

# eFLINT: a Domain-Specific Language for Executable Norm Specifications

L. Thomas van Binsbergen  
Centrum Wiskunde & Informatica  
Amsterdam, The Netherlands  
ltvanbinsbergen@acm.org

Robert van Doesburg  
Leibniz Institute, University of Amsterdam / TNO  
Amsterdam, The Netherlands  
robertvandoesburg@uva.nl

Lu-Chi Liu  
University of Amsterdam  
Amsterdam, The Netherlands  
l.liu@uva.nl

Tom van Engers  
Leibniz Institute, University of Amsterdam / TNO  
Amsterdam, The Netherlands  
vanengers@uva.nl

## ABSTRACT

Software systems that share potentially sensitive data are subjected to laws, regulations, policies and/or contracts. The monitoring, control and enforcement processes applied to these systems are currently to a large extent manual, which we rather automate by embedding the processes as dedicated and adaptable software services in order to improve efficiency and effectiveness. This approach requires such *regulatory services* to be closely aligned with a formal description of the relevant norms.

This paper presents eFLINT, a domain-specific language developed for formalizing norms. The theoretical foundations of the language are found in transition systems and in Hohfeld's framework of legal fundamental conceptions. The language can be used to formalize norms from a large variety of sources. The resulting specifications are executable and support several forms of reasoning such as automatic case assessment, manual exploration and simulation. Moreover, the specifications can be used to develop regulatory services for several types of monitoring, control and enforcement. The language is evaluated through a case study formalizing articles 6(1)(a) and 16 of the General Data Protection Regulation (GDPR). A prototype implementation of eFLINT is discussed and is available online.

---

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).

GPCE, *Generative Programming: Concepts & Experiences*, 2020

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8174-1/20/11...\$15.00

<https://doi.org/10.1145/3425898.3426958>

## CCS CONCEPTS

• **Software and its engineering** → **Domain specific languages**; • **Computing methodologies** → *Reasoning about belief and knowledge*; • **Security and privacy** → Privacy-preserving protocols.

## KEYWORDS

normative modeling, domain-specific language, policy enforcement, GDPR, executable specifications

## ACM Reference Format:

L. Thomas van Binsbergen, Lu-Chi Liu, Robert van Doesburg, and Tom van Engers. 2020. eFLINT: a Domain-Specific Language for Executable Norm Specifications. In *Proceedings of the 19th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences (GPCE '20), November 16–17, 2020, Virtual, USA*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3425898.3426958>

## 1 MOTIVATION

Governmental institutions provide services to citizens and companies that are primarily defined in laws and regulations. However, in practice there is often no clear connection between the software systems that support or provide these services and the laws and regulations that govern them. Similarly, business processes are subjected to laws and (inter)national regulations as well as internal policies, branch-wide codes and contracts. In both government and business, a direct connection between a software's implementation and the norms that govern the software's operations is highly desirable. A direct connection makes the software easier to validate and increases the software's maintainability with respect to following changes in regulations and policies. Moreover, with a direct connection it is possible to explain the actions taken within a software system to stakeholders in terms of the relevant norms. Our approach is to automate the required monitoring, control and enforcement processes,

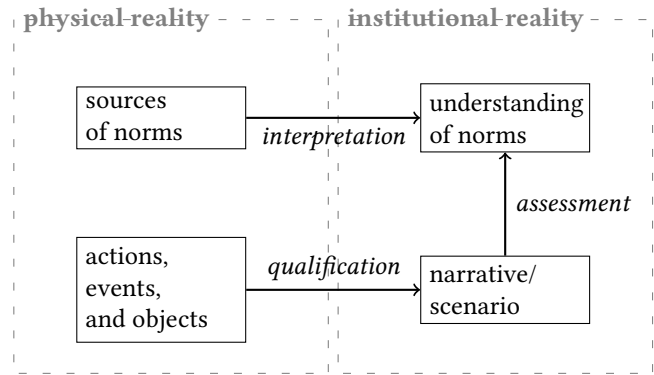
that currently are predominantly manual processes, as dedicated and adaptable regulatory services. To improve trust, the regulatory services are based on formal specifications of the relevant norms that can be verified in isolation.

This paper presents eFLINT, a domain-specific language (DSL) for formalizing norms as executable specifications. Compared to existing languages, eFLINT is novel in several respects and is most similar to languages based on the event calculus such as Symboleo [28] and InstAL [20]. A significant body of work exists concerning the formalization, analysis and enforcement of specific kinds of norms [14] such as policies for access control [29], network policies [1] (e.g. firewall configurations) and contracts [27, 28]. Instead, eFLINT is designed for describing a wide variety of normative sources such as laws, regulations, policies and contracts. Other formal languages for expressing norms are based on deontic logics [12], action logic [15] and defeasible logic [10, 18]. Some of these languages are not suited to capture some important aspects of norms such as the actors bound by the norms and the activities regulated by these norms. An important aspect of eFLINT is that the language is action-based and that the normative positions of actors are derived from the actions they can perform (permissions) or are expected to perform (duties) at a given moment in time. Moreover, the language supports the legal concept of power – the ability to grant or remove permissions or duties to/of actors. The benefit of the action-based approach is that checking the compliance of a scenario or software implementation is simplified because scenarios and software implementations are inherently action-based. Together, these features enable eFLINT for various types of applications requiring online or offline compliance-checking, monitoring, traceability and explainability.

This paper contributes by presenting eFLINT, discussing its use in a variety of applications, reflecting on its design and placing it in a wider context. The language is introduced through an example in Section 3. In section 4 we explain how eFLINT is used for offline and online compliance checking. Section 5 formalizes the parts of articles 6 and 16 of the GDPR (General Data Protection Regulation) that relate to ‘consent’ and the ‘right to rectification’ as a case study. After a reflection on the features and design of the language, the language is compared to relevant alternatives in Section 6.

## 2 LEGAL FOUNDATIONS

In this section we summarize the normative theory that underpins eFLINT. The theory is explained in reference to legal case analysis, involving the processes of *interpretation*, *qualification* and *assessment*, visualized in Figure 1. The diagram distinguishes between physical reality (left-hand side) and



**Figure 1: Schematic overview of the processes of interpretation, qualification and assessment.**

the institutional reality of Searle’s social theory [26] (right-hand side) as the reality in which actors interact physically with objects and each other on the one hand, and certain abstractions over that physical reality on the other hand. The institutional abstractions are twofold: a general understanding of the norms relevant to a case (top right) and the understanding of the case itself as a narrative (bottom right, henceforth called scenario): a series of actions, events and observations that may or may not be compliant.

The process of assessment determines whether a particular scenario is compliant with a particular understanding of the norms. In order to assess a scenario, it is first necessary to interpret the sources of norms one deems relevant to the case (top left), such as legal documents, regulations and policy descriptions. The process of qualification attributes institutional meaning to certain actions performed by actors, events and objects (bottom left). Such qualification is always context-bound which makes the qualification process a subtle interplay between the observations and the interpreted norms applicable to those observations. For example, the action of raising one’s arm is typically qualified as “requesting permission to ask a question” in a classroom but as “placing a bid” during an auction.

The normative aspect of eFLINT is based on the legal<sup>1</sup> framework constructed by Hohfeld for analyzing courthouse activities [13]. The first core aspect of Hohfeld’s framework is the observation that the ‘normative position’ of an individual, such as the individual being deemed to having a ‘duty’ or ‘power’, is always with respect to another individual. For example, if person X has the duty to do A, then there is a person Y having a ‘claim’ to A being done (and benefiting

<sup>1</sup>Although Hohfeld developed the framework for the analysis of courthouse activities, we apply its concepts more generally and speak of “normative positions” rather than “legal positions” and “normative relations” rather than “legal relations” in this paper.

from  $A$ ). In this example,  $X$  and  $Y$  are said to be in a *duty-claim* relation with respect to action  $A$ . The second core aspect of Hohfeld’s framework is the explicit consideration of change caused by an actor exercising a ‘power’ (performing a certain action), possibly having an impact on the actor with the correlative ‘liability’ position. For example, if  $X$  and  $Y$  are in a *power-liability* relation with respect to  $A$ , then  $X$  has the power to do  $A$  and  $Y$  is liable to – in the sense of ‘being bound to’ – the effects of  $A$ .

eFLINT sets itself apart from other formal languages for norm specifications by integrating both core aspects of Hohfeld’s framework, i.e. describing ‘normative relations’ rather than individual positions and allowing normative relations to change over time by the effects of actions and events.

### 3 LANGUAGE OVERVIEW

The dual nature of institutional reality is reflected in the design of eFLINT. An interpretation<sup>2</sup> is formalized as a collection of type declarations. A scenario is formalized as a series of statements. The statements describe a trace in the transition system induced by the declarations of act-types and event-types. Figure 2 gives the (simplified) abstract syntax of eFLINT specifications as sequences of type declarations. The operators that form Boolean expressions and instance expressions<sup>3</sup> have been omitted; they will be introduced alongside the case study in Section 5. An example specification is given by the listings in Figure 4. Both this example and the GDPR specification of Section 5 are available and can be tested in the web-interface for eFLINT [32]. The example captures the norm that a child has the power to ask a legal parent (i.e. a natural or adoptive parent) for help with their homework, resulting in a duty for the parent to help. Types can be re-defined by subsequent type declarations. This is convenient from a perspective of reuse: a generic interpretation can be used by several applications by letting each application specialize certain types to the domain of the application.

Figure 3 gives the abstract syntax of scripts as sequences of statements and queries. An example script is given by the listings of Figure 5. In the example scenario, Alice is a natural parent of Bob, Bob asks Alice for help, but Alice only helps when the homework is already due, causing the `help-with-homework` duty to be violated. For automatic case assessment, it is convenient to use four input files: a file containing the type-declarations of a generic model, a file re-declaring types according to a concrete or specialized domain, a file containing initialization statements and a file that contains the statements of a scenario. This separation across files is practical as there is a conceptual one-to-many relation between these files (in the order listed above). For

<sup>2</sup>We use “interpretation” also for the result of the process of interpretation.

<sup>3</sup>An instance expression computes instances of declared types.

$x \in$	<b>type_ids</b>	::=	...
$s \in$	<b>strings</b>	::=	...
$z \in$	$\mathbb{Z}$	::=	$\{0, 1, \dots, -1, -2, \dots\}$
$i \in$	<b>instance_exprs</b>	::=	...
$b \in$	<b>boolean_exprs</b>	::=	...
$\delta \in$	<b>domains</b>	::=	<b>strings</b>   <i>string_set</i> ( $s_1, \dots, s_n$ )   $\mathbb{Z}$   <i>int_set</i> ( $z_1, \dots, z_n$ )   <i>product</i> ( $x_1, \dots, x_n$ )   ...
$fdc \in$	<b>fact_decls</b>	::=	<i>fdecl</i> ( $x, \delta, b_0^?$ )
$adc \in$	<b>act_decls</b>	::=	<i>adec</i> ( $x_0, x_1, x_2, x^*, c^*, b_0^?$ )
$edc \in$	<b>event_decls</b>	::=	<i>edec</i> ( $x_0, x^*, c^*, b_0^?$ )
$ddc \in$	<b>duty_decls</b>	::=	<i>ddec</i> ( $x_0, x_1, x_2, x^*, b^*, b_0^?$ )
$c \in$	<b>post_conditions</b>	::=	<i>create</i> ( $i$ )   <i>terminate</i> ( $i$ )
$dc \in$	<b>decls</b>	::=	$fdc$   $adc$   $edc$   $ddc$   ...
	<b>specifications</b>	::=	$dc^*$

Figure 2: Abstract syntax of eFLINT specifications.

	<b>elems</b>	::=	$s$   $z$   <i>tuple</i> ( $v_1, \dots, v_n$ )   ...
$v \in$	<b>instances</b>	=	<b>elems</b> $\times$ <b>type_ids</b>
$\sigma \in$	<b>configs</b>	=	$\mathcal{P}$ ( <b>instances</b> )
$t \in$	<b>stmts</b>	::=	<i>create</i> ( $i$ )   <i>terminate</i> ( $i$ )   <i>trigger</i> ( $i$ )
$q \in$	<b>queries</b>	::=	<i>query</i> ( $b$ )
	<b>scripts</b>	::=	( $t$   $q$ ) <sup>*</sup>

Figure 3: Abstract syntax of eFLINT scenarios.

example, many scenarios can start from the same initial state.

The declaration of a type determines the set of values (instances) of that type, inherited either from an atomic type (e.g. strings or integers) or a composite record-type (represented as tuples in the abstract syntax). Record-types define relations over concepts (when they have two or more fields, e.g. `natural-parent`) or establish predicates over a concept (when they have one field, e.g. `homework-due`). An institutional model of the physical world *at a particular moment in time* is represented by the values, referred to as a *facts*, deemed to hold true at that moment<sup>4</sup>. Sets of facts are the configurations in the transition system induced by the specification. If a record of type `homework-due` is in configuration  $\sigma$ , then the person at the `child` field of the record is deemed to having their homework due according to  $\sigma$ . The accuracy of this fact depends on the accuracy of the qualifications that caused it to hold, e.g. the `+homework-due(Bob)` statement in Figure 5.

<sup>4</sup>The facts of eFLINT are the fluents of the event calculus, see Section 6.

```

Fact person Identified by String
Placeholder parent For person
Placeholder child For person
Fact natural-parent Identified by parent * child
Fact adoptive-parent Identified by parent * child
Fact legal-parent Identified by parent * child
Holds when adoptive-parent(parent, child)
           || natural-parent(parent, child)
Act ask-for-help
  Actor child
  Recipient parent
  Creates help-with-homework(parent, child)
  Holds when legal-parent(parent, child)
Fact homework-due Identified by child
Duty help-with-homework
  Holder parent
  Claimant child
  Violated when homework-due(child)
Act help
  Actor parent
  Recipient child
  Terminates help-with-homework(parent, child)
  Holds when help-with-homework(parent, child)

```

```

Fact person Identified by Alice, Bob, Chloe, David

```

**Figure 4: Type declarations capturing (normative) concepts (top) and a domain of discourse (bottom).**

```

+natural-parent(Alice, Bob).
+adoptive-parent(Chloe, David).

ask-for-help(Bob, Alice).
+homework-due(Bob). // homework deadline passed
?Violated(help-with-homework(Alice, Bob)).
help(Alice, Bob).

```

**Figure 5: A script consisting of an initial state (top) and a scenario with a query (bottom). Duty help-with-homework is violated after homework is due.**

Type declarations may have an optional derivation clause (**Holds when**). A derivation clause is a Boolean expression<sup>5</sup> ( $b_0^?$  in the abstract syntax) that computes which instances of the type hold true in a given configuration. The derivation clause of `legal-parent` determines that, for every parent  $P$  and child  $C$ , `legal-parent(P,C)` holds true if and only if `adoptive-parent(P,C)` or `natural-parent(P,C)` holds true. A fact-type is either a ‘derived fact’ or ‘postulated fact’, depending on whether it is declared with a derivation clause. To ensure consistency, only postulated facts can be created or terminated by statements.

<sup>5</sup>This is a simplification. Derivation clauses are fully explained in Section 5.

*act-, event- and duty-type declarations.* Act-, event- and duty-types are fact-types with additional meaning. An act-type declaration consists of a *performing actor-type* (**Actor**), a *recipient actor-type* (**Recipient**) and optional further related types (**Related to**). An action<sup>6</sup> – an instance of an act-type – is a record value with a field for each of the associated types (a *tuple*  $(v_1, \dots, v_n)$ ). An action  $A$  is said to be enabled in  $\sigma$  if  $A$  holds true according to  $\sigma$ , i.e. if  $A$  is a member of  $\sigma$ . If  $A$  is enabled in  $\sigma$ , then the performing actor  $X$  and the recipient actor  $Y$  are in a power-liability relation with respect to  $A$  in  $\sigma$ . The post-conditions associated with an act-type ( $c^*$  in the abstract syntax) determine the effects its instances have when executed by a statement. The effects are to create or terminate facts, thus giving rise to a new configuration (and possibly updated normative relations).

The institutional view on the world can also change by physical actions for which there is no institutional counterpart. For example, there is no direct institutional counterpart to ‘changing address’ in the GDPR even though relocations influence the accuracy of personal data. Moreover, the world also changes due to natural events such as earthquakes, fires and the passage of time. For these reasons, eFLINT distinguishes between actions and events. An event-type declaration (not in the example) is essentially an act-type declaration without a performing actor or recipient actor.

A duty-type declaration contains the type of the *duty holder*, the type of the *claimant* and optional further related types. A duty – an instance of a duty-type – is a record value with a field for each of the types. A duty-type declaration has zero or more violation-conditions ( $b^*$  in the abstract syntax). If a duty holds true in configuration  $\sigma$ , then the holder and claimant of the duty are in a duty-claim relation in  $\sigma$ . The holder of a duty is in violation of the duty (and the claimant has a valid claim) according to  $\sigma$  if the duty holds true in  $\sigma$  and if one of the violation-condition holds true in  $\sigma$ . To avoid violating a duty, an actor must perform an action that terminates the duty (e.g. `help`) before one of the violation conditions holds.

*transitions and compliance.* In summary, a set of facts forms a configuration representing an institutional view on the physical world at a particular moment in time. For any given configuration it is possible to determine the normative relations between actors. Actions and events change configurations when executed, potentially modifying the normative relations between actors. Since actions and duties are also facts, an actor may have the power to assign duties or to grant powers to others. The post-conditions of action- and event-types give rise to a transition system. The execution of an action or event triggers a transition by removing and/or

<sup>6</sup>Instances of act-types are institutional actions. There is not necessarily an institutional actions for every physical action or vice versa.



inserting facts from/to the current configuration. For example, the statement `ask-for-help(Bob,Alice)`. (*trigger(i)* in the abstract syntax) executes the action `ask-for-help(Bob,Alice)`, causing the creation of a duty for Alice. Individual facts can also be created and terminated (*create(i)* and *terminate(i)* in the abstract syntax). For example, `+homework-due(Bob)` creates the fact that Bob’s homework is due (termination is written with a minus symbol). The statements of a script form a trace (sequence of transitions) in the transition system. A query (e.g. `?Violated(help-with-homework(Alice,Bob))`) is a Boolean expression that is evaluated in the context of the current configuration. The language also supports invariants: queries that have to hold true in every reachable configuration. For example, the following fragment shows an invariant declaration that captures that one cannot be their own legal parent.

```
Invariant parentship:
  Not(Exists person : legal-parent(person, person))
```

Invariant violations are reported as soon as they arise in order to discover inconsistencies in (applications of) specifications.

A trace may be *action-compliant* and/or *duty-compliant*. A trace is action-compliant if every transition on the trace is labeled with an event or action that is enabled in the source configuration of the transition. A trace is duty-compliant if no duties are violated in any of its configurations. These notions are independent: a trace can be action-compliant, duty-compliant, action- and duty-compliant or neither. Assessing a scenario is deciding whether it produces an action- and duty-compliant trace, given an initial configuration. The example scenario is action-compliant but not duty-compliant as a violation occurs after the second statement.

A trace records the normative positions and relations of all actors as they evolve over time and therefore provides sufficient information to determine important details about violations such as when they occurred and which actors were responsible. This a crucial aspect: by recording traces, eFLINT makes it possible to reproduce the entire decision making process and to explain the decisions that have been made.

## 4 IMPLEMENTATION

The previous section explained eFLINT informally through an example. This section gives an overview of the different applications supported by the eFLINT prototype implementation at the time of writing. The implementation is available online [31]. A simple web-interface for automatic case assessment is available online as well [32]. The details of the expression language used by the eFLINT implementation have been omitted in the previous section and are discussed alongside the GDPR case study in the next section. The transition system semantics of eFLINT depends only on the expression language in that there are Boolean expressions and instance

expressions. Alternative expression languages can thus be used by alternative implementations, e.g. using objects rather than records to structure data. Similarly, our implementation has integers and strings as atomic values, but other types of atoms, such as floating points, are easily added.

*automated assessment.* One of the executables of the implementation receives a specification and script and determines whether the scenario in the script is action- and duty-compliant and whether all the queries are successful. The output is a sequence of violations and failed queries or a JSON object that also includes the produced trace. The JSON output has been used to develop the aforementioned web-interface [32]. Besides editing, the interface can be used to analyze traces by inspecting the contents of, and the changes to, configurations. The web-interface has been used in a MSc-level course on ‘Policy Making and Rule Governance’ at the University of Amsterdam, with user-feedback feeding directly into the design of the language. Automatic assessment is an important tool during the development of eFLINT specifications as it facilitates testing and debugging. Once a specification has been adopted, the primary purpose of automatic assessment is to analyze concrete cases that have been observed or hypothetical cases that might arise. Both are crucial, not only in the development of the specification, but also as feedback to lawmakers and policymakers. As part of future work we intend to add model checking to our implementation, expanding the set of tools through which confidence in the correctness of a specification is obtained. As mentioned, eFLINT already has safety properties in the form of invariants.

*exploration.* To further support the aforementioned use cases, the eFLINT implementation also enables manually exploring the transition system induced by a specification. The tool can run as a Read-Eval-Print Loop (REPL) loaded with a specification and a script producing an initial state. At the top-level, the REPL accepts declarations, queries and statements, which can be mixed freely and produce immediate feedback. It is also possible to delete or re-declare types, enabling on-the-fly updates to the specifications. After every statement, the REPL reports violations of action- or duty-compliance, changes to the current configuration and any invariants that were not upheld. The user can backtrack to a previously visited configuration to explore an alternative scenario. An example interaction with the REPL is shown in Figure 6. The REPL is loaded with the specification of Figure 4. The interaction shows Chloe helping David in the first explored branch. After backtracking, another branch is explored in which homework is due before Chloe has helped.

During the first interaction in Figure 6, the REPL responds with the information that the fact `legal-parent(Chloe,David)` has been added to the configuration. This fact is derived

```

Available commands:
:<INT>          trigger action or event <INT>
:force <INT>    force action or event <INT>
:revert <INT>   revert to configuration <INT>
:display :d     show the current configuration
:options :o     show available actions & events
:help :h       show these commands
:quit :q       end the exploration
or just type a <PHRASE>
#0 > +natural-parent(Chloe,David)
+legal-parent(Chloe,David)
+natural-parent(Chloe,David)
+ask-for-help(David,Chloe)
enabled actions & events:
1. ask-for-help(David,Chloe)
#1 > :1
+help(Chloe,David)
+help-with-homework(Chloe,David)
enabled actions & events:
1. ask-for-help(David,Chloe)
2. help(Chloe,David)
#2 > :2
-help(Chloe,David)
-help-with-homework(Chloe,David)
enabled actions & events:
1. ask-for-help(David,Chloe)
#3 > ?Violated(help-with-homework(Chloe,David))
query failed
#3 > :revert 2
enabled actions & events:
1. ask-for-help(David,Chloe)
2. help(Chloe,David)
#2 > +homework-due(David)
violated duty!: help-with-homework(Chloe,David)

```

**Figure 6: Example interaction with the eFLINT REPL.**

from the fact `natural-parent(Chloe,David)` postulated by the user. In order to make this derivation, the implementation has enumerated all possible instances of `legal-parent` and evaluated the derivation clause for each. Enumerating all instances of types is possible when working with a small<sup>7</sup>, finite domain. However, when checking the compliance of a running system, an application discussed below, an open-ended domain is typically required. The ability to redefine and specialize types in eFLINT enables us to reuse specifications across applications in which some require a finite and others an open-ended domain. In the example of Figure 4, an open-ended domain (the declaration of `person` initially does not list its instances) is replaced by a finite domain. Reusing specifications in both types of applications is also made possible by the pragmatic design choice to give different behavior to the enumeration operator (`ForEach`, introduced in the next section) depending on whether it enumerates instances of a finite or infinite type. In the latter case, the operator only

<sup>7</sup>In order not to suffer from combinatorial explosion.

enumerates the instances of the type that hold true in the current configuration. Note that a domain is finite if all of the atomic types have finite sets of instances, because then, by induction, all record types are finite too.

*normative actors.* Benefiting from the principled approach to REPLs presented in [33], the back-end of the REPL has been reused to form the basis of a TCP server. The server is loaded with a specification and waits for incoming declarations, statements and queries on a given port. The server is used to integrate eFLINT specifications in arbitrary software systems. To develop and experiment with regulatory services for enforcement, we use the Akka framework<sup>8</sup> for actor-oriented programming in Scala. Actor-oriented programming can be used to develop or model complex, distributed systems. The components of a software system are implemented as actors, with message-passing as the only form of communication between them.

Central to our approach is the notion of a ‘normative actor’ that administers an eFLINT specification. A normative actor is created with one or more files containing declarations and statements and starts its own server instance. As the server is loaded, the input files are executed in order, so that later files may specialize types introduced by earlier files.

A survey of various software architectures that incorporate policy enforcement mechanisms [22, 23, 37, 38] has revealed at least the four types of enforcement listed below. Our implementation of normative actors in Scala enables these four types of enforcement.

- *Ex-ante enforcement of permissions:* ensuring that an actor has permission to execute a particular action before it is performed and blocking the action if there is no permission
- *Ex-ante enforcement of positive duties:* informing actors of their duties to perform certain actions
- *Ex-post enforcement of violations of prohibitions:* applying some form of resolution to the observation that an action has been performed which was not enabled
- *Ex-post enforcement of violated duties:* applying some form of resolution to a violated duty

A normative actor responds to eFLINT statements and queries received as messages. The response to a query is simply whether the query holds true. Actors can send queries to normative actors to let the responses guide their behavior. For example, an actor can ensure action-compliance by checking whether an action is enabled before performing it.

Receiving a statement causes the normative actor to update its internal state (a configuration) by executing the statement. The resulting transition might impact other actors in the system and they will be informed by the normative

<sup>8</sup><https://akka.io>

actor accordingly. If a duty is created or violated by the transition, the holder and claimant of the duty are informed. Similarly, if an action is enabled by the transition, the performing and recipient actor of the action are informed. If the transition was triggered by a disabled action, the performing and recipient actor of the action are informed of this violation. The actors receiving such messages can react in several meaningful ways. For example, the claimant of a violated duty might have the power to notify an authority that, in turn, has the power to place a penalty on the holder of the violated duty. Another common use case is for the claimant of a new (but not yet violated) duty to start a timer that runs out when the claimant thinks the duty should have been fulfilled (terminated). The example of Figure 4 can be extended with a teacher that places the `complete-homework` duty on children. When the timer expires, the claimant (teacher) sends a message to the normative actor to communicate this observation (in the form of a fact, e.g. `homework-due`) possibly causing duties to be violated (e.g. the duty `complete-homework`, but perhaps also `help-with-homework`). In our experiments with GDPR, the event `rectification-delay`, creating the fact `undue-rectification-delay` (both discussed in Section 5), is an event triggered by the use of a timer.

The actor sending a statement to a normative actor may be a neutral observer and, for example, not the performer or recipient of an action. The normative actor responds to the observer with a summary of the effects of the transition, similar to the output produced by the command-line REPL. Normative actors can thus be used in a variety of ways.

A system can have one or more monitoring actors making qualifications based on the observed communication between other actors. These qualifications are sent to normative actors that administer the norms considered by the monitoring actor. In this use case, the communications of normative actors can be restricted to monitoring actors only.

In a multi-agent system (MAS), normative actors can be internalized by agents, playing the role of a (moral, social, or legal) conscience. An agent communicates with its internal normative actors to possibly update its beliefs, desires and intentions. Every normative actor of an agent embodies the particular interpretation of a set of norms adopted by the agent. In this case, the communications of normative actors should be restricted to their encompassing agent.

## 5 CASE STUDY: GDPR

The eFLINT specification developed in this section captures the following aspects of the General Data Protection Regulation (GDPR) [19]: the requirement to receive a data subject’s consent prior to data processing and the subject’s ‘right to rectification’. The purpose of this section is to dive deeper

into the details of the language, such as its expression language, and to demonstrate its expressivity in connection to a realistic case. The presented code snippets form the GDPR component of a larger case study regarding the Know Your Customer (KYC) requirements placed on financial institutions. This case study is performed in collaboration with the banks ABN AMRO and ING. In order to improve the accuracy of customer risk assessment, the banks are willing to share customer data under certain conditions. The banks wish to keep certain data secret, e.g. if the data provides a competitive advantage. Moreover, the banks want to demonstrate compliance with the GDPR. The goal of the wider case study is to experiment with architectures for regulated data sharing systems. Academically, the case study is interesting in that it requires reasoning about multiple norm specifications, each with their own ontologies of concepts. Moreover, satisfying a duty in one policy might cause a violation in another, demonstrating the need for priorities between norms.

The KYC case consists of three eFLINT specifications of less than a 100 lines of code for which one or more normative actors (see previous section) are created. Besides the GDPR specification, the case involves an internal policy specification and a sharing agreement. Every bank has their own specialization of the internal policy. The eFLINT specifications are kept small, formalizing only specific rules and norms to focus on their interaction at the level of policy design and the level of component behavior and implementation. The system consists of actors representing banks, employees and clients. These actors communicate with normative actors loaded with instances of the three mentioned eFLINT specifications. This communication is indirect and happens via intermediate actors that convert domain knowledge (about banks, employees and clients) to institutional knowledge (about institutional facts, powers and duties) and vice versa. These intermediate actors are derived from a high-level specification (not shown here). An interesting aspect of our approach is that normative actors are created for pairs of interacting banks and clients, with each normative actor loaded with a version of the GDPR specification specialized to that bank and client. This specialized variant of the GDPR specification can be seen as a contract between an individual bank and an individual client, which is monitored by the normative actor. The normative actor informs the bank and the client of any duties and violations.

### 5.1 Concept definitions

The following code fragment below captures the GDPR concepts ‘subject’ (a natural person) and ‘data’.

```
Fact subject
Fact data
Fact subject-of Identified by subject * data
```

A fact-type declaration without an **Identified by** clause defaults to **Identified by String**.

*expressions.* Consider the following expression.

```
(Exists subject: subject-of(subject, data))
```

Expressions are literals, variables or operators and constructors applied to other expressions. Type names can occur as variables in expressions, such as `subject` and `data` above. Type names can also occur as constructors in expressions, for example `subject-of` above. There is no ambiguity between a type name occurring as a variable or as a constructor because only constructors are followed by (zero or more) formal arguments within parentheses.

Constructor application can be written in two styles. The first style – familiar from functional and logic programming – requires as many arguments as the number of fields of the constructed record and the arguments must be written in the same order as the fields are written in the type-declaration. An example is `subject-of(subject, data)`. In the second style, field names are explicitly mentioned. For example, in `subject-of(subject=subject, data=data)` the name `subject` occurs as a field name on the left-hand side of the equal symbol and as a variable on the right-hand side. In this style, formal arguments can be written in any order and can also be omitted. If a formal argument for field  $x$  is omitted, then it defaults to  $x = x$ . The constructor application of this example can thus be written as `subject-of()`. If the variable `subject-of` is bound to a record, then `subject-of.subject` evaluates to the value of the field `subject` of that record (projection).

*accumulators.* The example expression shows that eFLINT is based on first-order logic with existential and universal quantification via the **Exists** and **Forall** operators. The inner expression of a quantifier, appearing behind the colon, must be a Boolean expression. However, the sub-expression `subject-of(subject, data)` of the example is an instance expression when taken out of context. The static semantics of eFLINT rewrites this expression to the Boolean expression `Holds(subject-of(subject, data))`. The **Holds** operator checks whether the instance computed by its operand holds true in the current configuration.

The example expression is equivalent to the following:

```
Or(ForEach subject: subject-of(subject, data))
```

The semantics of **ForEach** are to evaluate its inner expression multiple times, each time binding its binders (the comma-separated variables before the colon) to a different combination of instances of their respective types. For example, when the above expression is evaluated, the inner expression is evaluated once for every possible binding of the variable `subject` to an instance of the type `subject`. The behavior of **ForEach** differs depending on whether all its binders refer to

types with finite numbers of instances. If this is the case, then the binders are bound to all possible combinations of instances of their types. If one or more of the types does not have a finite amount of instances, only the instances of these types that hold true in the current configuration are enumerated instead. As mentioned in the previous section, this design decision has been made to simultaneously accommodate applications with finite and open-ended domains.

The **ForEach** operator is non-deterministic in that it computes multiple values. However, non-deterministic expressions are only allowed in certain places, such as in the post-conditions of actions and as the operand of an accumulator. An accumulator is an operator that reduces a sequence of values to a single result (thus turning a non-deterministic expression in a deterministic one). The **Or** accumulator evaluates to **True** if any of its inputs is **True**. Similarly, **And** evaluates to **True** if all its inputs are **True**. Occurrences of **Exists** and **Forall** desugar to an application of **ForEach** inside an application of **Or** or **And** respectively. Other examples of accumulators are **Count** and **Sum** for counting instances and summing integers.

*derivation clauses.* The example expression was taken from the following fact-type declaration with a derivation clause.

```
Fact personal-data Identified by data Holds when
(Exists subject: subject-of(subject, data))
```

The derivation clause determines that `data` is personal data if it has a `subject`, which closely resembles the definition of “personal data” in Article 4(1) of the GDPR.

Derivation clauses come in two forms: a **Holds when** clause with a Boolean expression or a **Derived from** clause with an instance expression. The instance expression of a **Derived from** clause produces all the instances of the type that are deemed to hold true by the clause. Any variables not explicitly bound by occurrences of **Exists**, **Forall** or **ForEach** in this instance expression are implicitly bound by an outermost **ForEach**. A **Holds when** clause is syntactic sugar for a **Derived from** clause. The above fragment is equivalent to the following:

```
Fact personal-data Identified by data Derived from
(ForEach data: personal-data() When
(Exists subject: subject-of(subject, data)))
```

The **When** operator is used in non-deterministic expressions to filter out unwanted results. The operator evaluates to the result of its first operand, but only if its second operand evaluates to **True**.

In general, a **Holds when** clause with Boolean expression `[t]` for a fact-type `[x]` with fields `[a]`, `[b]` and `[c]` desugars to **Derived from** `(ForEach [a],[b],[c]: [x]() When [t])`. A **Holds when** clause can therefore only be part of a fact-type declaration with a record type. The desugaring shows that the expression of a **Holds when** clause is evaluated in the context



of a record instance to determine whether that instance holds true. In the example, this is the instance `personal-data(data = data)` for every instance of `data`.

## 5.2 Consent

This subsection formalizes Article 6(1)(b) on consent [19]. The following fragment defines the related concepts of data controller, data processor, purpose and consent.

```
Fact controller
Fact processor
Fact processes Identified by
  processor * data * controller * purpose
Fact purpose
Fact consent Identified by
  subject * controller * purpose
Fact accurate-for-purpose Identified by
  data * purpose
```

A controller is a legal entity collecting and processing the data of a data subject. A processor is a legal entity storing or processing data on behalf of a controller. An important aspect of the GDPR is that a controller communicates the purpose for which it is collecting and processing data and that the subject gives explicit consent to processing<sup>9</sup> the data for that purpose. If the record of type `consent` that contains subject *S*, controller *C* and purpose *R* holds true in a configuration, then this means that *S* has given consent to *C* to collect and process their data for purpose *R*. If the record of type `processes` that contains processor *P*, controller *C*, data *D* and purpose *R* holds true in a configuration, then this means that *P* processes the data *D* on behalf of *C* for the purpose *R*.

The presented formulation abstracts over operations on data by capturing all changes to data as replacements. When the postal address of a bank's client changes, the currently held postal address is no longer accurate for the purpose of building a KYC client profile. However, there is no definition of an action for clients (subjects) to change their addresses because relocation is not a GDPR notion. Instead, an instance of the `data-change` event, defined below, is triggered when a change in address is observed.

```
Event data-change
  Related to data, new-data, purpose
  When data != new-data
    && subject-of() && subject-of(data = new-data)
  Terminates accurate-for-purpose(data, purpose)
  Creates accurate-for-purpose(new-data, purpose)
  Holds when accurate-for-purpose(data, purpose)
```

```
Placeholder new-data For data
```

The `When` clause introduces an instance constraint (omitted from the abstract syntax in Section 3), establishing a connection between the fields of the declared type. In this example,

<sup>9</sup>Consent is one of several grounds for lawfully processing personal data.

the constraint determines that the new data is indeed new and has the same subject as the old data. An instance constraint keeps the `Foreach` operator from enumerating 'invalid' instances of the declared type.

The definition of `data-change` refers to multiple instances of `data`. To support multiple such references, variables can have 'decoration' in the form of integer numbers or prime characters at the end of their name (e.g. `data1`, `data2`, and `data'`). The user can also add their own variable names with a `Placeholder` declaration. In the fragment above, the name `new-data` is introduced as a placeholder for instances of the fact-type `data`. Note that the constructor application `subject-of()` has `data` as an implicit argument because this is the name of one of its fields.

The act-type declaration of `give-consent`, shown below, describes the power of subjects giving consent to controllers. The `Related to` clause reflects that consent is given for a specific purpose.

```
Act give-consent
  Actor subject
  Recipient controller
  Related to purpose
  Creates consent()
  Holds when !consent()
```

The derivation clauses and post-conditions of an act-type are evaluated in an environment that binds the field names of the type. For example, to determine the effects of the instance `give-consent(Alice,Bank,KYC)`, the variable `subject` is bound to `Alice`, the variable `controller` is bound to `Bank` and the variable `purpose` is bound to `KYC`. These bindings are used as the (implicit) arguments to the constructor application `consent()`. The derivation clause of `give-consent` prevents repeated execution of the same instance of `give-consent` (although one can argue that the power to give consent is unconditioned).

The act-type `collect-personal-data` given below determines that consent must have been given by the subject for the purpose for which the data is being collected and that the data must be accurate for this purpose. If, in physical reality, multiple processors process the collected data, then the institutional action `collect-personal-data` is to be executed multiple times, with different processor arguments.

```
Act collect-personal-data
  Actor controller
  Recipient subject
  Related to data, processor, purpose
  When subject-of()
  Creates processes()
  Holds when consent() && accurate-for-purpose()
```

## 5.3 The right to rectification

Article 16 of the GDPR states [19]:

*The data subject shall have the right to obtain from the controller without undue delay the rectification of inaccurate personal data concerning him or her. [...]*

The interpretation we formalize consists of:

- (1) the power for the data subject to demand rectification when the data processed by the controller is inaccurate for the purpose it is processed
- (2) a duty held by the controller to rectify the data without undue delay when the subject demands it
- (3) the power of the controller to terminate this duty once their processors use accurate data

The fragment below formalizes (1) as an act-type.

```
Act demand-rectification
  Actor subject
  Recipient controller
  Related to purpose
  Creates rectification-duty()
  Holds when (Exists data, processor:
    subject-of() && processes() &&
    !accurate-for-purpose())
```

The action `demand-rectification` is enabled for a subject if there is a processor that, on behalf of the controller, processes inaccurate personal data of the subject. The effect of the action is to place the duty (point (2) above) on the controller.

The duty is defined by the following duty-type declaration.

```
Duty rectification-duty
  Holder controller
  Claimant subject
  Related to purpose
  Violated when undue-rectification-delay()
```

```
Fact undue-rectification-delay Identified by
  controller * purpose * subject
```

Since accuracy depends on purpose, the duty-claim relation between controllers and subjects is related to a purpose.

Like the conditions of act-types, a violation condition is evaluated in an environment binding the fields of the types. The `rectification-duty` is violated whenever there is undue delay in between the demand for rectification and the rectification taking place. The usage of the term “undue delay” in the article leaves room for discussion (deliberately). In other words, potential cases of delay are subject to qualification. The event `rectification-delay`, defined below, can be triggered to indicate that the duty has been violated because of an undue delay. The fact `undue-rectification-delay()` is only created if the rectification duty (still) exists.

```
Event rectification-delay
  Related to controller, purpose, subject
  Creates undue-rectification-delay()
  When rectification-duty()
```

Events without an explicit derivation clause hold true by default, i.e. they have an implicit clause `Holds when True`. Actions without a derivation clause are considered ‘postulated facts’, i.e. they can be created or terminate by other actions.

The application of `When` in the creating post-condition of `rectification-delay` shows that post-conditions can be non-deterministic expressions evaluating to zero or more instances. All instances computed for a creating (terminating) post-condition are created (terminated). The post-condition of `rectification-delay` computes zero or one instances of `undue-rectification-delay()` depending on whether the expression `rectification-duty()` holds true. A post-condition can evaluate to more than one instance by the implicit or explicit use of `ForEach`. Any variables that are not bound in a post-condition are bound by `ForEach` implicitly. Similarly, any variables not bound in derivation clauses or violation conditions are bound by `Exists` implicitly. If a creating post-condition evaluates to an instance that already holds true, then this instance is not added to the configuration a second time (a configuration is a set). Conversely, if a terminating post-condition evaluates to an instance that does not hold true, then terminating this instance has no effect.

The controller can relieve itself of the `rectification-duty` when its processors use accurate data for the given purpose.

```
Act rectified-personal-data
  Actor controller
  Recipient subject
  Related to purpose
  Terminates rectification-duty()
  ,undue-rectification-delay()
  Holds when processors-accurate()
```

The predicate `processors-accurate` determines whether all processors that process data on behalf of a controller for a specific subject and purpose use accurate data:

```
Fact processors-accurate
  Identified by controller * subject * purpose
  Holds when (Forall processor, data :
    accurate-for-purpose()
    When processes() && subject-of())
```

## 5.4 Reflections

The GDPR case study presented in this section demonstrates that eFLINT can be used to formalize, rather concisely, norms described in significant, real-world regulations. Although not shown in this paper, the formalization can be used to assess concrete cases and can be used to reason about the compliance of running systems. The web-interface for eFLINT provides an example scenario and demonstrates case assessment with the GDPR specification [32]. In future work we

intend to give a comprehensive account of a generic data-sharing architecture that involves normative actors for regulatory services, showing in particular how multiple eFLINT specifications co-exist and how eFLINT is used to enforce compliance of software with respect to multiple regulations, policies and contracts.

In the design of eFLINT, focus has been on the possibility to simultaneously use specifications in isolation – typically with a finite domain – as well as in running systems – typically with an open-ended domain. To this end, the language has a reactive, REPL-oriented design, based on [33], supporting manual exploration and enabling external systems to trigger actions and events. This motivation guided other design choices as well, such as the semantics of `ForEach` and the ability to redefine types to form specialized domains. The ability to redefine types makes it possible to reuse a norm specification across different applications in which certain concepts, such as `data` of the GDPR, are concretized differently.

Omitting formal arguments in constructor applications has significantly improved the brevity of the GDPR specifications. Moreover, by hiding the structure of data, clauses of type declarations read almost like natural text and less like programming instructions. For example, the derivation clause `Holds when consent()` reads more natural than `Holds when consent(subject, controller, purpose)`. The downside is the (mental) effort required to reproduce the full constructor applications whenever a detailed reading of the code is necessary. This effort can be greatly reduced with IDE support, e.g. by showing the declaration of a type when holding the cursor over a constructor.

The distinction between derived facts and postulated facts is a crucial feature of eFLINT. Derivation clauses effectively describe logical implications of the kind that are common in logical programming. However, the consequents are not limited to Boolean formulas. For example, in the following code, the type `wealth` always has exactly one instance corresponding to the sum of all people’s savings.

```
Fact person
Fact balance Identified by Int
Fact balance-of Identified by person * balance
Fact wealth Identified by Int Derived from
Sum(ForEach balance-of : balance-of.balance)
```

An important aspect of postulated facts is that their validity is established externally. In the current implementation, facts are proactively provided by an external system. However, the implementation can be extended to support on-demand requests. Determining the validity of a fact can then be delegated to a relevant authority, e.g. a health-care provider confirming the validity of a receipt to support an insurance claim. Further reflections are in comparison to related work.

## 6 RELATED WORK

The aspects of norms automated by software are roughly divided into structural aspects – such as production, instantiation and publication – and dynamic aspects – such as implementation, modification, termination, monitoring and enforcement. On the structural side, standards have been developed for the digital representation of contracts (e.g. Oasis eContracts [6]) and for referring to sources of law (e.g. MetaLex, Akomo Ntoso, Juriconnect and ECLI). Type-declarations in eFLINT can easily be extended with references to sources, e.g. using MetaLex [2], making it possible to relate components of traces in the transition system to sources, further enhancing explainability. A predecessor of eFLINT demonstrated this capability and has been used to discover inconsistencies in sources of norms [34]. Natural language processing can assist with the interpretation process, introducing the necessary structure to allow contracts to be analyzed by software [4, 35]. However, legal experts and policy makers are not all programmers, and certain information required to compute with norms is not present in sources, e.g. the structure of, and relations between, values. This observation has motivated and influenced the development of eFLINT.

The computational theory underneath eFLINT is most similar to event calculus [16] and its simplified variants [24]. In the variants of [5, 21], time is represented as a sequence of (distinct) time points. At each time point, certain *fluents* are stated to hold true by the *HoldsAt*(*f*, *t*) predicate (with *f* a fluent and *t* a time point). Events initiate or terminate fluents at certain time points according to the predicate *Initiates*(*e*, *f*, *t*) or *Terminates*(*e*, *f*, *t*) (with *e* an event). The judgment *Happens*(*e*, *t*) states that event *e* takes place at time *t*. A collection of axioms determines that initiated fluents hold true until they are terminated, that terminated fluents do not hold until they are initiated (again) and that the correct fluents are initiated and terminated after events happen. The facts of eFLINT correspond to the fluents of the event calculus and time points can be seen as identifiers for configurations, i.e. *HoldsAt*(*f*, *t*) states that fact *f* is in configuration *t*. The creating and terminating post-conditions of actions and events in eFLINT correspond to judgments involving the *Initiates* and *Terminates* predicates respectively (for every possible time point). The derivation clauses of fact-type declarations introduce rules of the form  $HoldsAt(f_1, t) \wedge \dots \wedge HoldsAt(f_k, t) \implies HoldsAt(f_{k+1}, t)$ . Expressions in eFLINT are thus interpreted as logical formulae with `ForAll` and `Exists` implying the usage of first-order logic. Halpern and Weissman discuss the use of first-order logic to describe and reason about policies, identifying in particular a number of first-order languages with different characteristics regarding the complexity of reasoning [11].

A scenario in eFLINT corresponds to a set of concrete judgments of the form  $Happens(e, t)$ . However,  $Happens(e, t)$  can be stated for multiple instances of  $e$  and the same  $t$ , i.e. multiple events can happen simultaneously, which is not true for the actions and events of eFLINT. Perhaps eFLINT can be extended with declarations of composite events for this purpose. The notions of action-violation and duty-violation used by eFLINT can be formulated as logical predicates. A translation from eFLINT to answer set programming based on the connection with the event calculus is being considered.

Several formal languages for norms are based on extensions to the event calculus such as Symboleo [28] and Instal [20]. Besides the event calculus, Symboleo and eFLINT are also related in their Hohfeldian foundations. In Symboleo obligations and powers are instantiated by ‘triggers’. In eFLINT, derivation clauses can be used for this purpose. However, in eFLINT expressions are always evaluated in the current configuration, whereas Symboleo is true to the event calculus in that expressions concern the entire timeline.

Significant work exists about the formalization, analysis and enforcement of specific kinds of policies such as policies for access control and network policies [1] (e.g. firewall configurations), of which a survey is given by Jabal et al [14]. The eFLINT language is instead used to describe a wide variety of normative sources such as laws, regulations, policies and contracts. The Margrave Policy Analyzer tool<sup>10</sup> can be used to reason about access control and firewall configurations, supporting several formalisms such as the widely adopted XACML for access control [29]. The tool implements several types of reasoning, including Change-Impact Analysis [8].

Governatori et al. give a detailed overview on the interpretation and lifecycle of contracts in [9]. In terms of the elements described in [9], eFLINT is used to make interpretations explicit, including *implied terms* and norms that follow from the *integration* of the contract in a wider context. With the example of undue rectification delay and accuracy for purpose, we have shown how *open-textured terms* are treated by explicit qualification in eFLINT. The automated assessment of scenarios, explained in Section 4, can assist with *dispute resolution*. Section 4 also explains how normative actors can be used for *monitoring* and *enforcement*. The implementation also enables dynamic *modification* of contracts by removing and redefining (act- and duty-) types on-the-fly. Several avenues are being considered to support the *implementation* of contracts formalized in eFLINT, including smart contracts.

The idea of smart contracts was first introduced by Szabo [30] as software (or hardware) that facilitates the exchanges of digital items of value between two or more parties, with a security mechanism in place that ensures the

exchange at a risk low enough for all parties. With the advent of blockchain technology [17], smart contracts are now typically understood as scripts that facilitate the execution of ‘transactions’. The underlying blockchain technology forms a distributed ledger establishing consensus between potential witnesses about the history of transactions. A popular smart contract language is Solidity [7], running on the Ethereum platform [3, 36]. The Flint language (unrelated to eFLINT) offers a safer alternative to Solidity for writing smart contracts running on Ethereum [25]. The Marlowe DSL is used to develop smart contracts at a higher level of abstraction that run on the Cardano platform [27].

## 7 CONCLUSIONS

We have presented eFLINT, a novel domain-specific modeling language for formalizing norms found in laws, regulations, policies, contracts and (data-sharing) agreements. The action-oriented nature of the language makes it possible to use the language in a variety of ways, including case analysis and monitoring the compliance of a running system. Pragmatic design decisions have been made to allow specifications to be reused for both types of reasoning and across applications. In particular, the generic concepts encountered in laws and regulations can be formalized at a high level of abstraction, whilst applications of these formalizations can effortlessly redefine the concepts to match the domain of the application. Our approach involves the explicit qualification of physical reality as institutional facts, actions, events and duties. The normative positions of actors evolve over time as actions are performed and events take place. The resulting traces can be used to diagnose violations and to provide explanations about the decisions made based on the norms.

*Acknowledgements.* The authors thank colleagues and reviewers for their comments and suggestions on earlier versions of this paper. This work is supported by the NWO project (628.009.014) Secure Scalable Policy-enforced Distributed Data Processing (SSPDDP), part of the NWO research program Big Data: Real Time ICT for Logistics.

## REFERENCES

- [1] E. S. Al-Shaer and H. H. Hamed. 2004. Modeling and Management of Firewall Policies. *IEEE Transactions on Network and Service Management* 1, 1 (2004), 2–10. <https://doi.org/10.1109/TNSM.2004.4623689>
- [2] A. Boer, R. Hoekstra, E. De Maat, F. Vitali, M. Palmirani, and B. Ratai. 2010. Metalex (Open XML Interchange Format for Legal and Legislative Resources). Technical Report CWA, Vol. 15710:2010. European Committee for Standardization (CEN).
- [3] V. Buterin. 2018. Ethereum White Paper.
- [4] I. Chalkidis, I. Androutsopoulos, and A. Michos. 2017. Extracting Contract Elements. In *Proceedings of the 16th Edition of the International Conference on Artificial Intelligence and Law (ICAIL 2017)*. ACM, 19–28. <https://doi.org/10.1145/3086512.3086515>

<sup>10</sup><http://www.margrave-tool.org/>



- [5] M. Charalambides, P. Flegkas, G. Pavlou, A. K. Bandara, E.C. Lupu, A. Russo, N. Dulay, M. Sloman, and J. Rubio-Loyola. 2005. Policy Conflict Analysis for Quality of Service Management. In *6th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY 2005)*, 6-8 June 2005, Stockholm, Sweden. IEEE, 99–108. <https://doi.org/10.1109/POLICY.2005.23>
- [6] OASIS LegalXML eContracts TC. 2007. eContracts Version 1.0 Committee Specification.
- [7] Ethereum. 2016. Solidity Documentation Online. <https://solidity.readthedocs.io>. [Online, accessed 7 October 2020].
- [8] K. Fisler, S. Krishnamurthi, L.A. Meyerovich, and M.C. Tschantz. 2005. Verification and Change-Impact Analysis of Access-Control Policies. In *Proceedings of the 27th International Conference on Software Engineering (St. Louis, MO, USA) (ICSE 2005)*. Association for Computing Machinery, New York, NY, USA, 196–205. <https://doi.org/10.1145/1062455.1062502>
- [9] G. Governatori, F. Idelberger, Z. Milosevic, R. Riveret, G. Sartor, and X. Xu. 2018. On legal contracts, imperative and declarative smart contracts, and blockchain systems. *Artificial Intelligence and Law* 26, 4 (2018), 377–409. <https://doi.org/10.1007/s10506-018-9223-3>
- [10] G. Governatori, M.J. Maher, G. Antoniou, and D. Billington. 2004. Argumentation Semantics for Defeasible Logic. *Journal of Logic and Computation* 14, 5 (10 2004), 675–702. <https://doi.org/10.1093/logcom/14.5.675>
- [11] J.Y. Halpern and V. Weissman. 2008. Using First-Order Logic to Reason about Policies. *ACM Transactions on Information and System Security* 11, 4 (2008), 21:1–21:41. <https://doi.org/10.1145/1380564.1380569>
- [12] H. Herrestad. 1993. Norms and Formalization. In *Proceedings of the 3th International Conference on Artificial Intelligence and Law (ICAIL 1993)*. ACM, 175–184. <https://doi.org/10.1145/112646.112667>
- [13] W.N. Hohfeld. 1913. Some Fundamental Legal Conceptions as Applied in Judicial Reasoning. *Yale Law Journal* 23(1) (1913), 59–64.
- [14] A.A. Jabal, M. Davari, E. Bertino, C. Makaya, S. Calo, D. Verma, A. Russo, and C. Williams. 2019. Methods and Tools for Policy Analysis. *Comput. Surveys* 51, 6, Article 121 (2019). <https://doi.org/10.1145/3295749>
- [15] A.J.I. Jones and M. Sergot. 1996. A Formal Characterisation of Institutionalised Power. *Logic Journal of the IGPL* 4, 3 (06 1996), 427–443. <https://doi.org/10.1093/jigpal/4.3.427>
- [16] R.A. Kowalski and M.J. Sergot. 1986. A Logic-based Calculus of Events. *New Generation Computing* 4, 1 (1986), 67–95. <https://doi.org/10.1007/BF03037383>
- [17] S. Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).
- [18] D. Nute. 2003. Defeasible Logic. In *Web Knowledge Management and Decision Support*, O. Bartenstein, U. Geske, M. Hannebauer, and O. Yoshie (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 151–169.
- [19] Council of the EU. 2016. General Data Protection Regulation.
- [20] J. Padget, E. Elakehal, T. Li, and M. De Vos. 2016. *InstAL: An Institutional Action Language*. Law, Governance and Technology Series, Vol. 30. Springer Verlag, 101.
- [21] A. Russo, R. Miller, B. Nuseibeh, and J. Kramer. 2002. An Abductive Approach for Analysing Event-Based Requirements Specifications. In *Logic Programming, 18th International Conference, ICLP 2002, Copenhagen, Denmark, July 29 - August 1, 2002, Proceedings (LNCS, Vol. 2401)*, P.J. Stuckey (Ed.). Springer, 22–37. [https://doi.org/10.1007/3-540-45619-8\\_3](https://doi.org/10.1007/3-540-45619-8_3)
- [22] M. Ruta, F. Scioscia, S. Ieva, G. Capurso, A. Pinto, and E. Di Sciascio. 2018. A Blockchain Infrastructure for the Semantic Web of Things. In *Proceedings of the 26th Italian Symposium on Advanced Database Systems (CEUR Workshop Proceedings, Vol. 2161)*. CEUR-WS.org.
- [23] M. Ruta, F. Scioscia, S. Ieva, G. Capurso, and E. Di Sciascio. 2017. Semantic Blockchain to Improve Scalability in the Internet of Things. *Open Journal of Internet of Things* 3 (2017), 46–61.
- [24] F. Sadri and R.A. Kowalski. 1995. Variants of the Event Calculus. In *Logic Programming, Proceedings of the Twelfth International Conference on Logic Programming, Tokyo, Japan, June 13-16, 1995*, L. Sterling (Ed.). MIT Press, 67–81.
- [25] F. Schrans, D. Hails, A. Harkness, S. Drossopoulou, and S. Eisenbach. 2019. Flint for Safer Smart Contracts. <https://arxiv.org/pdf/1904.06534.pdf>.
- [26] J.R. Searle. 1996. *The construction of social reality*. Penguin Books.
- [27] P.L. Seijas, A. Nemish, D. Smith, and S. Thompson. 2020. Marlowe: implementing and analysing financial contracts on blockchain. In *Workshop on Trusted Smart Contracts (Financial Cryptography 2020)*.
- [28] S. Sharifi, A. Parvizimosaed, D. Amyot, L. Logrippo, and J. Mylopoulos. 2020. Symbolo: Towards a Specification Language for Legal Contracts. In *28th IEEE Int. Requirements Engineering Conf. (RE 2020)*. IEEE.
- [29] OASIS Standard. 2013. eXtensible Access Control Markup Language (XACML) Version 3.0. <http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html>.
- [30] N. Szabo. 1997. Formalizing and Securing Relationships on Public Networks. *First Monday* 2, 9 (1997). <https://doi.org/10.5210/fm.v2i9.548>
- [31] L.T. van Binsbergen. 2020. Haskell prototype implementation of the eFLINT language. <https://gitlab.com/eflint/haskell-implementation>. [Online, accessed 7 October 2020].
- [32] L.T. van Binsbergen and G. Sileno. 2020. Web-interface for automatic case assessment in eFLINT. <http://ltvanbinsbergen.nl/publications/eflint/online/gpce2020/>. [Online, accessed 8 October 2020].
- [33] L. T. van Binsbergen, M. Verano Merino, P. Jeanjean, T. van der Storm, B. Combemale, and O. Barais. 2020. A principled approach to REPL interpreters. In *Proceedings of the 2020 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (GPCE 2020)*. ACM. <https://doi.org/10.1145/3426428.3426917>
- [34] R. van Doesburg, T. van der Storm, and T. van Engers. 2016. CALCULEMUS: Towards a Formal Language for the Interpretation of Normative Systems. In *AI4J Workshop at ECAI 2016 (AI4J 2016)*. 73–77.
- [35] T. van Engers and R. van Doesburg. 2015. At Your Service, On the Definition of Services from Sources of Law. In *Proceedings of the 15th International Conference on Artificial Intelligence and Law (ICAIL 2015)*. ACM, 221–0225. <https://doi.org/10.1145/2746090.2746115>
- [36] D.D. Wood. 2014. Ethereum: a secure decentralised generalised transaction ledger.
- [37] H. Zhou, X. Ouyang, J. Su, C. de Laat, and Z. Zhao. 2019. Enforcing trustworthy cloud SLA with witnesses: A game theory-based model using smart contracts. *Concurrency and Computation: Practice and Experience* (2019), e5511. <https://doi.org/10.1002/cpe.5511>
- [38] M. Zichichi, M. Contu, S. Ferretti, and V. Rodríguez-Doncel. 2020. Ensuring Personal Data Anonymity in Data Marketplaces through Sensing-as-a-Service and Distributed Ledger. In *Proceedings of the 3rd Distributed Ledger Technology Workshop Co-located with ITASEC (CEUR Workshop Proceedings, Vol. 2580)*, F. Chiaraluce and L. Mostarda (Eds.). CEUR-WS.org.