# Checking Traces for Regulatory Conformance* 

Nikhil Dinesh, Aravind Joshi, Insup Lee, and Oleg Sokolsky<br>Department of Computer Science<br>University of Pennsylvania<br>Philadelphia, PA 19104-6389, USA<br>\{nikhild, joshi,lee, sokolsky\}@seas.upenn.edu


#### Abstract

We consider the problem of checking whether the operations of an organization conform to a body of regulation. The immediate motivation comes from the analysis of the U.S. Food and Drug Administration regulations that apply to bloodbanks - organizations that collect, process, store, and use donations of blood and blood components. Statements in such regulations convey constraints on operations or sequences of operations that are performed by an organization. It is natural to express these constraints in a temporal logic. There are two important features of regulatory texts that need to be accommodated by a representation in logic. First, the constraints conveyed by regulation can be obligatory (required) or permitted (optional). Second, statements in regulation refer to others for conditions or exceptions. An organization conforms to a body of regulation if and only if it satisfies all the obligations. However, permissions provide exceptions to obligations, indirectly affecting conformance. In this paper, we extend linear temporal logic to distinguish between obligations and permissions, and to allow statements to refer to others. While the resulting logic allows for a direct representation of regulation, evaluating references between statements has high complexity. We discuss an empirically motivated assumption that lets us replace references with tests of lower complexity, leading to efficient trace-checking algorithms in practice.


## 1 Introduction

Regulations, laws and policies that affect many aspects of our lives are represented predominantly as documents in natural language. Mechanically checking compliance with these regulations and policies is an area of growing importance [1-3].

In this paper, we will consider one such regulation, the Food and Drug Administration's Code of Federal Regulations (FDA CFR) [4] that governs the operations of U.S. bloodbanks. The CFR is developed by experts in the field of medicine, and regulates the tests that need to be performed on donations of blood before they are used.

Bloodbanks are organizations that perform collection, testing, storage, and distribution of blood donations and are required to conform to the regulation (CFR). The operations of a bloodbank are logged in a database that keeps track of donations that are collected by the bloodbank, tests that are performed on them and, ultimately, the

[^0]way each donation is used. Our goal is to check in an efficient manner that the operations as recorded in the database are compliant with the CFR, and to raise an alarm if a non-compliant action is detected. To achieve this goal, we first need to settle on an approach to formalize regulatory documents, and then consider the feasibility of checking database logs with respect to the formalized regulations.

As we illustrate with examples in Section 2, the basic structure of regulatory statements is to declare that a certain action can take place when certain conditions apply. At a first glance, it seems that such statements can be encoded as logical clauses, where a set of preconditions imply a postcondition. However, there are two complications that need to be addressed. First, regulations convey permissions and obligations, which have to be reflected in the formal description and handled accordingly during the checking. Second, a common phenomenon in regulatory texts is for sentences to function as conditions or exceptions to others. This function of sentences makes them dependent on others for their interpretation, and makes the translation to logic difficult. We call this the problem of references to other laws, and it is the central focus of this paper.

In Section 2, we argue that a logic to represent regulation should provide mechanisms for statements to refer to others, and to make inferences from the sentences referred to. We then turn to formalization of regulatory documents and regulated operations. In Section 3.1, we define an abstract model for representing the operations of an organization, followed in Section 3.2 by a predicate-based linear temporal logic to express normative statements in regulation. Formal definitions of conformance are given. We then extend the logic to allow sentences to refer to others, in Section 3.3.

Section 4 describes the checking process. We adapt the methodology of the rulebased formalism Eagle [5] to handle references. In order to check statements with references, we need to compute a fixed point, propagating information between references from one statement to another until we get a consistent evaluation. The evaluation of references has high complexity. We identify a condition, motivated by a case study of the CFR, under which references can be replaced by tests of lower complexity. We also discuss a prototype checking tool.

Section 5 concludes with a discussion of future research directions and a survey of related work.

## 2 Motivation

In this section, we consider a representative sample of the CFR and argue that a logic to represent regulation should provide a mechanism for sentences to refer to others.
Example. Below we present shortened versions of sentences from the CFR Section 610.40, which we will use as a running example throughout the paper.
(1) Except as specified in (2), every donation of blood or blood component must be tested for evidence of infection due to Hepatitis B.
(2) You are not required to test donations of source plasma for evidence of infection due to Hepatitis B.

Statement (1) conveys an obligation to test donations of blood or blood component for Hepatitis B, and (2) conveys a permission not to test a donation of source plasma
(a blood component) for Hepatitis B. To assess an organization's conformance to (1) and (2), it suffices to check whether "all non-source plasma donations are tested for Hepatitis B". In other words, (1) and (2) imply the following obligation:
(3) Every non-source plasma donation must be tested for evidence of infection due to Hepatitis B.

There are a variety of logics in which one can capture the interpretation of (3), as needed for conformance. For example, in first-order logic, one can write $\forall x:(d(x) \wedge$ $\neg s p(x)) \Rightarrow \operatorname{test}(x)$, where $d(x)$ is true iff $x$ is a donation, $s p(x)$ is true iff $x$ is a source plasma donation, and $\operatorname{test}(x)$ is true iff $x$ is tested for Hepatitis B. Thus, to represent (1) and (2) formally, we inferred that they implied (3) and (3) could be represented more directly in a logic.

Now suppose we have a sentence that refers to (1):
(4) To test for Hepatitis B, you must use a screening test kit.

The reference is more indirect here, but the interpretation is: "if (1) requires a test, then the test must be performed using a screening test kit". A bloodbank is not prevented from using a different kind of test for source plasma donations. (4) can be represented by first producing (3), and then inferring that (3) and (4) imply the following:
(5) Every non-source plasma donation must be tested for evidence of infection due to Hepatitis B using a screening test kit.

It is easy to represent the interpretation of (5) directly in a logic. However, (5) has a complex relationship to the sentences from which it was derived, i.e., (1), (2) and (4). The derivation takes the form of a tree:


We argue that constructing a single derived obligation from multiple statements should be avoided. On the one hand, the derived obligation can become very complex. The full version of statement (1) in the CFR contains six exceptions, and these exceptions in turn have statements that qualify them further. A statement can be used as an exception to multiple other statements, and it is easy to see that the derived obligation can be exponentially larger than the original set of statements. We advocate an approach that allows us to introduce references into the syntax of the logic, and resolve references during evaluation.

## 3 Formalization of Regulatory Documents

In this section, we extend linear temporal logic (LTL) to distinguish between obligations and permissions, and allow references between statements. We begin, in Section 3.1, by representing a bloodbank as a run or trace. Section 3.2 extends LTL to distinguish between obligations and permissions, leading to definitions of conformance. We then extend the logic to allow sentences to refer to others (Section 3.3).

### 3.1 Model for Regulated Operations

Given the need to demonstrate conformance to the regulation in case of an audit, regulated organizations such as bloodbanks keep track of their operations in a database, for example donor information and the tests they perform. Such a system can be thought of abstractly as a relational structure evolving over time. At each point in time (state), there are a set of objects (such as donations and donors) and relations between the objects (such as an association between a donor and her donations). The state changes by the creation, removal or modification of objects. We represent this as a run.

Definition 1 (A Run of a System). Given a set $O$ (of objects) and countable sets $\Phi_{1}, \ldots, \Phi_{n}$ (where $\Phi_{j}$ is a set of predicate names of arity $j$ ), a run of a system $R\left(O, \Phi_{1}, \ldots, \Phi_{n}\right)$, abbreviated as $R$, is a tuple $\left(r, \pi_{1}, \ldots, \pi_{n}\right)$ where:

- $r: N \rightarrow S$ is a sequence of states. $N$ is the set of natural numbers, and $S$ is a set of states.
- $\pi_{j}: \Phi_{j} \times S \rightarrow 2^{O^{j}}$ is a truth assignment to predicates of arity $j$. Given $p \in \Phi_{j}$, we will say that $p\left(o_{1}, \ldots, o_{j}\right)$ is true at state s iff $\left(o_{1}, \ldots, o_{j}\right) \in \pi_{j}(p, s)$.

Given a run $R$ and a time $i \in N$, the pair $(R, i)$ is called a point (statements in linear temporal logic are evaluated at points). Given the predicate names $\left(\Phi_{1}, \ldots, \Phi_{n}\right)$, the corresponding space of runs is denoted by $\mathcal{R}\left(\Phi_{1}, \ldots, \Phi_{n}\right)$, abbreviated as $\mathcal{R}$.

| Time | Objects | Predicates |
| :--- | :--- | :--- |
| 1 | $o_{1}$ | $d\left(o_{1}\right), \operatorname{sp}\left(o_{1}\right), \neg$ test $\left(o_{1}\right)$ |
| 2 | $o_{1}$ | $d\left(o_{1}\right), \operatorname{sp(o_{1}),\neg \operatorname {test}(o_{1})}$ |
|  | $o_{2}$ | $d\left(o_{2}\right), \neg \operatorname{sp}\left(o_{2}\right), \neg$ test $\left(o_{2}\right)$ |
| 3 | $o_{1}$ | $d\left(o_{1}\right), \operatorname{sp}\left(o_{1}\right)$, test $\left(o_{1}\right)$ |
|  | $o_{2}$ | $d\left(o_{2}\right), \neg \operatorname{sp}\left(o_{2}\right), \neg$ test $\left(o_{2}\right)$ |

Table 1. A run of a bloodbank

Table 1 shows a possible run of a bloodbank. First, an object $o_{1}$ is entered into the system. $o_{1}$ is a donation of source plasma ( $d\left(o_{1}\right)$ and $\operatorname{sp}\left(o_{1}\right)$ are true). When a donation is added, its test predicate is initially false. Then, an object $o_{2}$ is added, which is a donation but not of source plasma. In the third step, the object $o_{1}$ is tested.

### 3.2 Logic for Regulatory Conformace

Predicate-based Linear Temporal Logic (PredLTL) The logic that we define in this section is a restricted fragment of first-order modal logic. The restriction is that we allow formulas with free variables, but no quantification over objects. Formulas will be interpreted using the universal generalization rule, i.e., over all assignments to free variables. The restrictions are similar in spirit to logic programs, which have been observed to be sufficiently expressive for the generic statements in regulation [6, 7].

Definition 2 (Syntax). Given sets $\Phi_{1}, \ldots, \Phi_{n}$ (of predicate names) and a set of variables $X$, the language $L\left(\Phi_{1}, \ldots, \Phi_{n}, X\right)$, abbreviated as $L$, is the smallest set such that:

- $p\left(y_{1}, \ldots, y_{j}\right) \in L$ where $p \in \Phi_{j}$ and $\left(y_{1}, \ldots, y_{j}\right) \in X^{j}$.
- If $\varphi \in L$, then $\neg \varphi \in L$ and $\square \varphi \in L$. If $\varphi, \psi \in L$, then $\varphi \wedge \psi \in L$.

Disjunction $\varphi \vee \psi=\neg(\neg \varphi \wedge \neg \psi)$ and implication $\varphi \Rightarrow \psi=\neg \varphi \vee \psi$ are derived connectives. The temporal operator is understood in the usual way: $\square \varphi$ ( $\varphi$ holds and will always hold (globally)). $\diamond \varphi$ ( $\varphi$ will eventually hold) is defined as $\neg \square \neg \varphi$.

We now extend the syntax to express normative statements in a body of regulation, by distinguishing between obligations and permissions.

Definition 3 (Syntax of Regulation). Given a finite set of identifiers ID, a body of regulation Reg is a set of statements such that for each id $\in I D$, there exist $\varphi, \psi \in L$ such that either: id.o: $\varphi \leadsto \psi \in \operatorname{Reg}$, or id.p: $\varphi \leadsto \psi \in \operatorname{Reg}$
id. $\mathbf{o}: \varphi \leadsto \psi$ (id.p: $\varphi \leadsto \psi$ ) is read as: "it is obligated (permitted) that the precondition $\varphi$ leads to the postcondition $\psi$ ".

Definition 4 (Semantics). Given a run $R=\left(r, \pi_{1}, \ldots, \pi_{n}\right), \varphi \in L$, and a variable assignment $v: X \rightarrow O$, the relation $(R, i, v) \models \varphi$ is defined inductively as follows:

- $(R, i, v) \models p\left(y_{1}, \ldots, y_{j}\right)$ iff $\left(v\left(o_{1}\right), \ldots, v\left(o_{j}\right)\right) \in \pi_{j}(p, r(i))$.
- The semantics of conjunction and negation is defined in the usual way.
- $(R, i, v) \models \square \varphi$ iff for all $k \geq i:(R, k, v) \models \varphi$.

We extend the semantic relation to regulatory staments. We take $\models$ to stand for "conforms to":

- $(R, i, v) \models$ id.o: $\varphi \leadsto \psi$ iff $(R, i, v) \models \varphi \Rightarrow \psi(\Rightarrow$ is implication $)$
- $(R, i, v) \models$ id. $\mathbf{p}: \varphi \sim \psi$. Runs vacuously conform to permissions. Permissions will become relevant when references from obligations are present (Section 3.3).

Consider again our example from Section 2. We use three predicates defined as follows. $d(x)$ is true iff $x$ is a donation. $s p(x)$ is true iff $x$ consists of source plama. test $(x)$ is true iff $x$ is tested for Hepatitis B.

Statement (3) is represented as: 3.o: $d(x) \wedge \neg s p(x) \leadsto \diamond$ test $(x)$. Statement (2) can be represented as: 2.p: $d(y) \wedge s p(y) \leadsto \neg \diamond t e s t(y)$. However, statement (1) cannot be represented directly.

The deontic concepts of obligation and permission are treated as properties of sentences. Only obligations matter for conformance. If a non-source plasma donation is not tested, there is a problem. On the other hand, a bloodbank may choose to test a donation of source plasma or not. In assessing conformance, the function of a permission is to serve as an exception to an obligation, and in this indirect manner it becomes relevant. We will give a semantics to this function of permissions in Section 3.3. Such a treatment of permissions has its basis in the legal theory of Ross [8].

In the formulation here, obligations and permissions are top-level operators and cannot be negated. This restriction can be removed in several ways, e.g., using a manyvalued interpretation. However, we avoid this to simplify presentation. A more crucial restriction is that iterated deontic constructs cannot be expressed directly, i.e., sentences of the form "required to allow x " or "allowed to require x.". One has to decide what top-level obligations or permissions are implied by these constructs. To our knowledge, handling iterated constructs is an open problem in deontic logic [9].

Conformance of a run $R$ is defined using the notion of validity. $\varphi$ is valid at the point $(R, i),(R, i) \models \varphi$, iff for all variable assignments $v$ : $(R, i, v) \models \varphi . \varphi$ is valid in $R, R \models \varphi$ iff for all $i:(R, i) \models \varphi$.

Definition 5 (Run Conformance). Given a body of regulation Reg and a run $R$ representing the operations of an organization, we say that $R$ conforms to the regulation iff for all obligations id.0: $\varphi \leadsto \psi \in \operatorname{Reg}$ we have $R \models i d . \mathbf{o}: ~ \varphi \leadsto \psi$.

The definition of conformance is given in terms of obligations. We now extend the logic to allow sentences to refer to others making permissions relevant to conformance.

### 3.3 References to Other Laws

In this section, we describe the logical machinery we use to express and handle references to laws. We give an example-driven account here, followed by a formal account in the context of a runtime checking algorithm in Section 4.1.

We extend the syntax with an inference predicate $\operatorname{by}_{\mathrm{Id}}(\varphi)$, where Id is a set of identifiers. by ${ }_{\mathrm{Id}}(\varphi)$ is read as "by the laws in Id $\varphi$ holds". There are two restrictions: (a) $\varphi$ is a statement in PredLTL (Definition 2) and (b) the predicate by ${ }_{\text {Id }}(\varphi)$ can appear only in preconditions of laws. These restrictions are similar to those that apply to justifications in default logic [11]. In the examples that we discuss, the set Id has a single element, i.e., a statement refers to a single other law. In general, laws refer to sets of statements, e.g., "except as specified in this section".

Consider again our example statements (1) and (2), which can now be represented as follows:

- 1.0: $d(x) \wedge \neg \operatorname{by}_{\{2\}}(\varphi(x)) \leadsto \diamond t e s t(x)$, and
- 2.p: $d(y) \wedge s p(y) \sim \neg \diamond \operatorname{test}(y)$

In the formula above, the subformula by ${ }_{\{2\}}(\varphi(x))$ is understood as "by the law (2) the formula $\varphi(x)$ holds". It remains to define the formula $\varphi(x)$. Intuitively, this should be the negation of the postcondition of (1). In other words, if $\neg \diamond$ test $(x)$ follows from (2), then the postcondition of (1) need not hold. ${ }^{1}$
1.0: $d(x) \wedge \neg \operatorname{by}_{\{2\}}(\neg \diamond$ test $(x)) \leadsto \diamond \operatorname{test}^{(x)}$

We interpret by $\left\{_{2\}}(\neg \diamond\right.$ test $(x))$, by letting formulas have output. In other words, when the precondition of an obligation or permission is true at a point, the point is annotated with the postcondition. Given a point $(R, i)$ and a variable assignment $v$, first we consider the formula 2.p: $d(y) \wedge s p(y) \leadsto \neg \diamond t e s t(y)$. We evaluate this as follows:

[^1]- If $(R, i, v) \models d(y) \wedge s p(y),(R, i)$ is annotated with $2: \neg \diamond t e s t(v(y))$. Observe that the annotation happens regardless of whether $(R, i, v) \models \neg \diamond$ test $(y)$ and the variable is replaced with the object assigned to it.
- Otherwise, there is no annotation.

Given a variable assignment $v$ and a PredLTL formula $\varphi, v(\varphi)$ is the formula obtained by replacing all variables $x$ by an identifier for the object $v(x)$. Note that $v(\varphi)$ is equivalent to a propositional LTL formula, as the variables have been replaced by constant symbols. We now define annotations:

Definition 6 (Annotation). Given a run $R$, a set of identifiers $I D$, a variable assignment $v$, and a body of regulation Reg, an annotation is a statement id: $v(\psi)$ such that $i d \in I D$ and id. $\mathbf{x}: ~ \varphi \leadsto \psi \in \operatorname{Reg}$ (which is either an obligation or a permission). The set of annotations is denoted by $A(R, I D$, Reg $)$, abbreviated $A$.

Definition 7 (Annotation Function). Given a run $R$, an annotation function $\alpha: N \rightarrow$ $2^{A}$ assigns a set of annotations to each point. Given a set of identifiers ID and $I d \subseteq$ $I D$, we use $\alpha . I d(i)$ to denote the set of annotations id: $\psi \in \alpha(i)$ such that $i d \in I d$.

| Time | Objects | Predicates | Annotations |
| :--- | :--- | :--- | :--- |
| 1 | $o_{1}$ | $d\left(o_{1}\right), \operatorname{sp}\left(o_{1}\right), \neg \operatorname{test}\left(o_{1}\right)$ | $2: \neg \diamond$ test $\left(o_{1}\right)$ |
| 2 | $o_{1}$ | $d\left(o_{1}\right), \operatorname{sp}\left(o_{1}\right), \neg \operatorname{test}\left(o_{1}\right)$ | $2: \neg \diamond \operatorname{test}\left(o_{1}\right)$ |
|  | $o_{2}$ | $d\left(o_{2}\right), \neg \operatorname{sp}\left(o_{2}\right), \neg \operatorname{test}\left(o_{2}\right)$ | $1: \diamond$ test $\left(o_{2}\right)$ |
| 3 | $o_{1}$ | $d\left(o_{1}\right), \operatorname{sp}\left(o_{1}\right)$, test $\left(o_{1}\right)$ | $2: \neg \diamond$ test $\left(o_{1}\right)$ |
|  | $o_{2}$ | $d\left(o_{2}\right), \neg \operatorname{sp}\left(o_{2}\right), \neg \operatorname{test}\left(o_{2}\right)$ | $1: \diamond$ test $\left(o_{2}\right)$ |

Table 2. A run and its annotations

Table 2 shows a run of a bloodbank augmented with annotations. As we discussed in Section 3.1, $o_{1}$ is a donation of source plasma which is tested at time 3 and $o_{2}$ is a non-source plasma donation which has not been tested. Unless the run is extended to test $o_{2}$ as well, it does not conform with the regulation according to Definition 5.

Since the precondition of statement (2) is true for the assignment of $y$ to $o_{1}$, we have the annotation 2: $\neg \diamond$ test $\left(o_{1}\right)$ at all points. However, since $o_{2}$ is not a donation of source plasma, there is no correponding annotation.

Now consider the formula by $\{2\}(\neg \diamond t e s t(x))$. This is evaluated as follows:

- Evaluate 2.p: $d(y) \wedge s p(y) \leadsto \neg \diamond t e s t(y)$ at $(R, i)$ w.r.t. all variable assignments.
- Let $\psi_{2}$ be the conjunction of the annotations produced by the formula for (2), i.e., $\psi_{2}=\bigwedge \varphi$ for all $\varphi \in \alpha .\{2\}(i)$.
- $(R, i, v) \models \operatorname{by}_{\{2\}}(\neg \diamond \operatorname{test}(x))$ iff $\models \psi_{2} \Rightarrow \neg \diamond \operatorname{test}(v(x))$.

Notice that the last step requires a validity check, but it is a validity check in (propositional) LTL. Validity in LTL is coNP-complete when the only modality is globally, and

PSPACE-complete with the until modality [12]. In Section 4, we discuss cases where the size of the validity tests grows large, and we explore a restriction that lets us avoid validity tests during checking.

Returning to the run in Table 2, the states are annotated with 2: $\neg \diamond$ test $\left(o_{1}\right)$ and $\models$ $\neg \diamond \operatorname{test}\left(o_{1}\right) \Rightarrow \neg \diamond \operatorname{test}\left(o_{1}\right)$, since $\varphi \Rightarrow \varphi$ is a propositional tautology. So $(R, i, v) \models$ $\operatorname{by}_{\{2\}}(\neg \diamond$ test $(x))$ when $v(x)=o_{1}$.

We can evaluate 1.o: $d(x) \wedge \neg \operatorname{by}_{\{2\}}(\neg \diamond$ test $(x)) \leadsto \diamond$ test $(x)$ similarly by annotating states with $\diamond \operatorname{test}(x)$ if the precondition holds. In Table 2 , this results in an annotation of 1: $\diamond \operatorname{test}\left(o_{2}\right)$ on the appropriate states. If $o_{2}$ is never tested, the run will be declared non-conforming (by Definition 5), but the annotation will remain. This lets a law which depends on (1) draw the correct inference.

The semantic evaluation outlined above works only when the references are acyclic, since an order of evaluation needs to be defined. To handle cycles, we move to a threevalued logic where the third (middle) value stands for undetermined. Initially, all statements are undetermined, and there are no annotations. At each step we assign truth values and annotations, using truth values and annotations from the previous step, until we reach a fixed point. In a companion paper [10], which focusses on the design of the logic, we prove that there is a least fixed point, which can be computed in an iterative fashion. In this paper, we use the existence of the least fixed point to derive a runtime checking algorithm.

## 4 Runtime Checking of Specifications with References

### 4.1 An Algorithm for Evaluating Specifications with References

We augment the evaluation procedure of the rule-based formalism Eagle [5] to handle references. In Eagle, formulas in LTL are evaluated by transforming them into other formulas, and discharging the remainder (if any) at trace end. The update calculus used in [5] provides a general treatment of past modalities and data dependencies. To simplify presentation, we will work directly with the formulas in the logic.

The key idea is to treat the predicate $\mathrm{by}_{\mathrm{Id}}(\varphi)$ as kind of eventuality. As we discussed in Section 3.3, to evaluate $\operatorname{by}_{\mathrm{Id}}(\varphi)$ at time $i$, we need to check the annotations obtained from the laws in Id at time $i$. If the preconditions of the laws in Id are temporal, we need to wait until they are evaluated before the annotations are obtained. So, we need to keep annotations for a time $i$ until all subformulas $\operatorname{by}_{\mathrm{Id}}(\varphi)$ for time $i$ have been evaluated. Given $\operatorname{by}_{\operatorname{Id}}(\varphi)$ and a time $i$, we attempt to evaluate it using the current set of annotations. If we cannot determine the truth value, $\operatorname{by}_{\mathrm{Id}}(\varphi)$ is transformed into $\operatorname{by}_{\operatorname{Id}}(\varphi, i)$ (read as " $\operatorname{by}_{\operatorname{Id}}(\varphi)$ is true at time $i$ "), and evaluated at subsequent times.

Following [13], we use a three-valued logic with values from $\mathcal{B}^{3}=\{\top, \perp, ?\}$, with the meaning true, false, and undetermined, respectively. For notational simplicity, we use truth values as terms in preconditions:

Definition 8 (Syntax of Preconditions). Given sets $\Phi_{1}, \ldots, \Phi_{n}$ (of predicate names), $a$ set of variables $X$, and a finite set of identifiers $I D$, the language $L^{\prime}\left(\Phi_{1}, \ldots, \Phi_{n}, X, I D\right)$, abbreviated as $L^{\prime}$, is the smallest set such that:

- If $t \in \mathcal{B}^{3}$, then $t \in L^{\prime} . p\left(y_{1}, \ldots, y_{j}\right) \in L^{\prime}$, where $p \in \Phi_{j}$ and $\left(y_{1}, \ldots, y_{j}\right) \in X^{j}$.
- If $\varphi \in L^{\prime}$, then $\neg \varphi \in L^{\prime}$ and $\square \varphi \in L^{\prime}$. If $\varphi, \psi \in L^{\prime}$, then $\varphi \wedge \psi \in L^{\prime}$
- If Id $\subseteq I D$ and $\varphi \in L\left(\Phi_{1}, \ldots, \Phi_{n}, X\right)$ (Definition 2), then $\operatorname{by}_{\mathrm{Id}}(\varphi) \in L^{\prime}$. In addition, for all natural numbers $i \in N, \operatorname{by}_{\mathrm{Id}}(\varphi, i) \in L^{\prime}$

The syntax of regulatory statements (Definition 3) is modified so that the preconditions of laws are statements from $L^{\prime}$. The set $L^{\prime}$ together with a set of regulatory statements Reg is denoted by $L^{+}=L^{\prime} \cup R e g$. Given a set of objects $O, V(X, O)$ denotes the set of all variable assigments, i.e., functions $v: X \rightarrow O$.

We can now adapt the Eagle procedure of transforming formulas. The transformation function uses two annotation functions $\alpha$ and $\alpha^{\prime}$ such that for all $i, \alpha(i) \subseteq \alpha^{\prime}(i)$. $\alpha(i)$ is the set of annotations obtained from laws with true preconditions, while $\alpha^{\prime}(i)$ is set of annotations from laws with true or undetermined preconditions. The truth of $\operatorname{by}_{\mathrm{Id}}(\varphi)$ is determined using $\alpha$, and falsity is determined using $\alpha^{\prime}$.
Definition 9 (Transformation function). Given a set of objects $O$ and annotation functions $\alpha$ and $\alpha^{\prime}$ such that $\alpha(i) \subseteq \alpha^{\prime}(i)$ for all $i \in N$, the transformation function $\tau_{\left(\alpha, \alpha^{\prime}\right)}: L^{+} \times S \times N \times V(X, O) \rightarrow L^{+}$is defined as follows:

- $\tau_{\left(\alpha, \alpha^{\prime}\right)}(t, s, i, v)=t$ if $t \in \mathcal{B}^{3}$.
- $\tau_{\left(\alpha, \alpha^{\prime}\right)}\left(p\left(y_{1}, \ldots, y_{j}\right), s, i, v\right)=\top$ if $\left(v\left(y_{1}\right), \ldots, v\left(y_{j}\right)\right) \in \pi_{j}(p, s)$.
$\tau_{\left(\alpha, \alpha^{\prime}\right)}\left(p\left(y_{1}, \ldots, y_{j}\right), s, i, v\right)=\perp$ otherwise.
- $\tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi \wedge \psi, s, i, v)=\tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, s, i, v) \wedge \tau_{\left(\alpha, \alpha^{\prime}\right)}(\psi, s, i, v)$.
$\tau_{\left(\alpha, \alpha^{\prime}\right)}(\neg \varphi, s, i, v)=\neg \tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, s, i, v)$
- $\tau_{\left(\alpha, \alpha^{\prime}\right)}(\square \varphi, s, i, v)=\tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, s, i, v) \wedge \square \varphi$
- $\tau_{\left(\alpha, \alpha^{\prime}\right)}(\operatorname{by} \operatorname{Id}(\varphi), s, i, v)=\tau_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by}_{\mathrm{Id}}(\varphi, i), s, i, v\right)$

$$
\begin{aligned}
& \quad \tau_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by} \mathrm{y}_{\mathrm{Id}}(\varphi, j), s, i, v\right)=\left\{\begin{array}{l}
\top \text { if } j \leq i \text { and } \bigwedge \alpha \cdot I d(j) \wedge v(\neg \varphi) \text { is not satisfiable } \\
\perp \text { if } j \leq i \text { and } \bigwedge \alpha^{\prime} . I d(j) \wedge v(\neg \varphi) \text { is satisfiable } \\
\operatorname{by} \operatorname{Id}(\varphi, j) \text { otherwise }
\end{array}\right. \\
& -\tau_{\left(\alpha, \alpha^{\prime}\right)}(\text { id.o. }: \varphi \leadsto \psi, s, i, v)=\text { id.o: } \tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, s, i, v) \leadsto \tau_{\left(\alpha, \alpha^{\prime}\right)}(\psi, s, i, v) \\
& \tau_{\left(\alpha, \alpha^{\prime}\right)}(\text { id.p: } \varphi \leadsto \psi, s, i, v)=\text { id.p: } \tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, s, i, v) \leadsto \psi
\end{aligned}
$$

Note that the postcondition of permissions are not transformed, as their truth value is irrelevant. The only use of postconditions of permissions is to provide annotations. To update the annotation function, we need to know if a precondition has become true or false. We now define a function to map formulas to truth values:
Definition 10. Given a set of objects $O$ and annotation functions $\alpha$ and $\alpha^{\prime}$ such that $\alpha(i) \subseteq \alpha^{\prime}(i)$ for all $i \in N$, the function value ${ }_{\left(\alpha, \alpha^{\prime}\right)}: L^{+} \times S \times N \times V(X, O) \rightarrow \mathcal{B}^{3}$ is defined as follows:

- Truth values, predicates, conjunction and negation are handled in the usual way.
- $\operatorname{value}_{\left(\alpha, \alpha^{\prime}\right)}(\square \varphi, s, i, v)=\top$ if $s$ is the final state.
value $_{\left(\alpha, \alpha^{\prime}\right)}(\square \varphi, s, i, v)=$ ? otherwise.
- $\operatorname{value}_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by}_{\mathrm{Id}}(\varphi), s, i, v\right)=\operatorname{value}_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by}_{\mathrm{Id}}(\varphi, i), s, i, v\right)$

$$
\boldsymbol{v a l u e}_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by}_{\operatorname{Id}}(\varphi, j), s, i, v\right)=\left\{\begin{array}{l}
\top \text { if } j \leq i \text { and } \bigwedge \alpha \cdot \operatorname{Id}(j) \wedge v(\neg \varphi) \text { is not satisfiable } \\
\perp \text { if } j \leq i \text { and } \bigwedge \alpha^{\prime} . \operatorname{Id}(j) \wedge v(\neg \varphi) \text { is satisfiable } \\
? \text { otherwise }
\end{array}\right.
$$

$$
\begin{aligned}
\text { - } \left.^{\text {value }_{\left(\alpha, \alpha^{\prime}\right)}(i d . \mathbf{o}: ~} \varphi \leadsto \psi, s, i, v\right) & =\operatorname{value}_{\left(\alpha, \alpha^{\prime}\right)}(\varphi \Rightarrow \psi, s, i, v) . \\
\text { value }_{\left(\alpha, \alpha^{\prime}\right)}(i d . \mathbf{p}: ~ & \leadsto \psi, s, i, v)
\end{aligned}
$$

At the end of the trace, subformulas $\square \varphi$ are replaced by $\top$, but subformulas $\operatorname{by}_{\text {Id }}(\varphi, j)$ may still be undetermined. This is due to the fact that with circular references, we can create paradoxical statements $-\mathrm{id.o:} \neg \operatorname{by}_{\{\mathrm{id}\}}(\varphi) \leadsto \varphi$. This statement requires $\varphi$ to hold when it doesn't require $\varphi$, and is always undetermined.

```
Update \(\left(\operatorname{Reg}, \Phi, \alpha, \alpha^{\prime}, s, i\right)\) :
Input: The regulation Reg, the set of formulas to be updated \(\Phi\), the annotation functions
        \(\alpha\) and \(\alpha^{\prime}\), the state \(s\) and time \(i\)
Let \(\alpha(i)=\alpha^{\prime}(i)=\emptyset\);
for all id. \(\mathbf{x}: \varphi \leadsto \psi \in\) Reg and assignments \(v\) do
    Let \(\phi=\tau_{\left(\alpha, \alpha^{\prime}\right)}(\mathrm{id} . \mathbf{x}: \varphi \leadsto \psi, s, i, v)\);
    \(\Phi=\Phi \cup\{(\phi\), id: \(v(\psi), i, v)\}\), and \(\alpha^{\prime}(i)=\alpha^{\prime}(i) \cup\{\) id: \(v(\psi)\}\)
end
repeat
    for all (id. \(\mathbf{x}: ~ \varphi \leadsto \psi, a, j, v) \in \Phi\) do
        If value \((\varphi, s, i, v)=\mathrm{T}\), then \(\alpha(j)=\alpha(j) \cup\{a\} ;\)
        If value \((\varphi, s, i, v)=\perp\), then \(\alpha^{\prime}(j)=\alpha^{\prime}(j)-\{a\}\)
    end
    Let \(\Phi^{\prime}=\emptyset\);
    for all (id. \(\mathbf{x}: \varphi \sim \psi, a, j, v) \in \Phi \mathbf{d o}\)
            Let \(\phi=\tau_{\left(\alpha, \alpha^{\prime}\right)}(\) id. \(\mathbf{x}: ~ \varphi \leadsto \psi, s, i, v)\) and \(\varphi^{\prime}=\tau_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, s, i, v)\);
            If value \((\phi, s, i, v)=\) ? or value \(\left(\varphi^{\prime}, s, i, v\right)=\) ?, then \(\Phi^{\prime}=\Phi^{\prime} \cup\{(\phi, j)\}\);
            If value \((\phi, s, i, v)=\perp\), then raise alarm.
        end
    \(\Phi=\Phi^{\prime}\)
until \(\alpha\) and \(\alpha^{\prime}\) do not change ;
```

Algorithm 1: An algorithm for evaluating statements with references

We note that the function value ${ }_{\left(\alpha, \alpha^{\prime}\right)}$ does not determine a formula to be true or false as early as possible. To decide if a formula is true as early as possible, we need to check whether all possible suffixes to the trace satisfy the formula, as in [13]. In other words, we need to decide if the transformed formula is valid. In [10], we show that with references one can encode formulas in first-order logic as regulations, and as a result, the validity problem is undecidable for $L^{+}$. The satisfiability tests used to evaluate the inference predicates are in propositional LTL, and are decidable.

Fixed points are defined at the level of a run. Suppose we are given a body of regulation $R e g$, a run $R$ and annotation functions ( $\alpha_{1}, \alpha_{1}^{\prime}$ ). The result of evaluation gives us new annotations $\left(\alpha_{2}, \alpha_{2}^{\prime}\right)$ corresponding to laws that have true preconditions $\left(\alpha_{2}\right)$, and true or undetermined preconditions $\left(\alpha_{2}^{\prime}\right)$. We will say that $\left(\alpha_{1}, \alpha_{1}^{\prime}\right)$ is a fixed point iff $\left(\alpha_{1}, \alpha_{1}^{\prime}\right)=\left(\alpha_{2}, \alpha_{2}^{\prime}\right)$.

The function value $\left(\alpha, \alpha^{\prime}\right)$ is extended to runs. The definition remains identical except that for $\operatorname{by}_{\text {Id }}(\varphi, j)$ we do not require that $j \leq i$ to determine truth or falsity, and for the temporal operator:

$$
\operatorname{value}_{\left(\alpha, \alpha^{\prime}\right)}(\square \varphi, R, i, v)=\left\{\begin{array}{l}
\top \text { if for all } j \geq i, \text { value }_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, R, i, v)=\top \\
\perp \text { if there exists } j \geq i, \text { value }_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, R, i, v)=\perp \\
? \text { otherwise }
\end{array}\right.
$$

Definition 11 (Consistent Annotations). Given a body of regulation Reg and a run $R$ with a set of objects $O$, the pair of annotation functions $\left(\alpha, \alpha^{\prime}\right)$ is consistent iff for all $(i d . \mathbf{x}: \varphi \leadsto \psi, i, v) \in \operatorname{Reg} \times N \times V(X, O)$ :

If id: $v(\psi) \in \alpha(i) \cap \alpha^{\prime}(i)$, then value $_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, R, i, v)=\top$
If id: $v(\psi) \notin \alpha(i) \cup \alpha^{\prime}(i)$, then value ${ }_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, R, i, v)=\perp$
In addition, for all $i$, we require that $\alpha(i) \subseteq \alpha^{\prime}(i)$.
Definition 12 (Fixed Point). Given a body of regulation Reg and a run $R$ with a set of objects $O$, the pair of consistent annotation functions $\left(\alpha, \alpha^{\prime}\right)$ is a fixed point iff for all $(i d . \mathbf{x}: \varphi \sim \psi, i, v) \in \operatorname{Reg} \times N \times V(X, O)$ :

If value ${ }_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, R, i, v)=\top$, then id: $v(\psi) \in \alpha(i) \cap \alpha^{\prime}(i)$
If value $_{\left(\alpha, \alpha^{\prime}\right)}(\varphi, R, i, v)=$ ?, then id: $v(\psi) \in \alpha^{\prime}(i)-\alpha(i)$
Otherwise, id: $v(\psi) \notin \alpha(i) \cup \alpha^{\prime}(i)$
We say that $\left(\alpha_{1}, \alpha_{1}^{\prime}\right) \leq\left(\alpha_{2}, \alpha_{2}^{\prime}\right)$ if for all $i$, we have $\alpha_{1}(i) \subseteq \alpha_{2}(i)$. We now review some results that are proved in [10]. The partially ordered set of consistent annotations has a least fixed point and one or more maximal fixed points. Distinct fixed points arise if there are circular references. The converse is not necessarily true, i.e., there may be circular references and a unique fixed point. There is a smallest element in the set of consistent annotations $\left(\alpha_{0}, \alpha_{0}^{\prime}\right)$ such that for all $i, \alpha_{0}(i)=\emptyset$ and $\alpha_{0}^{\prime}(i)$ contains all annotations. The least fixed point can be obtained iteratively using ( $\alpha_{0}, \alpha_{0}^{\prime}$ ).

Algorithm 1 describes the procedure for computing the least fixed point in a runtime setting. In addition to $\alpha$ and $\alpha^{\prime}$, we maintain a set of tuples $\Phi$, where each element is a transformed regulatory statement, the associated annotation, time and variable assignment. Given (id.x: $\varphi \leadsto \psi, a, j, v) \in \Phi$, if $\varphi$ is determined to be true, the annotation $a$ is added to $\alpha(j)$. On the other hand, if $\varphi$ is determined to be false $a$ is removed from $\alpha^{\prime}(j)$. For all $j \in N, \alpha(j)$ increases monotonically, and $\alpha^{\prime}(j)$ decreases monotonically with each execution of the repeat loop, until a fixed point is reached.

### 4.2 Complexity Analysis by Example

The complexity of Algorithm 1 in each state of a run depends on two factors - the number of steps necessary to reach a fixed point, and the size of satisfiability problem instances that need to be handled in the evaluation of the predicate $\operatorname{by}_{\text {Id }}(\varphi)$. We discuss examples that illustrate these two aspects, by encoding the graph reachability problem in different ways. In the first example, the number of steps taken to reach the fixed point grows with the number of objects. In the second example, the size of the satisfiability instances grows with the number of objects.

Both examples operate on the same model, where a state in the run contains a description of a graph. Objects $o_{1}$ and $o_{2}$ represent nodes, and the predicate $\delta\left(o_{1}, o_{2}\right)$ is
true iff there is an edge between $o_{1}$ and $o_{2}$. In addition, $\delta^{+}\left(o_{1}, o_{2}\right)$ is true iff there is a path from $o_{1}$ to $o_{2}$. Suppose we wish to check whether $\delta^{+}$has been computed correctly.
Example 1. Consider a self-referential sentence:
id.o: $\delta(x, z) \vee\left(\delta(x, y) \wedge \operatorname{by}_{\{\mathrm{id}\}}\left(\delta^{+}(y, z)\right)\right) \leadsto \delta^{+}(x, z)$
The precondition of this sentence corresponds to the definition of a path. In other words, there is a path between $x$ and $z\left(\delta^{+}(x, z)\right)$, if there is an edge between $x$ and $y(\delta(x, y))$, and a path between $y$ and $z\left(\operatorname{by}_{\{\mathrm{id}\}}\left(\delta^{+}(y, z)\right)\right)$. Consider the sequence of annotations obtained in the least fixed point computation $-\alpha_{0}, \ldots, \alpha_{f}$. It is easy to see that id: $\delta^{+}\left(o, o^{\prime}\right) \in \alpha_{j}(i)$ iff there is a path of length at most $j$ from $o$ to $o^{\prime}$. Given a graph with $|O|$ nodes, there is a path from $o$ to $o^{\prime}$ iff there is a path of length at most $|O|$ from $o$ to $o^{\prime}$. As a result, the fixed point will be reached in at most $|O|$ steps. The worst-case number of steps needed to reach the fixed point is $\mathbf{O}\left(m \times|O|^{k}\right)$, where $m$ is the size of the regulation, and $k$ is the maximum number of variables appearing in a sentence.
Example 2. Consider now the following statements:

$$
\begin{aligned}
& \text { A.o: } \text { by }_{\{\mathrm{B}, \mathrm{C}\}}\left(\delta^{+}(x, y)\right) \leadsto \delta^{+}(x, y) \\
& \text { B.o. } \delta(x, y) \leadsto \delta^{+}(x, y) \\
& \text { C.o: } \top \leadsto\left(\delta^{+}(x, y) \wedge \delta^{+}(y, z)\right) \Rightarrow \delta^{+}(x, z)
\end{aligned}
$$

Note that A refers to C . The presence of implication in the postcondition of C is an important feature of this example. Let, for simplicity, the graph in the state be a chain. Since the precondition of C is always true, the first step of the fixed point computation yields an annotation that contains $\mathrm{C}:\left(\delta^{+}\left(o, o^{\prime}\right) \wedge \delta^{+}\left(o^{\prime}, o^{\prime \prime}\right)\right) \Rightarrow \delta^{+}\left(o, o^{\prime \prime}\right) \in \alpha_{1}(i)$ for all $o, o^{\prime}, o^{\prime \prime}$ in the graph. The next step of the evaluation will yield the fixed point, but the size of the validity test performed in this step is $\mathbf{O}\left(|O|^{3}\right)$, as Algorithm 1 uses all the available annotations. The worst-case size of the validity instances is in $\mathbf{O}\left(m \times|O|^{k}\right)$, and the time complexity of a step in computing the fixed point is $\mathbf{O}\left(2^{m \times|O|^{k}}\right)$.
Discussion. In both examples above, Algorithm 1 checks validity instances of size polynomial in $|O|$. However, there is a crucial difference in the maximum size of tests that are needed. In Example 1, $\operatorname{by}_{\{\mathrm{id}\}}\left(\delta^{+}\left(o, o^{\prime}\right)\right)$ is true iff id: $\delta^{+}\left(o, o^{\prime}\right) \in \alpha(i)$. In other words, at most one annotation is need to evaluate by $\left\{_{\{i d\}}\left(\delta^{+}\left(o, o^{\prime}\right)\right)\right.$. In Example 2, we do need validity tests of size $|O|$ to evaluate $\operatorname{by~}_{\{\mathrm{B}, \mathrm{C}\}}\left(\delta^{+}\left(o, o^{\prime}\right)\right)$. A case study of the CFR revealed that the references behaved like Example 1 in that a single annotation or copy of the referenced statement suffices to evaluate formulas $\operatorname{by}_{\mathrm{Id}}(\varphi)$. We call this the single copy property.
Definition 13 (Single Copy Property). Given a body of regulation Reg, by $\mathrm{Id}_{\mathrm{Id}}(\varphi, j)$ has the single copy property iff for all runs $R$, and consistent annotations ( $\alpha, \alpha^{\prime}$ ):

$$
t=\left\{\begin{array}{l}
\top \text { if } \psi \wedge v(\neg \varphi) \text { is not satisfiable for some } \psi \in \alpha \cdot \operatorname{Id}(j) \\
\perp \text { if } \psi \wedge v(\neg \varphi) \text { is satisfiable for all } \psi \in \alpha^{\prime} . \operatorname{Id}(j) \\
? \text { otherwise }
\end{array}\right.
$$

where, $t=\operatorname{value}_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by}_{\mathrm{Id}}(\varphi, j), s, i, v\right)$
While the single copy property allows us to reduce the size of the satisfiability tests, we need to perform $\mathbf{O}\left(m \times|O|^{k}\right)$ tests for each inference predicate. The question arises as to whether satisfiability tests can be avoided during checking. We answer this question positively in the following section.

### 4.3 Pre-computing Satisfiability

Algorithm 1 evaluates $\operatorname{by}_{\mathrm{Id}}(\varphi)$ using satsfiabilty tests. The size of the satisfiability tests depends on $\alpha . I d(i)$, which in turn depends on the number of objects. If $\operatorname{by}_{\operatorname{Id}}(\varphi)$ has the single copy property, we can consider smaller satisfiability tests. In this section, we show that the single copy property gives us a way to assess satisfiability symbolically, and use tests of lower complexity during checking.

The strategy we use is as follows. Given a body of regulation, we perform a compilation step which involves: a) testing satisfiability, and b) replacing the predicates by $_{\operatorname{Id}}(\varphi)$ by equivalent formulas in another logic. We begin by discussing two examples, and then formalize the compilation step.
Example 1: Consider our regulatory sentences:

$$
\begin{aligned}
& \text { - 1.0: } d(x) \wedge \neg \operatorname{by}_{\{2\}}(\neg \diamond \text { test }(x)) \leadsto \diamond \text { test }(x) \text {, and } \\
& \text { - 2.p: } d(y) \wedge s p(y) \sim \neg \diamond \text { test }(y)
\end{aligned}
$$

Consider a state at which $o_{1}, o_{2}, \ldots, o_{n}$ are source plasma donations. This would result in $\neg \diamond$ test $\left(o_{1}\right), \neg \diamond$ test $\left(o_{2}\right), \ldots, \neg \diamond$ test $\left(o_{n}\right)$ being available as annotations. To evaluate by $_{\text {Id }}\left(\neg \diamond\right.$ test $\left.\left(o_{i}\right)\right)$, Algorithm 1 uses all the annotations in the satifiability test. However, in this case, it suffices to check if $\neg \diamond t e s t\left(o_{i}\right)$ is present as an annotation. The other annotations are irrelevant. To check if $\neg \diamond$ test $\left(o_{i}\right)$ is present as an annotation, it suffices to evaluate the precondition of the referenced law, i.e., whether $d\left(o_{i}\right) \wedge s p\left(o_{i}\right)$ is true (whether $o_{i}$ is a donation of source plasma). Instead of evaluating $\operatorname{by}_{\mathrm{Id}}(\phi)$ using satisfiability tests, we will check if the precondition of a referenced law is true.

Informally, the compilation step involves anwering the question when does statement 2 provide an exception for statement 1 . Equivalently, when does $\neg \diamond$ test $(y)$ imply $\neg \diamond \operatorname{test}(x)$. The answer is only when $y=x$. We can then evaluate the precondition of 2 with $y$ replaced by $x$, i.e., $d(x) \wedge s p(x)$. This lets us replace statement 1 with 1.0 : $d(x) \wedge \neg(d(x) \wedge s p(x)) \leadsto \diamond \operatorname{test}(x)$, which is equivalent to 1.0: $d(x) \wedge \neg s p(x) \sim$ $\diamond t e s t(x)$. Observe that this is the derived obligation implied by statements 1 and 2, i.e., every non-source plasma donation must be tested.
Example 2: The example above is simple in two ways: a) the number of variables in both statements are the same, and $b$ ) the references are acyclic. We discuss the general case in the context of the reachability example we saw in the previous section:

$$
\text { id.o: } \delta(x, z) \vee\left(\delta(x, y) \wedge \operatorname{by}_{\{\mathrm{id}\}}\left(\delta^{+}(y, z)\right)\right) \leadsto \delta^{+}(x, z)
$$

We observe that the precondition is structurally similar to a procedure that checks if a path exists between two nodes $x$ and $z$. That is, if $\delta(x, z)$ then $\delta^{+}(x, z)$ is true. Otherwise, if there exists $y$ such that $\delta(x, y)$ and there is a path from $y$ to $z$, then $\delta^{+}(x, z)$ is true, otherwise false.

We will produce a formula which mimics the procedure. There are two pieces of machinery used by the procedure that are not directly available in the logic: a) an existential quantifier over objects (there exists $y$ ), and $\mathbf{b}$ ) a mechanism for recursion. To address this, let us consider a logic which extends PredLTL with existential quantifiers, and a function symbol $\mathrm{P}_{\mathrm{id}}$ for $i d \in I D$ ( P stands for precondition). $\mathrm{P}_{\mathrm{id}}$ takes as argument a substitution $\theta: X \rightarrow X$, which is a function from variables to variables. A substitution is represented a set of replacements $x / y$ (read as " $x$ is replaced by $y$ "), such that each variable has at most one replacement. We replace the formula above with:
id.o: $\delta(x, z) \vee\left(\delta(x, y) \wedge \exists y_{1}: \mathrm{P}_{\mathrm{id}}\left(\left\{x / y, y / y_{1}, z / z\right\}\right)\right) \leadsto \delta^{+}(x, z)$
It remains to give this formula a semantics. Given a variable assignment $v$ and a subtitution $\theta, \theta(v)$ denotes the variable assignment $v^{\prime}$ such that $v^{\prime}(x)=v(\theta(y))$. Given a run $R$, time $i$ and regulation Reg, the idea is to say that $(R, i, v) \models \mathrm{P}_{\mathrm{id}}(\theta)$ iff $(R, i, \theta(v)) \models \varphi$ where id. $\mathbf{x}: \varphi \leadsto \psi \in \operatorname{Reg}$. We now formalize the compilation procedure.
Compiling References into Precondition Tests: We begin by defining the syntax of compiled preconditions:

Definition 14 (Syntax of Compiled Preconditions). Given sets $\Phi_{1}, \ldots, \Phi_{n}$ (of predicate names), a set of variables $X$, and a finite set of identifiers ID, the language $L_{C}^{\prime}\left(\Phi_{1}, \ldots, \Phi_{n}, X, I D\right)$, abbreviated as $L_{C}^{\prime}$, is the smallest set such that:

- If $t \in \mathcal{B}^{3}, t \in L_{C}^{\prime}$. And, $p\left(y_{1}, \ldots, y_{j}\right) \in L_{C}^{\prime}$ where $p \in \Phi_{j}$ and $\left(y_{1}, \ldots, y_{j}\right) \in X^{j}$.
- If $\varphi \in L_{C}^{\prime}$, then $\neg \varphi \in L_{C}^{\prime}$ and $\square \varphi \in L_{C}^{\prime}$. If $\varphi, \psi \in L_{C}^{\prime}$, then $\varphi \wedge \psi \in L_{C}^{\prime}$
- If $\varphi \in L_{C}^{\prime}$, for all $y \in X$, we have $\exists y: \varphi \in L_{C}^{\prime}$.
- For all id $\in I d$ and substitutions $\theta: X \rightarrow X$, we have $\mathrm{P}_{\mathrm{id}}(\theta) \in L_{C}^{\prime}$. In addition, for all natural numbers $i \in N, \mathrm{P}_{\mathrm{id}}(\theta, i) \in L_{C}^{\prime}$

The syntax of regulatory statements (Definition 3) is modified so that the preconditions of laws are statements from $L_{C}^{\prime}$. The set $L_{C}^{\prime}$ together with a set of regulatory statements $R e g_{C}$ is denoted by $L_{C}^{+}=L_{C}^{\prime} \cup R e g_{C}$. We remind the reader that $L^{+}$and $L^{\prime}$ are the languages with the predicate $\mathrm{by}_{\mathrm{Id}}(\varphi)$.

The semantics of $L_{C}^{+}$is defined in a manner similar to $L^{+}$. Rather than using annotations ( $\alpha, \alpha^{\prime}$ ), we now evaluate statements w.r.t. two sets of assignment functions $\left(\gamma, \gamma^{\prime}\right) . \gamma(i, i d)$ (resply., $\gamma^{\prime}(i, i d)$ ) is a set of variable assignments for which the precondition of the law with identifier $i d$ is true (resply., true or undetermined). As with annotations, we require that for all $i \in N$ and $i d \in I D, \gamma(i, i d) \subseteq \gamma^{\prime}(i, i d)$. Given an assignment $v$ and a substitution $\theta, \theta(v)$ denotes the assignment $\overline{v^{\prime}}$ such that for all $y \in X$, we have $v^{\prime}(y)=v(\theta(y))$. We can now adapt the value function:

$$
\text { value }_{\left(\gamma, \gamma^{\prime}\right)}\left(\mathrm{P}_{\mathrm{id}}(\theta, j), R, i, v\right)=\left\{\begin{array}{l}
\top \text { if } \theta(v) \in \gamma(j, i d) \\
\perp \text { if } \theta(v) \notin \gamma^{\prime}(j, i d) \\
? \text { otherwise }
\end{array}\right.
$$

The definitions of consistency and fixed points (Definitions 11 and 12) are easily adapted, and we leave the details to the reader.

We now describe the compilation procedure. Given $\varphi \in L^{+}$, we use $X(\varphi)$ to denote the set of variables appearing in $\varphi$, and $\theta(\varphi)$ to denote the formula obtained by performing the substitution $\theta$ on $\varphi$. Consider by $\operatorname{Id}_{\mathrm{Id}}(\varphi, j)$, which has the single copy property, and variables disjoint from all regulatory statements:

- Let $\mathcal{S}(\varphi, i d)=\{\theta \mid$ id. $\mathbf{x}: \varphi \sim \psi \in \operatorname{Reg}$, and $\theta(\psi \Rightarrow \varphi)$ is valid $\}$.
- For all $\theta \in \mathcal{S}(\varphi, i d)$, let $\varphi_{C}(\theta, i d)=\exists z_{1}, \ldots, z_{m}: \mathrm{P}_{\mathrm{id}}(\theta, j)$, where the existentially quantified variables are in one-to-one correspondence with those in $X(\phi)-X(\psi)$. More formally, $z_{j} \notin X(\phi)-X(\psi)$ and $\theta$ is a one-to-one function from $\left\{z_{j} \mid 1 \leq\right.$ $j \leq m\}$ to $X(\phi)-X(\psi)$.
- $\varphi_{C}\left(\operatorname{by}_{\text {Id }}(\varphi, j), i d\right)=\bigvee\left\{\varphi_{C}(\theta, i d) \mid \theta \in \mathcal{S}(\varphi, i d)\right\}$, and
- $\varphi_{C}\left(\operatorname{by}_{\mathrm{Id}}(\varphi, j)\right)=\bigvee\left\{\varphi_{C}\left(\operatorname{by} \mathrm{Id}_{\mathrm{Id}}(\varphi, j), i d\right) \mid i d \in I d\right\}$

We note that the first step makes crucial use of the single copy property (SCP). In computing $\mathcal{S}(\varphi, i d)$, it suffices to find substitutions such that $\theta(\psi \Rightarrow \varphi)$ is valid. If the SCP does not hold, then we need to check if multiple copies of postconditions provide the necessary implication (as in Example 2, Section 4.2). For example, we need to check if $\theta\left(\psi_{1} \wedge \ldots \wedge \psi_{n} \Rightarrow \varphi\right)$, where $\psi_{1}, \ldots, \psi_{n}$ are copies of the postcondition of a law with the variables renamed. It can be shown that detecting whether the SCP holds is undecidable. In future work, we plan to investigate restrictions on postconditions that make SCP-detection decidable.

To prove the correctness of the compilation procedure, we use a notion of correspondence between annotations and assignments. Let us assume as given a body of regulation $\operatorname{Reg}$ (in $L^{+}$), a run $R$ and consistent annotations ( $\alpha, \alpha^{\prime}$ ). Rather than producing a regulation in $L_{C}^{+}$, we prove correctness by evaluating formulas in $L_{C}^{\prime}$ against Reg. We construct $\left(\gamma_{\alpha}, \gamma_{\alpha^{\prime}}^{\prime}\right)$ such that for all $i \in N$ and $i d \in I D, v \in \gamma_{\alpha}(i, i d)$ iff id: $v(\psi) \in \alpha(i)$, and $v \in \gamma_{\alpha^{\prime}}^{\prime}(i, i d)$ iff id: $v(\psi) \in \alpha^{\prime}(i)$. We can now show the following:

Lemma 1. Given a body of regulation Reg, a run $R$, consistent annotations ( $\alpha, \alpha^{\prime}$ ), and $\operatorname{by}_{\operatorname{Id}}(\varphi, j)$ which has the single copy property, for all $i \in N$ and assignments $v$ :
value $_{\left(\alpha, \alpha^{\prime}\right)}\left(\operatorname{by}_{\text {Id }}(\varphi, j), R, i, v\right)=\operatorname{value}_{\left(\gamma_{\alpha}, \gamma_{\alpha^{\prime}}^{\prime}\right)}\left(\varphi_{C}\left(\operatorname{by}_{\text {Id }}(\theta, j)\right), R, i, v\right)$
Proof. The proof follows straightforwardly from the construction of $\varphi_{C}\left(\operatorname{by}_{\mathrm{Id}}(\theta, j)\right)$ and the single copy property. We sketch one of the cases.

Suppose value $\left(\alpha, \alpha^{\prime}\right)\left(\operatorname{by}_{\text {Id }}(\varphi, j), R, i, v\right)=\mathrm{T}$. There exists id: $v^{\prime}(\psi) \in \alpha(i)$ such that $v^{\prime}(\psi) \wedge v(\neg \varphi)$ is not satisfiable, or equivalently $v^{\prime}(\psi) \Rightarrow v(\varphi)$ is valid. It follows that there exists a substitution $\theta$ such that $\theta(\psi \Rightarrow \varphi)$ is valid. By definition $v^{\prime} \in \gamma_{\alpha}(i)$, and hence, value $\left(\gamma_{\alpha}, \gamma_{\alpha^{\prime}}^{\prime}\right)\left(\mathrm{P}_{\mathrm{id}}(\theta, j), R, i, v^{\prime}\right)=\mathrm{T}$. We can then argue using the construction that value ${ }_{\left(\gamma_{\alpha}, \gamma_{\alpha^{\prime}}^{\prime}\right)}\left(\exists z_{1}, \ldots, z_{m}: \mathrm{P}_{\mathrm{id}}(\theta, j), R, i, v\right)=\top$, and as a result, value ${ }_{\left(\gamma_{\alpha}, \gamma_{\alpha^{\prime}}^{\prime}\right)}\left(\varphi_{C}\left(\operatorname{by}_{\mathrm{Id}}(\theta, j)\right), R, i, v\right)=\top$. The other cases are handled similarly.

Given Reg in which all subformulas $\mathrm{by}_{\mathrm{Id}}(\varphi)$ have the single copy property, we can now produce the regulation $R e g_{C}$ in $L_{C}^{+}$with all occurences of $\operatorname{by}_{\mathrm{Id}}(\varphi)$ replaced by $\varphi_{C}\left(\operatorname{by}_{\mathrm{Id}}(\varphi)\right)$. It follows from Lemma 1 that if $\left(\alpha, \alpha^{\prime}\right)$ is a fixed point w.r.t. Reg, then $\left(\gamma_{\alpha}, \gamma_{\alpha^{\prime}}^{\prime}\right)$ is a fixed point w.r.t. $R e g_{C}$. In addition, the truth values assigned to regulatory statements are identical.

The complexity of evaluation depends on the number of disjuncts in $\varphi_{C}(\operatorname{by} \operatorname{Id}(\varphi))$, which in turn depends on the size of the set: $\mathcal{S}(\varphi, i d) .|\mathcal{S}(\varphi, i d)| \leq(2 k)^{2 k}$, where $k$ is the maximum number of variables in a regulatory statement. $(2 k)^{2 k}$ is a bound on the number of equivalence classes, i.e., we have $2 k$ variables ( $k$ in $\varphi$ and $k$ in $\psi$ ) and at most one equivalence class for each variable. Hence, the size of $\varphi_{C}\left(\operatorname{by}_{\mathrm{Id}}(\varphi)\right)$ is $\mathbf{O}\left(m \times(2 k)^{2 k}\right)$, where $m$ is the number of regulatory statements. Each quantified precondition test can be evaluated in $\mathbf{O}\left(|O|^{k}\right)$ time, where $O$ is the set of objects. As a result, the time complexity for evaluating $\varphi_{C}\left(\operatorname{by} \operatorname{Id}^{(\varphi)}\right)$ is $\mathbf{O}\left(m \times(2 k)^{2 k} \times|O|^{k}\right)$. We now describe an evaluation of the system, using a prototype implementation which performs this compilation procedure.

### 4.4 Evaluation

We have developed a prototype implementation of the checker. We briefly describe two aspects of the implementation: (a) the interface between regulations and traces (schemas), and (b) the trace-checker.

Schemas form the interface between the regulation and trace. A schema is a set of class and type definitions. Classes can inherit from others, and have attributes which have atomic types, tuples or unions of types, pointers to other objects or sets of values.

Our current implementation of the trace-checker is static in the sense that the entire trace is stored on disk (in an NDBM database). The objects at each state belong to classes in a given schema. The regulation, which is type-checked against the same schema, is compiled using the techniques discussed in Section 4.3, and evaluated at each state. We do not have any special optimizations for speed. The objects are stored as strings, and reparsed every time they are loaded into memory. The checker evaluates each obligation w.r.t. all variable assignments, loading into memory a single variable assignment at a time.

We now describe a preliminary evaluation of the implementation. Our goal was to check if we could scale to traces with a large number of objects, rather than very long traces. We created a schema based on the CFR, capturing donors, donations of several types, and various tests. We then checked a number of synthetic (final) states for conformance. Given a schema, we generate a set of donors by choosing random values for atomic attributes. For each donor we generate a set of donations again choosing attribute values at random. Each donation is randomly tested as follows: with $p=0.3$ it is tested for all diseases with negative results, with $p=0.3$ it is test for diseases with a random result, and otherwise it is not tested.

On the regulatory side, we created logic formulas for a portion of the CFR 610.40. A total of 12 sentences, and a list of 6 disease names were used. Lists are frequent in regulation, and statements refer to particular list items. Of the 12 sentences, 7 were obligations and 5 were permissions. A total of 8 reference formulas $\left(\operatorname{by}_{\operatorname{Id}}(\varphi)\right)$ were used, and of these 3 referred to list items. The compilation step of removing the references took 26 seconds with a total of 96 satisfiability tests. Each statement had at most 2 variables (one for donations and the other for disease names).

We evaluated performance of the checker against a number of states. The number of disease names was 8 , and the number of donations varied. The time taken varied linearly with a number of donations. For states with $100,1000,5000$, and 10000 donations the conformance check took 12s, 130s, 500s, and 1042s respectively. The performance suggests that the approach is practical for checking short traces. However, more incremental algorithms are needed to deploy such specifications in a runtime setting.

## 5 Discussion and Conclusions

We have described a logic for representing regulatory documents for the application of conformance checking. The logic allows statements to refer to others for conditions or exceptions. While references give us a way to represent regulation directly, the evaluation of references during checking has high complexity. Algorithm 1 uses satisfiability tests of size polynomial in the number of objects. In Sections 4.2 and 4.3, we
described an emprically motivated assumption (the single copy property), which lets us replace satifiablity tests with tests of lower complexity. The evaluation of our prototype implementation suggests that the approach is suitable for conformance audits of medium-sized traces.

An important part of making this approach useful in practice is to provide support for translating the regulatory documents into their formal representation. Such support has to rely heavily on natural language processing techniques, which require substantial extension of current state of the art. We are actively pursuing this line of research. Preliminary results are reported in $[14,15]$.
Related Work. The use of logic to represent and reason about regulation has been of interest for several years. We begin by discussing the literature in relation to two issues: a) the representation of obligation and permission, and b) references between laws. We compare our work with other approaches to conformance checking, and place it in the context of previous work on run-time checking of LTL.

The goal of deontic logic is a formalization of concepts such as obligation, permission and rights. There are many systems of deontic logic, but the most common approach is to treat obligation and permission as modal operators [16, 17]. The logic developed here focusses on the problem of references between laws, and we believe that the representation of obligation and permission is an important but orthogonal issue. In future work, we plan to add a modal treatment of obligation and permission to our system.

The problem of references between laws has been observed in regulatory texts in different domains [18,2]. More generally, the function of sentences as conditions or exceptions to others has been studied in a variety of contexts. Alchourron and Makinson [19] proposed a hierarchical structure for a legal discourse, to handle exceptions to statements. This led to the development of input-output logic [20], which is closely related to default logic [11]. Previous work on applying default logic has been mainly in the context of computing extensions to a theory, in the manner of logic programs [7, $18,6]$. We believe that the application of these ideas in conformance checking is novel.

Conformance checking has been receiving increasing attention in recent years [13,21]. [1] represents business contracts as SQL queries. [3,21] use a logic on a UML description of a domain. While the approaches of $[1,3,21]$ are similar in spirit to ours, they do not provide a treatment of references. [2] discusses the problem of references in the context of privacy regulation, and the references are resolved manually.

Our work builds upon the well-established work on run-time checking of LTL and its extensions. We have adapted the calculus of Eagle [5] to handle references. Rulebased formalisms [5,22] are quite general, but the transformation of formulas at each state can be expensive. Automata-based approaches [13] offer a more efficient alternative at the price of generality. We are currently exploring ways to adapt the automatabased approach to our setting.

## References

1. Abrahams, A.: Developing and Executing Electronic Commerce Applications with Occurrences. PhD thesis, Univeristy of Cambridge (2002)
2. Breaux, T.D., Vail, M.W., Anton, A.I.: Towards regulatory compliance: Extracting rights and obligations to align requirements with regulations. In: Proceedings of the 14th IEEE International Requirements Engineering Conference. (2006)
3. Giblin, C., Liu, A., Muller, S., Pfitzmann, B., Zhou, X.: Regulations Expressed as Logical Models (REALM). In Moens, M.F., Spyns, P., eds.: Legal Knowledge and Information Systems. (2005)
4. U.S. Food and Drug Administration: Code of Federal Regulations. http://www.gpoaccess.gov/cfr/index.html
5. Barringer, H., Goldberg, A., Havelund, K., Sen, K.: Rule-based runtime verification. In: Proceedings of 5th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI 2004). Volume 2937 of LNCS. (January 2004) 44-57
6. McCarty, L.T.: A language for legal discourse - i. basic features. In: Proceedings of ICAIL. (1989)
7. Sergot, M., F.Sadri, Kowalski, R., F.Kriwaczek, P.Hammond, Cory, H.: The british nationality act as a logic program. Communications of the ACM 29(5) (1986) 370-86
8. Ross, A.: Directives and Norms. Routlege and Kegan Paul (1968)
9. Marcus, R.B.: Iterated deontic modalities. Mind 75(300) (1966)
10. Dinesh, N., Joshi, A., Lee, I., Sokolsky, O.: Reasoning about conditions and exceptions to laws in regulatory conformance checking. In Submission: http://www.cis.upenn.edu/ /nikhild/reasoning.pdf (2008)
11. Reiter, R.: A logic for default reasoning. In: Readings in nonmonotonic reasoning. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1987) 68-93
12. Sistla, A.P., Clarke, E.M.: The complexity of propositional linear temporal logic. ACM 32 (1985) 733-49
13. Bauer, A., Leucker, M., Schallhart, C.: Monitoring of real-time properties. In: Proceedings of the 26th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06). Volume 4337 of LNCS. (December 2006)
14. Dinesh, N., Joshi, A.K., Lee, I., Webber, B.: Extracting formal specifications from natural language regulatory documents. In: Proceedings of the Fifth International Workshop on Inference in Computational Semantics. (2006)
15. Dinesh, N., Joshi, A., Lee, I., Sokolsky, O.: Logic-based regulatory conformance checking. In: Proceedings of the 14th Monterey Workshop. (2007)
16. von Wright, G.H.: Deontic logic. Mind 60 (1951) 1-15
17. Aqvist, L.: Deontic logic. In Gabbay, D., Guenthner, F., eds.: Handbook of Philosophical Logic, Volume II: Extensions of Classical Logic. (1984) 605-614
18. Bench-Capon, T., Robinson, G., Routen, T., Sergot, M.: Logic programming for large scale applications in law: A formalisation of supplementary benefit legislation. In: Proceedings of the 1st International Conference on AI and Law. (1987)
19. Alchourron, C., Makinson, D.: Hierarchies of regulation and their logic. In Hilpinen, R., ed.: New Studies in Deontic Logic. (1981)
20. Makinson, D., van der Torre, L.: Input/output logics. Journal of Philosophical Logic 29 (2000) 383-408
21. Glasse, E., Engers, T.V., Jacobs, A.: Power: An integrated method for legislation and regulations from their design to their use in e-government services and law enforcement. In Moens, M.F., ed.: Digitale Wetgeving, Digital Legislation. (2003) 175-204
22. Barringer, H., Rydeheard, D., Havelund, K.: Rule systems for run-time monitoring: From Eagle to RuleR. In: Proceedings of the $7^{\text {th }}$ Workshop on Runtime Verification. Volume 4839 of LNCS. (March 2007) 111-125

[^0]:    * This research was supported in part by NSF CCF-0429948, ARO W911NF-05-1-0158, and ONR MURI N00014-04-1-0735.

[^1]:    ${ }^{1}$ When $\operatorname{by}_{\mathrm{Id}}(\varphi)$ appears in the precondition of a law, $\varphi$ need not be the negation of the postcondition. An example is statement (4) in Section 2, which can be represented as:
    4.o: by $_{\{1\}}(\diamond \operatorname{test}(z)) \leadsto \diamond \operatorname{scr}(z)$, where $\operatorname{scr}(z)$ is true iff $z$ is tested using a screening test.

