

# Pacta sunt servanda: smart legal contracts in *Stipula*

Silvia Crafa  
crafa@math.unipd.it  
University of Padova  
Italy

Cosimo Laneve  
cosimo.laneve@unibo.it  
University of Bologna – Inria FOCUS  
Italy

Giovanni Sartor  
Giovanni.Sartor@unibo.it  
University of Bologna – European  
University Institute  
Italy

## ABSTRACT

There is a growing interest in running legal contracts on blockchain systems to exploit the advantages of such systems in terms of disintermediation, transparency and immutability. At the same time, it is important to understand to what extent smart contracts programs may capture legal content. To fulfil these needs, we undertake a foundational study of legal contracts and we distill four main features: agreement, permissions, violations and obligations. We therefore design *Stipula*, a domain specific language that assists lawyers in programming legal contracts through specific patterns. The language is based on a small set of abstractions that express the features of legal contracts and that are amenable to be executed on blockchain systems. *Stipula* is equipped with a formal semantics and an observational equivalence, that provide for a clear account of the contracts' behaviour, and allows us to sketch an implementation on top of the Ethereum platform, pinpointing the critical differences between the two settings.

## CCS CONCEPTS

• Theory of Computation; • Models of Computation; • Currency; • Distributed computing models;

## KEYWORDS

Legal contracts, smart contracts, domain-specific language, operational semantics, legal bismulation

### ACM Reference Format:

Silvia Crafa, Cosimo Laneve, and Giovanni Sartor. 2021. Pacta sunt servanda: smart legal contracts in *Stipula*. In ????. ACM, New York, NY, USA, 14 pages. <https://doi.org/????>

## 1 INTRODUCTION

A legal contract is “an agreement which is intended to give rise to a binding legal relationship or to have some other legal effect” [23]. The parties are in principle free to determine the content of their contracts (*party autonomy/freedom of contract*): the law recognizes their intention to achieve the agreed outcomes and secures the enforcement of such outcomes (*legally binding effect*). A contract produces the intended effects, declared by the parties, only if it is legally valid: the law may deny validity to certain

clauses (e.g. excessive interests rate) and/or may establish additional effects that were not stated by the parties (e.g. consumer's power to withdraw from an online sale, warranties, etc.). The intervention of the law is particularly significant when the contractor (usually the weaker party, such as the worker in an employment contract or the consumer in an online purchase) agrees without having awareness of all clauses in the contract, nor having the ability to negotiate them, due to the existing unbalance of power.

Blockchain-based smart contracts have been advocated for digitally encoding legal contracts, so that the execution and enforcement of contractual conditions may occur automatically, without human intervention. This technology enables cutting performance costs and increasing outcomes certainty. Moreover, since the execution of smart contracts is trackable and irreversible, parties that to not trust each other can use them with no intermediation of a trusted third party. On this account, several governments have recently recognised that smart contracts and, more general, programs operating over distributed ledgers may indeed have legal value [2, 4, 6]. However, the assimilation of smart contracts to legally binding contracts, or rather the double nature of smart contracts as computational mechanisms and as legal contract raises both legal and technological issues.

In this paper, after discussing the main issues raised by the literature, we propose a technology that may contribute to addressing them. The overall aim is to facilitate the transparency of smart contracts as well as the mapping of computational operations into legal-institutional outcomes, thus limiting or mitigating the problems concerning the implementation of smart contracts and preventing disputes between the parties.

Specifically, we put forward *Stipula*, a new domain specific language for the creation of smart legal contracts operating over a blockchain. *Stipula* is pivoted on a small number of abstractions that are useful to capture the distinctive elements of legal contracts, that is *permissions*, *prohibitions*, *obligations*, fungible and non fungible *assets exchanges*, risk (*viz. alea*), *escrows* and *securities*. All these normative elements are expressed by a strictly regimented behaviour in legal contracts: permissions and empowerments correspond to the possibility of performing an action at a certain stage, prohibitions correspond to the interdiction of doing an action, while obligations are recast into commitments that are checked at a specific time limit. Moreover, the set of normative elements changes over time according to the actions have been done (or not). To model these changes, *Stipula* commits to a *state-aware programming* style, inspired by the state machine pattern widely used in smart contracts (*c.f.* Solidity [1] and Obsidian [3]). This technique allows one to enforce the intended behaviour by prohibiting, for instance, the invocation of a function before another specific function is called.

A second distinctive feature of *Stipula* is the *event* primitive, a programming abstraction that is used to issue an obligation and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

???, ???, ???

© 2021 Association for Computing Machinery.

ACM ISBN ???...\$???

<https://doi.org/????>

schedule a future statement that automatically executes a corresponding penalty, if the obligation is not met. This allows one to implement legal obligations and commitments in terms of the future execution of a (state-aware) computation at a specific point in time.

A third peculiarity of *Stipula* is the *agreement operator*, which marks that the contract’s parties have reached a consensus on the contractual arrangement they want to create. In legal contracts, this phase corresponds to the subscription of the contract, where there are parties that are going to *set* the contractual conditions and others that *accept* them. Technically, the operation is a multiparty synchronization, which is implemented in current smart contract languages by ad-hoc protocols (see Section 6).

The fourth key feature regards *assets*, which are first-class linear concepts in *Stipula*. In particular, *Stipula* has an explicit, and thus conscious, management of linear resources, such as currency-like values and tokens. The transfer of such resources must preserve the total supply: the sender of the asset must always relinquish the control of the transferred asset. These assets, which are pervasive in the most successful blockchain applications and the induced economy, are necessary in legal contracts, as well: currency is required for payment but also for escrows, and tokens, both fungible and non-fungible, are useful to model securities and provide a digital handle on a physical good. The design choice of explicitly marking asset movements with an ad hoc syntax in *Stipula* promotes a safer, asset-aware, programming discipline that reduces the risk of the so-called double spending, the accidental loss or the locked-in assets. Therefore, without providing the flexibility/expressive power of resource-aware languages like Move [11] or Nomos [10], *Stipula* uses a simple and powerful core of primitives that support the writing of digital legal contracts even by non ICT experts.

In Section 2 we give an interdisciplinary discussion about smart legal contracts, focusing on the interlace between the on-chain and off-chain elements required by the digitalization of juridical acts. The syntax of *Stipula* is formally defined in Section 3 where a simple example – the Bike-Rental contract – is used to describe the concepts. The semantics of *Stipula* is defined in Section 4, sticking to an operational approach that specifies the runtime behaviour of a legal contracts by means of transitions. In Section 5, following a standard technique in concurrency theory [19], we develop an observational equivalence that provides for an equational theory of smart legal contracts. The equivalence is based on a notion of *bisimulation* that is suitable to blockchain-based observations, and that equates contracts differing for hidden elements, such as names of states, and singles out conditions for identifying contracts that send two assets in different order. The study of the implementation of the distinctive elements of *Stipula*, namely agreements, assets and events, on a blockchain systems using a Solidity-like target language is undertaken in Section 6. The specification in *Stipula* of a set of archetypal acts, ranging from renting to (digital) licenses and to bets, is reported in Section 7

We end our contribution by discussing the related work in Section 8 and delivering our final remarks in Section 9. The appendix, which has been added for reviewer sake, contains a technical part.

## 2 SMART LEGAL CONTRACTS

A substantial debate has taken place on whether the parties’ decision to execute a smart contract having certain computational effects may count as legal contract establishing corresponding legal effect, e.g. [12, 16, 18, 20]. A simple answer to this question comes from the principle of “*freedom of form*” in contracts, which is shared by modern legal systems: *parties are free to express their agreement using the language and medium they prefer, including a programming language*. Therefore, by this principle, smart contracts may count as legal contracts.

However, the problem is whether smart contracts preserve the essential elements of legal ones. In this respect, a legal contract is meant to bring about the *institutional effects* intended by the parties, that is establishing new obligations, rights, powers and liabilities between them or to transfer rights (such as rights to property) from one party to the other. These institutional effects are guaranteed by the possibility of activating judicial enforcements. That is, each party may start a lawsuit if he believes that the other party has failed to comply with the contract. In this case, the judge will have to interpret the contract, ascertain the facts of the case, and determine whether there has indeed been a contractual violation. Accordingly, the defaulting party may be enjoined to comply or pay damages. Whether and to what extent we may consider that a smart contract produces judicially enforceable legal effects is more debatable, given that smart contract modify especially in absence of international technical standards and transnational legal frameworks.

We recognise that no easy and comprehensive solution is yet at hand for the issue we have just mentioned. However, we believe that it can be at least mitigated if a strict, and understandable mapping is established between executable instructions and institutional-normative effects. To this aim, we observe that the lifecycle of a legal contract goes through a number of phases: (a) formation and negotiation, (b) contract storage/notarizing, (c) performance, enforcement and monitoring, (d) possible modification (e) dispute resolution, and (f) termination. Software-based solutions can be valuable in all these phases, but the specific features of blockchain-based smart contracts make them convenient only in some of them. For example, the negotiation of contractual conditions might require a degree of privacy that conflicts with that of the blockchain, which naturally runs on transnational infrastructures, thus crossing several, possibly different, legal systems and jurisdictions (c.f. General Data Protection Regulation is valid only in Europe). Similarly the dispute resolution can take advantage of online services, but can hardly be fully ported on-chain. It is also very problematic to amend the behaviour of a running smart contract to match a change in the parties’ will.

On the other hand, the blockchain is perfectly suited to phases (b) and (c). Both the content of the legal contract and the expression of agreement of the parties can be notarized and code can be used to express the contractual clauses and automatically enforce them through the runtime execution, which is verified and recorded by the nodes of the blockchain.

Nevertheless, if smart contracts are legally binding, then it is necessary to ensure that the parties are fully aware of the computational effects of their code. Only in this case, there may be a genuine agreement over the content of the contract. Thus transparency and some degree of readability by contractors that have no or little

computer expertise becomes a key requirement. Relating to this point, we acknowledge that similar problems also exist in natural language contracts, which are usually signed without the parties being aware of all of their clauses and judicial enforcement being too costly or complex to be practicable. Usually, when concluding online purchases of goods or services, most consumers just click the “accept” button, without even trying to read the clauses (which are often lengthy and full of legal jargon).

We believe that *Stipula* provides some important advantages as a language for specifying smart legal contracts. Its primitives are concise and abstract enough to be easily accessible to lawyers. Its formal operational semantics promises that the execution of contracts does not lead to unexpected behaviours and is amenable to automatic verification. To mention a distinctive feature of *Stipula*, the agreement primitive helps to deal with some legal issues, since it clearly identifies the moment when (some) legal effects are triggered and the parties who are involved. For instance, the set of parties involved in the agreement might include an Authority that is charged to monitor off-chain constraints, such as obligations of diligent storage and care, or the obligations of using goods only as intended, taking care of litigations and dispute resolution. Moreover, untrusted players involved in a bet contract can rely on the agreement to explicitly define the data source providing the outcome of the aleatory value associated to the bet.

### 3 THE STIPULA LANGUAGE

*Stipula* features a minimal set of primitives that are primary in legal contracts, such as agreements, field updates, conditional behaviour, timed events, value and asset transfer, functions and states like Finite State Machines (FSMs).

We use countable sets of: *contract names*, ranged over by  $C, C', \dots$ ; names of externally owned accounts, called *parties*, ranged over by  $A, A', \dots$ ; *function names* ranged over  $f, g, \dots$ . Parties represent the users involved in the contract, *i.e.* addresses in blockchain systems. *Assets* and generic contract’s *fields* are syntactically set apart since they have different semantics. Then we assume a countable set of *asset names*, ranged over by  $h, h', \dots$ , and a set of *field names*, ranged over  $x, x', \dots$ . We reserve  $z, z', y, y'$ , for function parameters and  $A, A', \dots$  for parameters that are parties. Finally, we will use  $Q, Q', \dots$ , to range over contract states. A smart legal contract in *Stipula* is written

```

legal_contract C {
  assets h1, ..., hk
  fields x1, ..., xk'
  agreement (A1, ..., An) {
    Ai =SET=> xi1, ..., xiri // i ∈ I
    Aj =OK=> xj1, ..., xjsj // j ∈ J
  } => @Q
  F
}

```

where  $C$  identifies the contract; its body contains assets and fields (without any typed information: *Stipula* is type-free), the *agreement code*, where  $I$  and  $J$  are subsets of  $1..n$ , and  $F$  is a sequence of functions.

The dichotomy between assets and fields is a key design choice of *Stipula*. Indeed, the relevance of first-class resources, *i.e.* linear

values that cannot be copied nor dismissed, is widely acknowledged in DLTs (*c.f.* Move [11] and Nomos [10]) to support a safer asset-aware contract programming. Legal contracts as well manage assets, such as money and tokens granting a digital access to (possibly physical) goods or services. Henceforth the decision to syntactically highlight the differences between operations on values and on assets.

The definition of *who* is going to participate to the contract and *what* are the terms of the contract in an explicit abstraction – the *agreement* – is another distinctive feature of *Stipula*. Technically, the agreement is the *constructor* of the contract, that specifies which parties may set the values for fields – operation =SET=> – and which ones may agree on such values – operation =OK=>. It also specifies the initial state of the contract. We assume a consistency criterium for these operations: (i) a field may be set by at most one party, (ii) parties may agree on fields set by others, and (iii) every party involved in the agreement must either set or agree on a possible empty sequence of fields, *e.g.*  $A =OK=> \emptyset$ . Observe that no asset can be set during the agreement, only fields. It is assumed that fields, assets and parties’ names do not contain duplicates.

Functions  $F$  and their bodies are written according to the syntax in Table 1. The function’s syntax highlights the constraint that only the party  $A$  can invoke the function  $f$ , and only when the contract is in state  $Q$ . Function’s parameters are split in two lists: the *formal parameters*  $z_1, \dots, z_m$  in brackets and the *asset parameters*  $y_1, \dots, y_{m'}$  in square brackets. The *precondition*  $(B)$  is a predicate on parameters of the functions and fields and assets of the contract that constrains the execution of the body of  $f$ . Finally the *body*  $\{S \ W\} \Rightarrow @Q'$  specifies the *statement part*  $S$ , the *event part*  $W$ , and the state  $Q'$  of the contract when the function execution terminates. Function’s parameters are assumed without duplicates, and empty lists of (asset) parameters are shortened by omitting empty parenthesis.

*Statements*  $S$  include the empty statement  $_$  and different types of assignments, followed by a continuation. Assignments use the two symbols  $\multimap$  and  $\rightarrow$  to differentiate updates of assets and of fields, respectively. The syntax of the two operators is taken from [10]. Assignments can be either *local*, that is referring to local fields or local assets, denoted by  $E \rightarrow x$  and  $E \multimap h, h'$ , respectively, or they can be *remote*, denoted by  $E \rightarrow A$  and  $E \multimap h, A$ , defining the sending of a value and an asset, respectively, to the address  $A$ . Asset assignments are ternary operations: the meaning of  $E \multimap h, h'$  is that the value of  $E$  is subtracted to the asset  $h$  and added to the asset  $h'$  – *resources stored in assets can be moved but cannot be destroyed*. The operational semantics will ensure that asset assignments can at most drain an asset, preventing assets with negative values. In the rest of the paper we will always abbreviate assignments such as  $h \multimap h, h'$  and  $h \multimap h, A'$  (which are very usual, indeed) into  $h \multimap h'$  and  $h \multimap A'$ , respectively. Statements also include *conditionals*  $(B)\{S\} \ S'$  that executes  $S$  if the predicate  $B$  is true and continues as  $S'$ .

*Events*  $W$  are sequences of *timed continuations*. A timed continuation is a term  $E \gg @Q\{S\} \Rightarrow @Q'$ , which is triggered at a time  $\mathbb{t}$  that is the value of  $E$ . When triggered, the continuation  $S$  will be executed only if the contract’s state is  $Q$ . At the end of the execution of  $S$ , the contract transits to  $Q'$ . That is, in *Stipula*, the programming abstractions of FSMs are used to schedule the future execution of

```

355 F ::= _ | @Q A: f(z1, ..., zm)[y1, ..., ym'] (B){S W} ⇒ @Q' F
356 S ::= _ | E → x S | E → h, h' S | E → A S | E → h, A S | (B){S} S
357 W ::= _ | E » @Q {S} ⇒ @Q' W
358 E ::= now | X | v | E op E
359 B ::= true | false | E rop E | B && B | B || B | ~B
360
361
362
363
364

```

**Table 1: Syntax of *Stipula*.**

a (state-aware) computation at a specific point in time. We will show that this notion of events is pivotal in the encoding of legal obligations and commitments.

*Expressions*  $E$  can be names of assets, fields and parameters, generically ranged over by  $X$ . The names  $v$  and  $u$  will be reserved for values; they include integer and boolean constants, asset constants like (non-fungible) tokens, strings, and a time datatype. We will be a little liberal with values and operations, generically denoted by  $E$  op  $E$ : they include standard arithmetic operations and operations on tokens, e.g. `use_once(token)` generates a usage-code providing a single access to the service or the good associated to the token asset. The operation  $\mathbb{t} + n$  sums  $n$  seconds to the time value  $\mathbb{t}$ . The identifier `now` stores the present time (when the code is executed). *Boolean expressions*  $B$  are the standard ones, where the operations `rop` are the relational operations (`=`, `>=`, etc.) The set of names occurring in  $E$  will be noted by  $fv(E)$ .

A *Stipula* program  $\mathcal{P}$  is a sequence of smart legal contracts definitions. The contracts are inactive as long as no group of addresses has interest to run them, by invoking the agreement code. We remark that in *Stipula* the code of a contract cannot invoke another contract: we postpone to future work the study of language extensions allowing cross references between legal contracts using inheritance and composition. Therefore, at the moment, since legal contracts are independent, there is no loss of generality in considering a program to be composed by a single smart legal contract definition.

*Example 3.1.* Consider the following simple contract for renting bikes:

```

394 1 legal_contract Bike_Rental {
395 2   assets balance
396 3   fields cost, time_limit, use_code
397 4
398 5   agreement (Lender, Borrower) {
399 6     Lender =SET=> time_limit, cost
400 7     Borrower =OK=> time_limit, cost
401 8   } ⇒ @Inactive
402 9
403 10  @Inactive Lender : offer (z) {
404 11    z → use_code
405 12  } ⇒ @Proposal
406 13
407 14  @Proposal Borrower : accept [y]
408 15    (y == cost) {
409 16      y → balance
410 17      use_code → Borrower
411 18      now + time_limit »
412 19      @Using {
413 20        //end-of-time usage
414 21        "End_Reached" → Borrower
415 22        balance → Lender
416 23      } ⇒ @End
417 24  } ⇒ @Using

```

```

24 @Using Borrower : end {
25   balance → Lender
26 } ⇒ @End
27
28 }

```

**Listing 1: The rent for free contract**

The agreement code specifies that there are two parties – the Lender and Borrower – and that the Lender sets the values for the time of usage (`time_limit`) and the cost, while the Borrower has to agree on such values. Then Lender sends a use-code that is stored in the contract (in the `use_code` field) and is not accessible to Borrower till he pays for the usage. The transition from `Inactive` to `Proposal` at the end of the offer function enables the Borrower to pay for the usage – function `accept` that takes in input an asset – so that the contract sends him the use-code (line 17) and he can use the bike till the time limit. This constraint is expressed by the event in lines 18-22. We have two remarks: first, the payment is not made to Lender but the asset is stored in the contract (in `balance`); second, the event will be triggered when the time expires. In this case a message to Borrower is sent, the payment is transferred to Lender, which will change bike’s use code so that the bike will be locked at the next Borrower’s stop. The function `end` can be invoked by Borrower to terminate the renting before time expires. The legal issues involved in a rent contract will be discussed in Section 7.

## 4 SEMANTICS

The meaning of *Stipula* primitives is defined in an operational way by means of a transition relation. Let  $C(\Phi, \ell, \Sigma, \Psi)$  be a *runtime contract* where

- $C$  is the contract name;
- $\Phi$  is the current state of the contract: it is either  $_$  (for no state) or a contract state  $Q$ ;
- $\ell$  is a mapping from fields and assets to values;
- $\Sigma$  is a possible residual of a function body or of an event handler, i.e.  $\Sigma$  is either  $_$  or a term  $SW \Rightarrow @Q$ ;
- $\Psi$  is a (possibly empty) multiset of *pending events* that have been already scheduled for future execution but not yet triggered. We let  $\Psi$  be  $_$  when there are no pending events, otherwise  $\Psi = W_1 | \dots | W_n$  such that each  $W_i$  is a single event expression (not a sequence), and its *time guard* is an expression that has already been evaluated into a time value  $\mathbb{t}_i$ .

Runtime contracts are ranged over by  $C, C', \dots$ . A *configuration*, ranged over by  $S, S', \dots$ , is a pair  $C, \mathbb{t}$ , where  $\mathbb{t}$  is a global clock. As anticipated, there is no loss of generality in considering programs made of a single running contract; a remark about the extension to

configurations containing several contracts is reported at the end of the section.

The transition relation of a *Stipula* program  $\mathcal{P}$  is  $\mathbb{S} \xrightarrow{\mu} \mathcal{P} \mathbb{S}'$ , where the label  $\mu$  is either empty, or  $A : f(\bar{u})[\bar{v}]$ , or  $v \rightarrow A$ , or  $v \dashv A$ , or  $\alpha$ , where  $\alpha = A_1 R_1 [\bar{x}_1 \mapsto \bar{v}_1] \parallel \dots \parallel A_k R_k [\bar{x}_k \mapsto \bar{v}_k]$ , with  $R_1, \dots, R_k \in \{\text{SET}, \text{OK}\}$ . In the following we will always omit the index  $\mathcal{P}$  because it is considered implicit. The formal definition of  $\mathbb{S} \xrightarrow{\mu} \mathbb{S}'$  is given in Tables 2 and 3, using the following auxiliary predicates and functions:

- $\llbracket E \rrbracket_{\ell}$  is a function that returns the value of  $E$  in the memory  $\ell$ . We omit the definition.
- $\llbracket W \rrbracket_{\ell}$  is the multiset of scheduled events obtained from the sequence  $W$  by replacing every time expression by its value in the time guards. That is  $\llbracket \_ \rrbracket_{\ell} = \_$  and  $\llbracket E \gg \text{@Q}\{S\} \gg \text{@Q}' \ W' \rrbracket_{\ell} = \llbracket E \rrbracket_{\ell} \gg \text{@Q}\{S\} \gg \text{@Q}' \ | \ \llbracket W' \rrbracket_{\ell}$ .
- Let  $\Psi$  be a multiset of pending events, and  $\mathbb{t}$  a time value, then the predicate  $\Psi$ ,  $\mathbb{t} \dashv$  is *true* whenever  $\Psi = \mathbb{t}_1 \gg \text{@Q}_1\{S_1\} \gg \text{@Q}'_1 \ | \ \dots \ | \ \mathbb{t}_n \gg \text{@Q}_n\{S_n\} \gg \text{@Q}'_n$  and, for every  $1 \leq i \leq n$ ,  $\mathbb{t}_i \neq \mathbb{t}$ , *false* otherwise.

Rule [AGREE] in Table 2 establishes the agreement of the parties involved in the contract. This operation is a multiparty synchronization and the label  $\alpha$  is intended to highlight this point. Some parties, namely  $A_i$  with  $i \in I$ , set the initial values of the contract's fields  $\bar{x}_i$ . At the same time the parties  $A_j$ , with  $j \in J$ , agree on this field initialisation – *c.f.* the last premise. The parties, represented in the legal contract by the parameters  $A_1, \dots, A_n$  are instantiated by the actual addresses  $A_1, \dots, A_n$ . We recall the consistency criterium stating that a field may be set by at most one party, the parties may agree on fields that they do not set, and all the  $n$  parties appear in the label  $\alpha$ .

Rule [FUNCTION] defines function invocations; the label specifies the address  $A$  performing the invocation and the function name  $f$  with the actual parameters. The transition may occur provided (i) the contract is the state  $Q$  that admits invocations of  $f$  from  $A$  and (ii) the contract is *idle*, *i.e.* the contract has no statement to execute – *c.f.* the left-hand side runtime contract – (iii) the precondition  $B$  is satisfied, and no event can be triggered – *c.f.* the premise  $\Psi$ ,  $\mathbb{t} \dashv$ . In particular, this last constraint expresses that events *have precedence* on possible function invocations. For example, if a payment deadline is reached and, at the same time, the payment arrives, it will be refused in favour of the event managing the deadline.

Rule [STATE CHANGE] says that a contract changes state when the statements execution terminates and the sequence of events  $W$  is added to the multiset of pending events, up to the evaluations of their time expressions, *i.e.* the occurrences of the identifier now are replaced by the current value of the clock.

Rule [EVENT MATCH] specifies that event handlers may run provided there is no statement to perform in the runtime contract, and the time guard of the event has exactly the value of the global clock  $\mathbb{t}$ . Observe that the timeouts of the events are evaluated in an eager way when the event is scheduled – *c.f.* rule [STATE CHANGE] – not when the event handler is triggered. Moreover, the state change performed at the end of the execution of the event handler is carried over again by the rule [STATE CHANGE], with an empty sequence  $W$ .

Rule [TICK] models the elapsing of time. This happens when the contract has no statement to perform and no event can be triggered. Intuitively, the implementation of *Stipula* on top of a blockchain

will bind the global clock to the timestamp of the current block. Therefore a sequence of semantic transitions performed in the same unit of time will correspond to a set of transactions inserted into the same block to be appended to the blockchain.

It is worth to notice that the foregoing rules imply that the complete execution of a function call does not affect the global time. This admits the paradoxical phenomenon that an endless sequence of function invocations does not make time elapse. While this is possible in theory, it is not in practice, since blocks can only include a finite number of transactions. Additionally, all the legal contracts we have analyzed are finite state, each state admits a single function invocation, and function invocations update the state in a noncircular way, thus preventing infinite sequence of function calls. An alternative choice would be to adjust the semantics so to increment the clock every time a maximal number of functions has been evaluated, thus forcing each block to contain at most a limited number of function invocations. We have preferred to stick to the simpler semantics.

Table 3 defines transitions due to the execution of statements. All these transitions are local to the runtime contract and time does not change. Therefore, for simplicity, we always omit the clock. We only discuss [ASSET\_SEND] and [ASSET\_UPDATE] because the other rules are standard. Rule [ASSET\_SEND] returns part of an asset  $h$  to the party  $A$ . This part, named  $v$ , is removed from the asset, *c.f.* the memory of the right-hand side runtime contract in the conclusion. In a similar way, [ASSET\_UPDATE] moves a part  $v$  of an asset  $h$  to an asset  $h'$ . For this reason, the final memory becomes  $\ell[h \mapsto \ell(h) - v, h' \mapsto \ell(h') + v]$ . We observe that assets, representing physical entities (coins, houses, goods, etc.) are never destroyed. The condition  $\ell(h) \geq v$  in the premises ensures that assets can never become negative.

The semantics of *Stipula* does not consider runtime errors, for instance an attempt to drain too much value from an asset results in a stuck configuration. We postpone to future work a precise account of runtime failures and contract errors, since it requires a deep interdisciplinary analysis of the legal issues involved in the execution of the exceptional cases.

The initial configuration of a *Stipula* program  $\mathcal{P}$  made of a single contract  $C$  is  $C(\_, \emptyset, \_, \_) , \mathbb{t}$ . The contract is *inactive* as long as no group of addresses has interest to run it, *c.f.* rule [AGREE]. The global clock can be any value, because it corresponds to the absolute time, defined by the timestamp of the current block in the blockchain system.

*Example 4.1.* Possible initial transitions of the *Bike\_Rental* contract in Example 3.1 are reported in Table 4. We assume that the actual names of parties are the same as the formal names (therefore we omit the mappings in the memories).

Let be  $\alpha = \text{Lender SET} [\text{time\_limit} \mapsto 3600, \text{cost} \mapsto 2] \parallel \text{Borrower OK} [\text{time\_limit} \mapsto 3600, \text{cost} \mapsto 2]$  and  $\ell = [\text{cost} \mapsto 2, \text{time\_limit} \mapsto 3600]$  (*i.e.* the cost of renting a bike is 2 euro per hour – the time is measured in seconds. Let also be  $\ell' = \ell[z \mapsto 123, \text{use\_code} \mapsto 123]$  and  $S \ W$  be the body of the function *accept*.

To sum up, a legal contract behaves as follows:

- (1) the first action is always an agreement, which moves the contract to an idle state;

591		650
592	[AGREE]	651
593	$\text{agreement}(A_1, \dots, A_n) \{ A_i = \text{SET} \Rightarrow \bar{x}_i \dots A_j = \text{OK} \Rightarrow \bar{x}_j \}^{i \in I, j \in J} \Rightarrow @Q \in C$ $\alpha = A_i \text{ SET } [\bar{x}_i \mapsto \bar{v}_i] \parallel \dots \parallel A_j \text{ OK } [\bar{x}_j \mapsto \bar{v}_j] \quad [\bar{x}_i \mapsto \bar{v}_i]^{i \in I} = [\bar{x}_j \mapsto \bar{v}_j]^{j \in J}$	652
594	$\frac{}{C(-, \emptyset, -, -), \mathbb{t} \xrightarrow{\alpha} C(Q, [A_i \mapsto A_i, \bar{x}_i \mapsto \bar{v}_i]^{i \in I}, -, -), \mathbb{t}}$	653
595	[FUNCTION]	654
596	$@Q A : f(\bar{z}) [\bar{y}] (B) \{ S W \} \Rightarrow @Q' \in C$	655
597	$\Psi, \mathbb{t} \mapsto$	656
598	$\frac{\ell(A) = A \quad \ell' = \ell[\bar{z} \mapsto \bar{u}, \bar{y} \mapsto \bar{v}] \quad \llbracket B \rrbracket_{\ell'} = \text{true}}{C(Q, \ell, -, \Psi), \mathbb{t} \xrightarrow{A: f(\bar{u})[\bar{v}]} C(Q; \ell', SW \Rightarrow @Q', \Psi), \mathbb{t}}$	657
599	[STATE CHANGE]	658
600	$\frac{\llbracket W \{ \mathbb{t} / \text{now} \} \rrbracket_{\ell} = \Psi'}{C(Q, \ell, - W \Rightarrow @Q', \Psi), \mathbb{t} \longrightarrow C(Q', \ell, -, \Psi'   \Psi), \mathbb{t}}$	659
601	[EVENT MATCH]	660
602	$\frac{\Psi = \mathbb{t} \gg @Q \{ S \} \Rightarrow @Q'   \Psi'}{C(Q, \ell, -, \Psi), \mathbb{t} \longrightarrow C(Q, \ell, S \Rightarrow @Q', \Psi'), \mathbb{t}}$	661
603	[TICK]	662
604	$\frac{\Psi, \mathbb{t} \mapsto}{C(Q, \ell, -, \Psi), \mathbb{t} \longrightarrow C(Q, \ell, -, \Psi), \mathbb{t} + 1}$	663

Table 2: The transition relation of *Stipula* – Part 1

609			664
610	[VALUE_SEND]	[ASSET_SEND]	665
611	$\frac{\llbracket E \rrbracket_{\ell} = v \quad \ell(A) = A}{C(Q, \ell, E \rightarrow A \Sigma, \Psi) \xrightarrow{v \rightarrow A} C(Q, \ell, \Sigma, \Psi)}$	$\frac{\llbracket E \rrbracket_{\ell} = v \quad \ell(h) \geq v \quad \ell(A) = A}{C(Q, \ell, E \rightarrow h, A \Sigma, \Psi) \xrightarrow{v \rightarrow A} C(Q, \ell[h \mapsto \ell(h) - v], \Sigma, \Psi)}$	666
612	[FIELD_UPDATE]	[ASSET_UPDATE]	667
613	$\frac{\llbracket E \rrbracket_{\ell} = v}{C(Q, \ell, E \rightarrow x \Sigma, \Psi) \longrightarrow C(Q, \ell[x \mapsto v], \Sigma, \Psi)}$	$\frac{\llbracket E \rrbracket_{\ell} = v \quad \ell(h) \geq v}{C(Q, \ell, E \rightarrow h, h' \Sigma, \Psi) \longrightarrow C(Q, \ell[h \mapsto \ell(h) - v, h' \mapsto \ell(h') + v], \Sigma, \Psi)}$	668
614	[COND_TRUE]	[COND_FALSE]	669
615	$\frac{\llbracket B \rrbracket_{\ell} = \text{true}}{C(Q, (B)\{S\} \Sigma, \Psi) \longrightarrow C(Q, \ell, S \Sigma, \Psi)}$	$\frac{\llbracket B \rrbracket_{\ell} = \text{false}}{C(Q, (B)\{S\} \Sigma, \Psi) \longrightarrow C(Q, \ell, \Sigma, \Psi)}$	670
616	[COND_TRUE]	[COND_FALSE]	671
617	[COND_TRUE]	[COND_FALSE]	672
618	[COND_TRUE]	[COND_FALSE]	673
619	[COND_TRUE]	[COND_FALSE]	674
620	[COND_TRUE]	[COND_FALSE]	675
621	[COND_TRUE]	[COND_FALSE]	676
622	[COND_TRUE]	[COND_FALSE]	677
623	[COND_TRUE]	[COND_FALSE]	678
624	[COND_TRUE]	[COND_FALSE]	679

Table 3: The transition relation of *Stipula* – Part 2

<p>Bike_Rental(-, <math>\emptyset</math>, -, -), <math>\emptyset</math></p> <p style="padding-left: 20px;">Lender: offer<sup>(123)</sup></p> <p style="padding-left: 20px;">Borrower: accept<sup>[2]</sup></p>	<p style="text-align: center;"><math>\xrightarrow{\alpha}</math></p> <p>Bike_Rental(Inactive, <math>\ell</math>, -, -), <math>\emptyset</math></p> <p>Bike_Rental(Inactive, <math>\ell</math>, -, -), 1</p> <p>Bike_Rental(Inactive, <math>\ell[z \mapsto 123]</math>, <math>z \rightarrow \text{use\_code} \Rightarrow @\text{Proposal}</math>, -), 1</p> <p>Bike_Rental(Inactive, <math>\ell'</math>, - <math>\Rightarrow @\text{Proposal}</math>, -), 1</p> <p>Bike_Rental(Proposal, <math>\ell'</math>, -, -), 1</p> <p>Bike_Rental(Proposal, <math>\ell'</math>, -, -), 2</p> <p>Bike_Rental(Proposal, <math>\ell'</math>, -, -), 3</p> <p>Bike_Rental(Proposal, <math>\ell'[y \mapsto 2]</math>, <math>S W \Rightarrow @\text{End}</math>, -), 3</p>	<p style="text-align: right;">[AGREE]</p> <p style="text-align: right;">[TICK]</p> <p style="text-align: right;">[FUNCTION]</p> <p style="text-align: right;">[FIELD_UPDATE]</p> <p style="text-align: right;">[STATE_CHANGE]</p> <p style="text-align: right;">[TICK]</p> <p style="text-align: right;">[TICK]</p> <p style="text-align: right;">[FUNCTION]</p>
--	--	--

Table 4: Initial transitions of *Bike\_Rental*

- (2) in an idle state, fire any ready event with a matching state. If there is one, execute its handler until the end, which is an idle state;
- (3) if there is no event to be triggered in an idle state, *either* tick or call any permitted function (*i.e.* with matching state and preconditions). A function invocation amounts to execute its body until the end, which is again an idle state.

Therefore, we observe that *Stipula* has three sources of non-determinism: (i) the order of the execution of ready event handlers, (ii) the order of the calls of permitted functions, and (iii) the delay of permitted function calls to a later time (thus, possibly, after other

event handlers). For example, the contract  $C$  with two functions  $@Q$   $A: f \{ \_ \} \Rightarrow @Q$  and  $@Q' A': g \{ \_ \} \Rightarrow @Q'$  behaves as either  $\xrightarrow{A: f} \xrightarrow{A': g}$  or  $\xrightarrow{A': g} \xrightarrow{A: f}$ , where  $\xrightarrow{\_}$  are transitions that make the time elapse (rule [TICK]). As another example, consider a contract  $C'$  with a function  $@Q A: f \{ \_ \text{ now } \gg @Q' \{ \text{hello} \rightarrow A \} \Rightarrow @Q' \} \Rightarrow @Q$  and a function  $@Q' A': g \{ \_ \} \Rightarrow @Q'$ . Then it may behave as either  $\xrightarrow{A: f} \xrightarrow{A': g} \xrightarrow{\text{hello}} \xrightarrow{A}$  or as  $\xrightarrow{A: f} \xrightarrow{A': g}$ , after which the action  $\text{hello} \rightarrow A$  is disabled, or as  $\xrightarrow{A': g}$ , which precludes the call of  $f$ .

*Remark.* The semantics of *Stipula* may be easily extended to configurations with several smart legal contracts. It is sufficient to

consider configurations as consisting of sets of runtime contracts and to change the rule [Tick]. To illustrate the general case, let  $C = C_1(\Phi_1, \ell_1, \Sigma_1, \Psi_1), \dots, C_n(\Phi_n, \ell_n, \Sigma_n, \Psi_n)$ . We define  $Events(C) \stackrel{\text{def}}{=} \Psi_1 \mid \dots \mid \Psi_n$ . The new rule [Tick] becomes

$$\frac{[Tick^*] \quad Events(C), \mathbb{t} \rightarrow}{C, \mathbb{t} \longrightarrow C, \mathbb{t} + 1}$$

## 5 STIPULA LAWS AND EQUATIONAL THEORY

The operational semantics of Tables 2 and 3 is too intensional because it defines the behaviour of a legal contract without showing any evidence of the differences between contracts. Nevertheless, it is the base for defining a more appropriate, extensional semantics using a standard technique based on observations [19]. According to this technique, two contracts cannot be separated if a party using them cannot distinguish one from the other. Said otherwise, a party can differentiate two contracts if he can *observe different interactions*. It turns out that defining the observations is a critical point of the overall technique, because it allows to fine-tune the discriminating power of the extensional semantics. The appropriate observations for smart legal contracts match the design principles of *Stipula*: let  $A$  be a party, then

- $A$  should observe the *agreement* that he is going to sign, because stipulating a different agreement may be unsatisfactory; therefore we envisage the observations  $A \text{ SET } [x \mapsto v]$  and  $A \text{ OK } [x \mapsto v]$ ;
- $A$  should also observe the *permission* or the *prohibition* to invoke a functionality at a given time  $\mathbb{t}$ , *i.e.* whether  $A : f(\vec{v})[\vec{w}]$  is possible at  $\mathbb{t}$  or not;
- $A$  should finally observe whether at  $\mathbb{t}$  he can *receive* a value or an asset, *i.e.* whether  $v \rightarrow A$  or  $v \dashv\rightarrow A$  are possible at  $\mathbb{t}$  or not.

The ordering of invocations and receives can be safely overlooked, as long as they belong to the same block of transactions, that is they are executed at the same global time. Notice also that the above observations allow a party to observe contract's *obligations*. Indeed, by shifting the observation at a specific point in time, one can observe the effects of executing the event that encodes a legal commitment, such as the issue of a sanction or the impossibility to do further actions. On the other hand, the foregoing notion of observations abstracts away from the names of the contract's assets and internal states.

We will use the following notations:

- Let  $R$  be either SET or OK, then we write  $A R [x \mapsto v] \in \alpha$  if  $\alpha$  contains  $A R [\vec{z} \mapsto \vec{w}]$  and  $[x \mapsto v]$  occurs in the sequence of assignments  $[\vec{z} \mapsto \vec{w}]$ . We write  $\alpha \sim \alpha'$  whenever  $A R [x \mapsto v] \in \alpha$  if and only if  $A R [x \mapsto v] \in \alpha'$ . Intuitively,  $\alpha$  and  $\alpha'$  express the same agreement up to reordering of the field assignments.
- Let  $\beta \xrightarrow{\mu} \stackrel{\text{def}}{=} \beta \xrightarrow{\mu} \beta$ , where  $\xrightarrow{\mu}$  stands for any number of  $\longrightarrow$  transitions, possibly zero.

The following definition of legal contract equivalence compares the observable behavior of contracts. It is defined over configurations, so to appropriately shift the time of the contract's observations. The equivalence is defined as a suitable bisimulation game

that is consistent with the idea that in blockchain systems the interactions are batched in blocks of transactions.

**Definition 5.1 (Legal Bisimulation).** A symmetric relation  $\mathcal{R}$  is a legal bisimulation between two configurations at time  $\mathbb{t}$ , written  $C_1, \mathbb{t} \mathcal{R} C_2, \mathbb{t}$ , whenever

- (1) if  $C_1, \mathbb{t} \xrightarrow{\alpha} C'_1, \mathbb{t}$  then  $C_2, \mathbb{t} \xrightarrow{\alpha'} C'_2, \mathbb{t}$  for some  $\alpha'$  such that  $\alpha \sim \alpha'$  and  $C'_1, \mathbb{t} \mathcal{R} C'_2, \mathbb{t}$ ;
- (2) if  $C_1, \mathbb{t} \xrightarrow{\mu_1} \dots \xrightarrow{\mu_n} C'_1, \mathbb{t} \longrightarrow C'_1, \mathbb{t} + 1$  then there exist  $\mu'_1 \dots \mu'_n$  that is a permutation of  $\mu_1 \dots \mu_n$  such that  $C_2, \mathbb{t} \xrightarrow{\mu'_1} \dots \xrightarrow{\mu'_n} C'_2, \mathbb{t} \longrightarrow C'_2, \mathbb{t} + 1$  and  $C'_1, \mathbb{t} + 1 \mathcal{R} C'_2, \mathbb{t} + 1$ .

Let  $\simeq$  be the largest legal bisimulation, called *bisimilarity*. When the initial configurations of contracts  $C$  and  $C'$  are bisimilar, we simply write  $C \simeq C'$ .

Being a symmetric relation, a legal bisimulation compares both contracts' permissions and prohibitions: if  $C$  permits an action (*i.e.* exhibits an observation), then  $C'$  must permit the same action, and if  $C$  prohibits an action (*i.e.* does not exhibit a function call or an external communication), then also  $C'$  must not exhibit the corresponding observation. Moreover, the bisimulation game enforces a *transfer property*, that is it shifts the time of observation to the future, so to capture and compare the changes of permissions/prohibitions and the (future) obligations. Observe that the equivalence abstracts away the ordering of the observations within the same time clock, since in a blockchain there is no strong notion of ordering between the transactions contained in the same block. Nevertheless, specific orderings of function invocations are important in *Stipula* contracts and the equivalence cannot overlook essential precedence constraints. For instance, the requirement that a function delivering a service can only be invoked after another specific function, say a payment. This is indeed the case for the legal bisimulation. To explain, consider the contract  $C$  with two functions  $@Q A : f\{-\} \Rightarrow @Q$  and  $@Q A' : g\{-\} \Rightarrow @Q$  and the contract  $C'$  with two functions  $@Q A : f\{-\} \Rightarrow @Q'$  and  $@Q' A' : g\{-\} \Rightarrow @Q$ . In  $C$  the functions can be called in any order, while in  $C'$  the function  $g$  can be invoked only after  $f$ . Accordingly,  $C \not\approx C'$  since there is a runtime configuration of  $C$  that exhibits the observation  $\xrightarrow{A':g}$ , while it is not the case for the contract  $C'$ , since at any time it can only exhibit  $\xrightarrow{A:f} \xrightarrow{A':g}$ .

The following theorem highlights the property that the internal state of the contract is abstracted away by the extensional semantics, which only observes the external contract's behavior. The proof is omitted because it is standard.

**THEOREM 5.2 (INTERNAL REFACTORING).** *Let  $C$  and  $C'$  be two contracts that are equal up-to a bijective renaming of states. Then  $C \simeq C'$ . Similarly, for bijective renaming of assets and contract names.*

Bisimilarity is also independent from *future* clock values. This allows us to garbage-collect events that cannot be triggered anymore because the time for their scheduling is already elapsed.

**THEOREM 5.3 (TIME SHIFT).**

- (1) If  $C, \mathbb{t} \simeq C', \mathbb{t}$  and  $\mathbb{t} \leq \mathbb{t}'$ , then  $C, \mathbb{t}' \simeq C', \mathbb{t}'$ .
- (2) If  $\mathbb{t} < \mathbb{t}'$  then  $C(Q, \ell, \Sigma, \Psi \mid \mathbb{t} \gg @Q \{S\} \Rightarrow @Q')$ ,  $\mathbb{t}' \simeq C(Q, \ell, \Sigma, \Psi)$ ,  $\mathbb{t}'$ .

We finally put forward a set of algebraic laws that formalize the fact that the ordering of remote communications can be safely overlooked, as long as they belong to the same transaction. The laws are defined over statements, therefore, let  $C[\ ]$  be a *context*, that is a contract that contains an hole where a statement may occur. We write  $S \simeq S'$  if, for every context  $C[\ ]$ ,  $C[S] \simeq C[S']$ .

**THEOREM 5.4.** *The following non-interference laws hold in Stipula (whenever they are applicable, we assume  $x \notin \text{fv}(E')$  and  $x' \notin \text{fv}(E)$  and  $h \notin \text{fv}(E')$  and  $h' \notin \text{fv}(E)$  and  $h'' \notin \text{fv}(E)$  and  $h''' \notin \text{fv}(E)$ ):*

$$\begin{aligned}
 E \rightarrow A \ E' \rightarrow A' &\simeq E' \rightarrow A' \ E \rightarrow A \\
 E \rightarrow x \ E' \rightarrow A &\simeq E' \rightarrow A \ E \rightarrow x \\
 E \rightarrow x \ E' \rightarrow x' &\simeq E' \rightarrow x' \ E \rightarrow x \\
 E \multimap h, A \ E' \rightarrow A' &\simeq E' \rightarrow A' \ E \multimap h, A \\
 E \multimap h, A \ E' \rightarrow x' &\simeq E' \rightarrow x' \ E \multimap h, A \\
 E \multimap h, h' \ E' \rightarrow A &\simeq E' \rightarrow A \ E \multimap h, h' \\
 E \multimap h, h' \ E' \rightarrow x' &\simeq E' \rightarrow x' \ E \multimap h, h' \\
 E \multimap h, A \ E' \multimap h'', A' &\simeq E' \multimap h'', A' \ E \multimap h, A \\
 E \multimap h, A \ E' \multimap h'', h''' &\simeq E' \multimap h'', h''' \ E \multimap h, A \\
 E \multimap h, h' \ E' \multimap h'', h''' &\simeq E' \multimap h'', h''' \ E \multimap h, h'
 \end{aligned}$$

## 6 IMPLEMENTATION

The semantics of Section 4 has not only a theoretical interest but it also drives us in the design of the implementation of *Stipula*. While a full discussion about this subject is out of the scope of the present work, we consider the three features that make the implementation of *Stipula* not trivial: agreements, assets, and events. (In this paper, we are overlooking issues regarding errors and backtracks.) Below, we analyze how they can be actually encoded in a lower level, Solidity-like language, that can be more directly implemented in the bytecode of some virtual machine running a mainstream smart contract language (e.g. Ethereum Virtual Machine [25], Move Virtual Machine [11]).

*Agreements.* The agreement code is technically a *multiparty synchronization* that expresses a consensus between the parties to start the contract with particular values of the (non-linear) fields. This construct can be implemented by resorting to a barrier-like protocol, where each party  $A_i$  may call, in whatever order, a specific function to propose the values he agrees on, and the barrier eventually checks the consistency of the proposed values before moving the contract to the initial state. The following snippet of Solidity-like code corresponds to the agreement code (see Section 3), where we assume that the fields of the contract are  $x_1, \dots, x_k$  with types  $T_1, \dots, T_k$ , respectively.

```

address A1, ... , An ;
T1 x1; ... ; Tk xk ;
enum State {Nothing, Q}
State state = State.Nothing ;
int counter = 1 ;

Tj1 auxi_xj1; ... ; Tjsi auxi_xjsi ; // 1 ≤ i ≤ n
bool use_oncei = true ; // 1 ≤ i ≤ n

function set_oki(Ti1 zi1, ..., Tir_i zir_i, Tj1 zj1, ..., Tjsi zjsi){
// 1 ≤ i ≤ n
    if (sender == Ai && use_oncei){
        use_oncei = false ;
        xi1 = zi1 ; ... ; xir_i = zir_i ;
        auxi_xj1 = zj1; ... ; auxi_xjsi = zjsi ;
        if (counter == n){
            if (∧_{1≤i≤n} (auxi_xj1 == xj1 && ... && auxi_xjsi == xjsi))

```

```

        ) state = State.Q ;
        else throw error ;
    } else counter = counter + 1 ;
} else throw error ;
}

```

Each function `set_oki` can be called only once by the party  $A_i$ , with two lists of parameters: the first  $r_i$  values are used to set the contract's fields  $x_{i_1}, \dots, x_{i_{r_i}}$ . The last  $s_i$  values are recorded into the auxiliary fields  $aux_{i-x_{j_1}}, \dots, aux_{i-x_{j_{s_i}}}$ , to express that  $A_i$  agrees *ex ante* with anyone setting the contract's fields  $x_{j_1}, \dots, x_{j_{s_i}}$  to the values  $z_{j_1}, \dots, z_{j_{s_i}}$ . When all the parties have done the agreement, *i.e.* `counter` is equal to  $n$ , a check on the consistency of `auxi_xh` is performed and, in case it succeeds, the contract becomes active moving to the state `Q`. The snippet also shows that contract's states can be easily mapped to enumerations, as usual in the Solidity State Machine pattern.

There is a discrepancy between the above code and the semantics of the agreement in Table 2. While the agreement has a transactional nature (it may occur as a whole or not), the above Solidity protocol takes time, *i.e.* it is performed in several blocks and, in any block, a failure may occur. In this case, there is a backtrack to the *initial state of the block* and not to the initial state of the protocol, as it happens in *Stipula*. Therefore, the error management should take care of removing partial values stored in the fields of the contract. Nevertheless, another source of discrepancy seems more awkward: the gas consumption. In fact, the successful termination of the agreement as well as the failed one have a cost in Ethereum, while it is not the case. To bridge this gap, we should design agreements with fees payed by parties that are used for the consensus. We have not yet studied these details, which are postponed to future work.

*Assets.* Assets are linear resources that cannot be duplicated or leaked: when a resource value is assigned, the location previously holding the value is emptied. The community of smart contracts has already recognized the relevance of having linear resources as first-class entities in programming languages because they can significantly simplify programming and the effort required for verification. *Stipula* features a simple abstraction to manage assets, which is used to represent both currency and indivisible tokens. To implement these assets we can resort to the popular token standards on Ethereum (ERC-20 for virtual currency and ERC-721 for non fungible tokens[13]). Alternatively, we can rely on the Move language, whose designers have featured programmable linear resources by constraining them to adhere to ad-hoc rules specified by its declaring module [7, 11]. Using a pseudo-code inspired to Move, we might define (divisible) assets  $h$  and  $h'$  as resources of type  $H$ , so that the operations  $E \multimap h, h'$  and  $E \multimap h, A$  may be encoded by `h.move(E, h')` and `h.withdraw(E, A)`, according to the definitions below:

```

resource H {
    T amount ;

    function move(T x, H h) {
        (x <= amount){ h = h + x ; amount = amount - x ; }
    }
    function withdraw(T x, address A) {
        (x <= amount){ A.send(x) ; amount = amount - x ; }
    }
    constructor(T x){ amount = x ; }
}

```



Events. The current blockchain technologies do not admit the record of statements that have to be performed afterwards in a future transaction (with one exception to our knowledge: Cardano, see Section 8). The standard technique adopted in Ethereum to circumvent this limitation is based on the Solidity's *event* construct and off-chain *oracle* services. To explain, the *Stipula* event  $E \gg @Q$   $\{ S \} \Rightarrow @Q'$  emitted by a function `foo` can be mapped to the Solidity code:

```

event R(uint time, address lc) ;
function call_back_R() external {
    if (state == State.Q) { S ; state = State.Q'; }
}
function foo(T1 u1, ..., Tn un){
    ...
    emit R(E, address(this)) ;
}

```

The Solidity function `foo` emits an event named `R` carrying the time `E` and the address of the issuing contract. Moreover, an external *DApp service* – an oracle – scans the blockchain looking for `R` events and calls the `call_back_R` function of the contract at the appropriate time `E`. Other more complex and safer protocols can be used. However a code external to the Ethereum blockchain is always necessary in order to trigger the scheduled event handler. A more satisfactory implementation of the *Stipula* events might adopt ideas taken from the implementation of timeouts in the Marlowe and Findel languages for financial contracts [21, 24] as discussed in Section 8.

## 7 EXPRESSIVITY OF STIPULA

*Stipula* has been devised for writing legal contracts in a formal and intelligible (to lawyers) way. In this section we analyze the expressivity of *Stipula* by writing the contracts for a set of archetypal acts, ranging from renting to (digital) licenses and to bets. We conclude this analysis with a table that summarizes the legal elements of the archetypal acts and the programming abstractions that we have used to express them in *Stipula*.

### 7.1 The free rent contract

The free rent is the simplest kind of legal contract. It involves two parties, the lender and the borrower, which initially agree about what good is rented, what use should be made of it, the time limit (or in which case it must be returned), the estimated of value and any defects in the good. Upon agreement, the delivery of the good triggers the legal bond, that is the borrower has the permission to use the good and the lender has the prohibition of preventing him from doing so. Note that there is no transfer of ownership, but only the right to use the good. The contract terminates either when the borrower returns the good, or when the time limit is reached. Litigations could arise when the borrower violates the obligations of diligent storage and care, the obligations of using the good only as intended, and not granting the use to a third party without the lender's consent. In these cases the lender may demand the immediate return of the object, in addition to compensation for the damage. On the other hand, the borrower is entitled to compensation if the good has defects that were known to the lender but that he did not initially disclose.

The free rent contract puts forward the following points:

- When a legal contract refers to a *physical* good, the smart contract needs a digital handle (an avatar) for that good. Many technological solutions, such as smart locks of IoT devices, are actually available. In *Stipula* we abstract from the specific nature of such a digital handle, and we simply represent it as an asset, which intuitively corresponds to a non fungible token associated to the physical good.
- The rent legal contract grants just the *usage* of a good without the transfer of ownership. Accordingly, while the communication of the token provides full control of the associated physical good, we assume an operation uses(token) (resp. use\_once(token) or uses(token,A)) that generates a usage-code providing access to the object associated to the token (resp. a usage-code only valid (once) for the party A).
- In a legal rent contract it is important to acknowledge the delivery of the good, since this is the action that triggers the legal bonds. We rely on assets and their *semantics* to implement this feature.

The *Stipula* code for the free rent of a locker is written in Listing 2. The two parties agrees on the time limit for the locker usage (`time_limit`) and the time limit to start the usage (`time_start`). Contract's states allow one sequence of actions: first Lender sends the number `n` of the locker and the token `t` associated to `n` by calling `box_proposal` (line 10). This action moves the contract to the temporary state `Proposal` and schedule the event in line 13. This event is essential to prevent the unique token associated to the locker to be indefinitely locked-in in the smart contract when the borrower never calls the `boxUse` function to finalize the delivery of the good. If Borrower calls the function `boxUse` (line 16) within the timeout `time_start`, then the number of the rented locker and the access code are returned. At the same time, a second timeout is installed to check the time limit for the locker usage, and the final state change to `Using` (line 21) vanishes the timeout installed in line 13. This second event is needed to prevent a never ending use of the locker. If `time_limit` is reached and the contract's state is still `Using`, then (lines 18-19) a message is sent to Borrower and the token is sent back to Lender, which becomes again in full control of the locker and can thus invalidate the access code held by the borrower. Otherwise, the rent contract can terminate because the borrower explicitly returns the good before the time limit. This is represented by a call of the function `returnBox` (line 23).

```

1 legal_contract Free_Rent {
2   assets token
3   fields numBox, t_start, t_limit
4
5   agreement (Lender, Borrower) {
6     Lender =SET=> t_start, t_limit
7     Borrower =OK=> t_start, t_limit
8   } => @Inactive
9
10  @Inactive Lender : boxProposal (n)[t] {
11    t -> token
12    n -> numBox
13    now + t_start >> @Proposal {
14      token -> Lender } => @End
15  } => @Proposal
16
17  @Proposal Borrower : boxUse {
18    (uses(token), numBox) -> Borrower
19    now + t_limit >> @Using {
20      "Time_Limit_Reached" -> Borrower
21  }
22
23  }

```

```

106321         token ← Lender
106422     } ⇒ @End
106523 } ⇒ @Using
106624
106625 @Using Borrower : returnBox {
106726     token ← Lender
106827 } ⇒ @End
106928 }
    
```

**Listing 2: The rent for free contract**

The smart legal contract of Listing 2 does not consider the compensations that would be needed to deal with the disputes due to breaches of the contract, such as the borrower’s diligent care during the locker’s usage. These violations require off-chain monitoring and a dispute resolution mechanism. The next example illustrates how this off-chain monitoring and enforcement can be combined with the on-chain code.

## 7.2 The Digital Licensee contract

Let us consider a contract corresponding to a licence to access a digital service, like a software or an ebook: the digital service can be freely accessed for a while, and can be permanently bought with an explicit communication within the evaluation period (for a similar example, see [16]). The licensing contractual clauses can be described as follows:

**Article 1.** Licensor grants Licensee for a licence to evaluate the Product and fixes (i) the *evaluation period* and (ii) the *cost* of the Product if Licensee will bought it.

**Article 2.** Licensee will pay the Product in advance; he will be reimbursed if the Product will not be bought with an explicit communication within the evaluation period. The refund will be the 90% of the cost because the 10% is payed to the Authority (see Article 3).

**Article 3.** Licensee must not publish the results of the evaluation during the evaluation period and Licensor must reply within 10 hours to the queries of Licensee related to the Product; this is supervised by Authority that may interrupt the licence and reimburse either Licensor or Licensee according to whom breaches this agreement.

**Article 4.** This license will terminate automatically at the end of the evaluation period, if the licensee does not buy the product.

Compared to the previous example, this contract involves payment and refund: an amount of currency is escrowed, and two parts of it will be sent to different parties, the Authority and either the Licensor or the Licensee. *Stipula* provides the general asset abstraction, together with a general operation to move just a (positive) subset of the asset to a different owner. This is exactly what is needed to deal with currency, therefore the *Stipula* licence contract holds two different assets: an indivisible token providing an handle to the digital service, and a balance that is a divisible asset corresponding to the amount of currency kept in custody inside the smart contract.

A further important feature of the contract is Article 3 that defines specific constraints about the *off-chain* behaviour of Licensor and Licensee. This exemplifies the very general situations where contract’s violations cannot be fully monitored by the on-chain software, such as the publication of a post in a social network, or the leakage of a secret password, or the violation of the obligation

of diligent storage and care. In all these cases, it is required a trusted third party, say an Authority, to supervise the disputes occurring from the off-chain monitoring and to provide a trusted dispute resolution mechanism. The code in Listing 3 illustrates the encoding of the off-chain monitoring and enforcement mechanism with the on-chain smart contract code in *Stipula*.

```

1  legal_contract Licence {
2  assets token, balance
3  fields cost, t_start, t_limit
4
5  agreement (Licensor, Licensee, Authority) {
6  Licensor =SET=> cost, t_start, t_limit
7  Licensee =OK=> cost, t_start, t_limit
8  Authority =OK=> _
9  } ⇒ @Inactive
10
11 @Inactive Licensor : offerLicence [t] {
12 t ← token
13 now + t_start >> @Proposal {
14 token ← Licensor } ⇒ @End
15 } ⇒ @Proposal
16
17 @Proposal Licensee : activateLicence [b]
18 (b == cost){
19 b ← balance
20 balance*0,1 ← balance, Authority
21 uses(token, Licensee) → Licensee
22 now + t_limit >> @Trial {
23 balance ← Licensee
24 token ← Licensor
25 } ⇒ @End
26 } ⇒ @Trial
27
28 @Trial Licensee : buy {
29 balance ← Licensor
30 token ← Licensee
31 } ⇒ @End
32
33 @Trial Authority : compensateLicensor {
34 balance ← Licensor
35 token ← Licensor
36 } ⇒ @End
37
38 @Trial Authority : compensateLicensee {
39 balance ← Licensee
40 token ← Licensor;
41 } ⇒ @End
42 }
    
```

**Listing 3: The contract for a digital licence**

The agreement of Listing 3 involves three parties: Licensor, which fixes the parameters of the contract, according to Article 1., Licensee, which explicitly agrees, and Authority, which does not need to agree upon the contracts’ parameters (*i.e.* the empty-set agreement noted  $\_$ ), but it is important that it is involved in the agreement synchronization, because it plays the role of the trusted third party that is entitled to call the functions `compensateLicensor` and `compensateLicensee`.

In `activateLicence`, the caller, *i.e.* the Licensee, is required to send an amount of assets equal to the fixed cost of the license. Notice the difference between line 19 and line 20: the first one is the move of a fraction of asset towards the authority, while the second is the simple communication to Licensee of a personal usage code associated to the token. Once entered in the `Trial` state, the contract can terminate in three ways: (i) the licensee

expresses its willingness to buy the licence by calling the function `call` which grants him the full token, or (ii) the time limit for the free evaluation period is reached, thus the scheduled event refunds the licensee and gives the token back to the licensor, or (iii) during the evaluation period a violation to Article 3 is identified and the authority pre-empts the license by calling either the function `compensateLicensor` or `compensateLicensee`. Observe that it is important that the code guarantees that, in all the possible cases, the assets, both the token and the balance, are not indefinitely locked-in the contract.

### 7.3 A bet contract

The bet contract is a simple example of a legal contract that contains an element of randomness (*alea*), i.e. where the existence of the performances or their extent depends on an event which is entirely independent of the will of the parties. The main element of the contract is a future, aleatory event, such as the winner of a football match, the delay of a flight, the future value of a company's stock.

A digital encoding of a bet contract requires that the parties explicitly agree on the source of data that will determine the final value of the aleatory event (the *Data Provider*), that is a specific online service, an accredited institution, or any trusted third party. It is also important that the digital contract provides precise time limits for accepting payments and for providing the actual value of the aleatory event. Indeed there can be a number of issues: the aleatory event does not happen, e.g. the football match is cancelled, or the data provider fails to provide the required value, e.g. the online service is down.

```

1209 1 legal_contract Alea {
1210 2     assets bet1, bet2
1211 3     fields alea_fact, val1, val2, data_source,
1212 4         fee, amount, t_before, t_after
1213 5
1214 6     agreement(Better1, Better2, DataProvider){
1215 7         DataProvider =SET=> fee, data_source
1216 8         DataProvider =OK=> t_after, alea_fact
1217 9         Better1 =SET=> alea_fact, amount, t_before, t_after
1218 10        Better1 =OK=> fee, data_source
1219 11        Better2 =OK=> alea_fact, amount, fee, t_before, t_after,
1220 12            data_source
1221 13    } => @Init
1222 14
1223 15    @Init Better1 : place_bet(x)[y]
1224 16        (y == amount){
1225 17            y -> bet1
1226 18            x -> val1
1227 19            t_before >> @First { bet1 -> Better1 } => @Fail
1228 20        } => @First
1229 21
1230 22    @First Better2: place_bet(x)[y]
1231 23        (y == amount){
1232 24            y -> bet2
1233 25            x -> val2
1234 26            t_after >> @Run {
1235 27                bet1 -> Better1
1236 28                bet2 -> Better2 } => @Fail
1237 29        } => @Run
1238 30
1239 31    @Run DataProvider : data(x,z)[]
1240 32        (x==alea_fact){
1241 33            (z==val1){
1242 34                // The winner is Better1
1243 35                fee -> bet2, DataProvider
1244 36                bet2 -> Better1
1245 37                bet1 -> Better1

```

```

37        }
38        (z==val2){
39            // The winner is Better2
40            fee -> bet1, DataProvider
41            bet1 -> Better2
42            bet2 -> Better2
43        }
44        (z != val1 && z != val2){
45            //No winner
46            bet1 -> DataProvider
47            bet2 -> DataProvider
48        }
49    } => @End

```

Listing 4: The contract for a bet

The *Stipula* code in Listing 4 corresponds to the case where Better1, respectively Better2, places `val1`, respectively `val2`, a bet corresponding to the agreed amount of currency, stored in the contract's assets `bet1` and `bet2` respectively<sup>1</sup>. Observe that both bets must be placed within an (agreed) time limit `t_before` (line 17), to ensure that the legal bond is established before the occurrence of the aleatory event. The second timeout, scheduled in line 24, is used to ensure the contract termination even if the `DataProvider` fails to provide the expected data, through the call of the function `data`.

Compared to the Digital Licence in Listing 3, the role of the `DataProvider` here is less pivotal than that of the Authority. While it is expected that Authority will play its part, `DataProvider` is much less than a peer of the contract. It is sufficient that it is an independent party that is entitled to call the contract's function to supply the expected external data. The crucial point of trust here is the `data_source`, not the `DataProvider`. In other terms, since the parties involved in the agreement need not to trust each other, it might happen that `DataProvider` supplies an incorrect value through the function `data`. In this case, the betters can appeal against the data provider since they agreed upon the data emitted by the `data_source`. As usual, any dispute that might render the contract voidable or invalid, e.g. one better knew the result of the match in advance, can be handled by adding to the code of an authority party, according to the pattern illustrated in the Digital Licence example.

### 7.4 Specification patterns in *Stipula*

The clauses of the foregoing legal acts have been specified in *Stipula* by means of patterns that, in our view, are common from the juridical point of view. This is summarized in Table 5.

## 8 RELATED WORKS

A number of projects have put forward legal markup languages, to wrap logic and other contextual information around traditional legal prose, and providing templates for common contracts. OpenLaw [26] also allows to reference Ethereum-based smart contracts into legal agreements, and automatically trigger them once the agreement is digitally signed by all parties. Signatures by all relevant parties are stored on IPFS (the Inter-Planetary File System) and the Ethereum blockchain. The Accord project [22] provides an open, standardized format for smart legal contracts, consisting of natural language and computable components. These contracts can

<sup>1</sup>For simplicity, this code requires Better1 to place its bet before Better2, however it is easy to add similar function to let the two bets be placed in any order.

Example	Legal clauses	Smart Encoding in <i>Stipula</i>
Free rent	Permission and Prohibition	States of the contract to allow or prevent the call of a function
	Obligation to return the good	Event: timeout that triggers a repercussion
	Access to a physical good without transfer of ownership	(non fungible) Token: transfer only a usage code associated to the Token, i.e. operation uses (token, L) → L
	Prohibition of preventing the usage	Token held in custody in the smart contract
Digital License	Permission, Prohibition, and Obligation	States and Events
	Currency and escrow	(divisible) Asset
	Usage and purchase	Token: either transfer only an associated, and personal, usage code, or transfer the token, i.e. uses (token, L) → L and token → L
	Off-chain constraints and Authority for dispute resolution	<i>Implicitly</i> trusted party included in agreement
Bet	Commitment and Obligation	Event
	Aleatory value	Event: timeout to decide the effective value
	Authority that decides the value	<i>Explicitly</i> trusted party included in agreement

**Table 5: Legal clauses of the archetypal legal acts and their encodings in *Stipula***

then be interpreted by machines but they do not necessarily operate on blockchains. These projects come with sets of templates for standard legal contracts, that can be customized by setting template’s parameters with appropriate values. In *Stipula*, rather than software templates, it is possible to define specific programming patterns that can be used to encode the building blocks of legal contracts. (see the Table 5). Lexon [14] uses context free grammars to define a programming language syntax that is at the same time human readable and automatically translated into, e.g. Solidity. Even if the high level Lexon code is very close to natural language, there is no real control over the code that is actually run: the semantics of the high level language is not defined, thus the actual behaviour of the contract is that of the automatically generated Solidity code, which might be much more subtle than that of the (much simpler and more abstract) Lexon source. Compared to the Solidity code of the Lexon examples in [15], the *Stipula* version of the same contracts is much clearer, thanks to primitives like agreement and asset movements. Thus, directly coding in *Stipula* appears to be safer than relying on the Lexon-Solidity pair. Nevertheless, it should be not difficult to design an automatic translation from Lexon to *Stipula*.

Our work aims at conducting a foundational study of legal contracts, in order to elicit a precisely defined set of building blocks that can be used to describe, analyse and execute (thus enforce) legal agreements. This is similar to what has been done in [17], which puts forward a set of combinators expressing financial and insurance contracts, together with a denotational semantics and algebraic properties that says what such contracts are worth. These ideas have been ported on the (Cardano) blockchain by the Marlowe language [5], which is a small domain specific language featuring constructs like participants, tokens, currency and timeouts to wait until a certain condition becomes true (similarly to *Stipula*).

In Marlowe, money cannot be locked forever in a contract because financial contracts have a finite lifetime, at the end of which

any remaining money is returned to the participants. In *Stipula*, the finite lifetime of a contract may be easily programmed by arranging its value during the agreement phase, and issuing a corresponding event in the initial state. Marlowe semantics is written in Haskell and is executed on the Cardano blockchain almost directly because the bytecode is very similar to Haskell (at the time of writing, Cardano developers are defining a bytecode language called Plutus). In the Cardano implementation of Marlowe, the lifetime is computed in terms of slot numbers, after which a transaction is issued that closes the contract (refund all the money in its account) [21]. It turns that Cardano leverages scripts that express conditions on the time [5]. This means that *Stipula* events can be implemented in Cardano by using slot numbers, as well. However, porting an imperative language like *Stipula* to Cardano amounts to writing an interpreter/compiler that manages contract executions following a continuation-passing style. Overall, this is an issue that merits future investigations.

The class-based programming style of *Stipula* is similar to that of Solidity, but there are many differences between the two. Actually, *Stipula* is much similar to Obsidian [9], which is based on state-oriented programming and explicit management of linear assets, whose usability has been experimentally assessed [8]. Obsidian has a type system that ensures the correct manipulations of objects according to their current states and that linearly typed assets are not accidentally lost. On the contrary, *Stipula* is untyped: the introduction of a type discipline is orthogonal and is postponed to a later stage where we plan to investigate static analysis techniques specifically suitable to the legal setting. As a cons, Obsidian has no agreement nor event primitives, therefore the consensus about the contract’s terms and the enforcing of legal obligations must be implemented in a much more indirect way.

## 9 CONCLUSIONS

This paper presents a domain-specific language for defining legal contracts that can be automatized on a blockchain system. *Stipula* features a distilled number of operations that enable the formalisation of the main elements of juridical acts, such as permissions, prohibitions, and obligations. *Stipula* is formal and we are proud of its semantics – the legal bisimulation – that allows one to equate contracts that differ for clauses (events) that can never be triggered or for the order of non-interfering communications. Furthermore, the *Stipula* semantics has also been used for sketching an implementation on top of the Ethereum blockchain.

We believe that a range of legal arrangements can be adequately translated into *Stipula*, using simple patterns for the key elements of legal contracts (see the archetypal examples in the Appendix). Nevertheless we acknowledge that legal contracts cannot be fully replaced by *Stipula* contracts, since the formers thoroughly use the flexibility and generality of natural languages, may appeal to complex and undetermined social-normative concepts (such as fairness or good faith), may need to be revised as circumstances change, may need intelligent enforcements in case on non-compliance, etc. It is matter of (our) future research to deeply investigate these interdisciplinary aspects and provide at least partial solutions.

From the computer science point of view, a number of issues deserve to be investigated in full detail: the extension of the language with operations for failures, devising linear type systems that enforce the partial correctness of *Stipula* codes, the implementation of the language, the definition of a (formal) translation in *Stipula* of a high level language, such as Lexon or part of it, in order to relieve lawyers from understanding computer science jargons.

Overall, we are optimistic that future research on *Stipula* can satisfactorily address the above issues because its model is simple and rigorous, which are, in our opinion, fundamental criteria for reasoning about legal contracts and for understanding their basic principles. In our mind, *Stipula* is the backbone of a framework where addressing and studying other, more complex features that are drawn from juridical acts.

## ACKNOWLEDGMENTS

Giovanni Sartor has been supported by the H2020 European Research Council (ERC) Project “CompuLaw” (G.A. 833647). Cosimo Laneve has been partly supported by the H2020-MSCA-RISE project ID 778233 “Behavioural Application Program Interfaces (BEHAPI)”. We are grateful to the law student Alessandro Parenti for his help in constructing the legal examples in Section 7.

## REFERENCES

- [1] Solidity Documentation: State Machine Common Pattern. <https://docs.soliditylang.org/en/v0.8.0/common-patterns.html#state-machine>.
- [2] Malta MDIA Act. At <https://mdia.gov.mt/wp-content/uploads/2018/10/MDIA.pdf>, 2018.
- [3] Obsidian: A safer blockchain programming language. Language Site at <http://obsidian-lang.com/>, 2018.
- [4] Smart contract legislation and enforceability in Italy. Gazzetta Ufficiale, Law of 11 febbraio 2019, n. 12, Art. 8 ter, at <https://www.gazzettaufficiale.it/eli/id/2019/02/12/19G00017/sg>, 2019.
- [5] Cardano Documentation. <https://docs.cardano.org/>, 2020.
- [6] Wyoming Regulation Act. At <https://www.wyoleg.gov/Legislation/2021/SF0038>, 2021.
- [7] Sam Blackshear, David L. Dill, Shaz Qadeer, Clark W. Barrett, John C. Mitchell, Oded Padon, and Yoni Zohar. Resources: A safe language abstraction for money. *CoRR*, 2020.

- [8] Michael J. Coblentz, Jonathan Aldrich, Brad A. Myers, and Joshua Sunshine. Can advanced type systems be usable? an empirical study of ownership, assets, and typestate in obsidian. *Proc. ACM Program. Lang.*, 4(OOPSLA):132:1–132:28, 2020.
- [9] Michael J. Coblentz, Reed Oei, Tyler Etzel, Paulette Koronkevich, Miles Baker, Yannick Bloem, Brad A. Myers, Joshua Sunshine, and Jonathan Aldrich. Obsidian: Typestate and assets for safer blockchain programming. *ACM Trans. Program. Lang. Syst.*, 42(3):14:1–14:82, 2020.
- [10] A. Das, S. Balzer, J. Hoffmann, F. Pfenning, and I. Santurkar. Resource-aware session types for digital contracts. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, pages 111–126, Los Alamitos, CA, USA, jun 2021. IEEE Computer Society.
- [11] Sam Blackshear et al. Move: A language with programmable resources. <https://developers.diem.com/main/docs/move-paper>, 2021.
- [12] Giusella Finocchiaro and Chantal Bompreszi. A legal analysis of the use of blockchain technology for the formation of smart legal contracts. *MediaLaws*, July 2020.
- [13] Ethereum Foundation. Token Standards. <https://ethereum.org/en/developers/docs/standards/tokens/>, 2015-21.
- [14] Lexon Foundation. Lexon Home Page. <http://www.lexon.tech>, 2019.
- [15] Lexon Foundation. Lexon Demo Editor. <http://demo.lexon.tech/apps/editor/>, 2020.
- [16] Guido Governatori, Florian Idelberger, Zoran Milosevic, Regis Riveret, Giovanni Sartor, and Xiwei Xu. On legal contracts, imperative and declarative smart contracts, and blockchain systems. *Artificial Intelligence and Law*, 26:377–409, 2018.
- [17] Simon L. Peyton Jones, Jean-Marc Eber, and Julian Seward. Composing contracts: an adventure in financial engineering, functional pearl. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP '00)*, Montreal, Canada, September 18-21, 2000, pages 280–292. ACM, 2000.
- [18] Giuliano Lemme. Blockchain, smart contracts, privacy, o del nuovo manifestarsi della volontà contrattuale. In Giuffrè Francis Lefebvre, editor, *Privacy digitale*, pages 293–323. Italy, 2019.
- [19] Robin Milner. *Communication and concurrency*. PHI Series in computer science. Prentice Hall, 1989.
- [20] Margaret Jane Radin. The deformation of contract in the information society. *Oxford Journal of Legal Studies*, 37, 2017.
- [21] Pablo Lamela Seijas, Alexander Nemish, David Smith, and Simon Thompson. Marlowe: implementing and analysing financial contracts on blockchain. 2020.
- [22] Open Source Contributors. The Accord Project. <https://accordproject.org>, 2018.
- [23] Research Group on EC Private Law (Acquis Group) Study Group on a European Civil Code. *Principles, Definitions and Model Rules of European Private Law: Draft Common Frame of Reference (DCFR), Outline Edition*. Sellier, 2009.
- [24] Alex Biryukov/Dmitry Khovratovich/Sergei Tikhomirov. Findel: Secure derivative contracts for ethereum. 10323, 2017.
- [25] Gavin Wood. Ethereum: A secure decentralised generalised transaction ledger. At <https://github.com/ethereum/yellowpaper/>, 2021.
- [26] Aaron Wright, David Roon, and ConsenSys AG. OpenLaw Web Site. <https://www.openlaw.io>, 2019.

## A TECHNICAL APPENDIX

**Theorem 5.4.** The following non-interference laws hold in *Stipula* (whenever they are applicable, we assume  $x \notin fv(E')$  and  $x' \notin fv(E)$  and  $h \notin fv(E')$  and  $h' \notin fv(E')$  and  $h'' \notin fv(E)$  and  $h''' \notin fv(E)$ ):

$$\begin{aligned}
 E \rightarrow A \quad E' \rightarrow A' &\simeq E' \rightarrow A' \quad E \rightarrow A & 1517 \\
 E \rightarrow x \quad E' \rightarrow A &\simeq E' \rightarrow A \quad E \rightarrow x & 1518 \\
 E \rightarrow x \quad E' \rightarrow x' &\simeq E' \rightarrow x' \quad E \rightarrow x & 1519 \\
 E \multimap h, A \quad E' \rightarrow A' &\simeq E' \rightarrow A' \quad E \multimap h, A & 1520 \\
 E \multimap h, A \quad E' \rightarrow x' &\simeq E' \rightarrow x' \quad E \multimap h, A & 1521 \\
 E \multimap h, h' \quad E' \rightarrow A &\simeq E' \rightarrow A \quad E \multimap h, h' & 1522 \\
 E \multimap h, h' \quad E' \rightarrow x' &\simeq E' \rightarrow x' \quad E \multimap h, h' & 1523 \\
 E \multimap h, A \quad E' \multimap h'', A' &\simeq E' \multimap h'', A' \quad E \multimap h, A & 1524 \\
 E \multimap h, A \quad E' \multimap h'', h''' &\simeq E' \multimap h'', h''' \quad E \multimap h, A & 1525 \\
 E \multimap h, h' \quad E' \multimap h'', h''' &\simeq E' \multimap h'', h''' \quad E \multimap h, h' & 1526
 \end{aligned}$$

**PROOF.** We prove the first equality. Let be  $S_1 = E \rightarrow A \quad E' \rightarrow A'$  and  $S_2 = E' \rightarrow A' \quad E \rightarrow A$ , and let be  $C_1 = C[S_1]$  and  $C_2 = C[S_2]$ , then we need to prove that  $C_1(-, \emptyset, -, -) \stackrel{\text{t}}{\simeq} C_2(-, \emptyset, -, -) \stackrel{\text{t}}{\simeq}$ . Let also  $C_i = C_i(\Phi, \ell, -, \Psi[S_i])$ , with  $i = 1, 2$ , be the runtime contract where the statement  $S_i$  occurs within a number of handlers of future events. We demonstrate that the symmetric closure of the

1535 following relation is a legal bisimulation:

$$\begin{aligned}
 & \{ (C_1(-, \emptyset, -, -), \mathbb{t}, C_2(-, \emptyset, -, -), \mathbb{t}) \} \\
 & \cup \{ (C_1(Q, \ell, -, \Psi[S_1]), \mathbb{t}', C_2(Q, \ell, -, \Psi[S_2]), \mathbb{t}') \\
 & \quad | \text{ for every } Q, \ell, \Psi[ ], \mathbb{t}' \}
 \end{aligned}$$

1540 Indeed, notice that the statement  $E \rightarrow A \ E' \rightarrow A'$  can only contribute to the behavior of  $C$  with a couple of transitions during the evaluation of the body of a function or the evaluation of an event handler. Therefore the statement must be completely executed with the same time clock, possibly a number  $k$  of times due to the multiple function calls and event handlers that are executed during the same time clock.

Formally, if  $C_1, \mathbb{t}' \xrightarrow{\mu_1} \dots \xrightarrow{\mu_n} C'_1, \mathbb{t}' \rightarrow C'_1, \mathbb{t}' + 1$ , then the sequence  $\mu_1 \dots \mu_n$  contains  $k$  occurrences of the pair  $v \rightarrow A, v' \rightarrow A'$ . Similarly, there exist  $\mu'_1 \dots \mu'_n$  and a configuration  $C'_2, \mathbb{t}'$  such that  $C_2, \mathbb{t}' \xrightarrow{\mu'_1} \dots \xrightarrow{\mu'_n} C'_2, \mathbb{t}' \rightarrow C'_2, \mathbb{t}' + 1$ , where the sequence  $\mu'_1 \dots \mu'_n$  is identical to  $\mu_1 \dots \mu_n$  but for the  $k$  occurrences of the pair  $v \rightarrow A, v' \rightarrow A'$  that has been swapped into  $v' \rightarrow A', v \rightarrow A$ . The argument also holds in the converse direction.  $\square$

1541  
1542  
1543  
1544  
1545  
1546  
1547  
1548  
1549  
1550  
1551  
1552  
1553  
1554  
1555  
1556  
1557  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565  
1566  
1567  
1568  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577  
1578  
1579  
1580  
1581  
1582  
1583  
1584  
1585  
1586  
1587  
1588  
1589  
1590  
1591  
1592  
1593

1594  
1595  
1596  
1597  
1598  
1599  
1600  
1601  
1602  
1603  
1604  
1605  
1606  
1607  
1608  
1609  
1610  
1611  
1612  
1613  
1614  
1615  
1616  
1617  
1618  
1619  
1620  
1621  
1622  
1623  
1624  
1625  
1626  
1627  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652