional dynamic logic for hyperproperties,” in *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, ser. LIPIcs, I. Konnov and L. Kovács, Eds., vol. 171. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pp. 50:1–50:22. [Online]. Available: <https://doi.org/10.4230/LIPIcs.CONCUR.2020.50>
- [27] B. Finkbeiner and M. Zimmermann, “The first-order logic of hyperproperties,” *CoRR*, vol. abs/1610.04388, 2016. [Online]. Available: <http://arxiv.org/abs/1610.04388>
- [28] N. Coenen, B. Finkbeiner, C. Hahn, and J. Hofmann, “The hierarchy of hyperlogics,” in *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*. IEEE, 2019, pp. 1–13. [Online]. Available: <https://doi.org/10.1109/LICS.2019.8785713>
- [29] J. Baumeister, N. Coenen, B. Bonakdarpour, B. Finkbeiner, and C. Sánchez, “A temporal logic for asynchronous hyperproperties,” in *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part I*, ser. Lecture Notes in Computer Science, A. Silva and K. R. M. Leino, Eds., vol. 12759. Springer, 2021, pp. 694–717. [Online]. Available: [https://doi.org/10.1007/978-3-030-81685-8\\_33](https://doi.org/10.1007/978-3-030-81685-8_33)

- [30] L. Bozzelli, A. Peron, and C. Sánchez, “Asynchronous extensions of hyperltl,” in *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*. IEEE, 2021, pp. 1–13. [Online]. Available: <https://doi.org/10.1109/LICS52264.2021.9470583>
- [31] J. O. Gutsfeld, M. Müller-Olm, and C. Ohrem, “Automata and fixpoints for asynchronous hyperproperties,” *Proc. ACM Program. Lang.*, vol. 5, no. POPL, pp. 1–29, 2021. [Online]. Available: <https://doi.org/10.1145/3434319>
- [32] L. Lamport and F. B. Schneider, “Verifying hyperproperties with TLA,” in *34th IEEE Computer Security Foundations Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25, 2021*. IEEE, 2021, pp. 1–16. [Online]. Available: <https://doi.org/10.1109/CSF51468.2021.00012>
- [33] I. Grishchenko, M. Maffei, and C. Schneidewind, “A semantic framework for the security analysis of ethereum smart contracts,” in *Principles of Security and Trust - 7th International Conference, POST 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings*, ser. Lecture Notes in Computer Science, L. Bauer and R. Küsters, Eds., vol. 10804. Springer, 2018, pp. 243–269. [Online]. Available: [https://doi.org/10.1007/978-3-319-89722-6\\_10](https://doi.org/10.1007/978-3-319-89722-6_10)
- [34] E. Albert, S. Grossman, N. Rinetzky, C. Rodríguez-Núñez, A. Rubio, and M. Sagiv, “Taming callbacks for smart contract modularity,” *Proc. ACM Program. Lang.*, vol. 4, no. OOPSLA, pp. 209:1–209:30, 2020. [Online]. Available: <https://doi.org/10.1145/3428277>