

An Algorithm to Generate Compliance Monitors from Regulations

Travis D. Breaux and Annie I. Antón

*Department of Computer Science
North Carolina State University
{tdbreaux, aianton}@eos.ncsu.edu*

Abstract

Developing software systems in heavily regulated industries requires methods to ensure systems comply with regulations and law. An algorithm to generate finite state machines (FSM) from stakeholder rights and obligations for compliance monitoring is proposed. Rights and obligations define what people are permitted or required to do; these rights and obligations affect software requirements and design. The FSM allows stakeholders, software developers and compliance officers to trace events through the invocation of rights and obligations as pre- and post-conditions. Compliance is monitored by instrumenting runtime systems to report these events and detect violations. Requirements and software engineers specify the rights and obligations, and our algorithm performs three supporting tasks: 1) identify ambiguities, 2) balance rights with obligations, and 3) generate finite state machines. Preliminary validation of the algorithm includes FSMs generated from U.S. healthcare regulations and tool support to parse these specifications and generate the FSMs.

1. Introduction

Automated Software Engineering is concerned with automating tasks within and across the software development lifecycle. Software requirements are one of the first artifacts to enter this lifecycle. Due to their informal conception through natural language, they are difficult to manipulate and integrate directly into the verification and testing of large software systems – the goal of requirements monitoring. In highly regulated industries within the United States, such as healthcare and finance, requirements monitoring is necessary to ensure systems comply with regulations and law. We propose an algorithm for automatically generating finite state machines (FSM) from semantic models of stakeholder rights and obligations. We illustrate our algorithm using model specifications extracted from the Privacy Rule in the U.S. Health Insurance Portability and Accountability Act (HIPAA) [13].

The FSMs correlate real world events (inputs and outputs) to stakeholder rights and obligations (states). To evaluate risk and compliance, software developers

can map these events to requirements and design specifications. Systems developers can use these event mappings to focus their verification efforts on those components most at-risk for non-compliance. After deployment, the FSMs can operate as compliance monitors by instrumenting software components to provide records of these events. As systems evolve to adapt to new organizational needs, these monitors ensure that new changes to existing software systems continue to operate within the regulatory framework of stakeholder rights and obligations.

It is estimated that healthcare organizations will spend \$17.6 billion over the next few years to bring their systems and procedures into compliance with the HIPAA [19]. Existing guidelines and standards not only fail to provide specific solutions, but also make compliance a significant challenge. According to a Ernst & Young survey of executives in over 1,300 international organizations, compliance with regulations and policy surpassed worms and viruses as the primary driver of information security policy 2005 [10]. The consequence of not complying with regulations is now forefront for those responsible for assuring that software systems containing sensitive information remain secure and protected.

The remainder of this paper is organized as follows: Section 2 reviews related work; Section 3 introduces the algorithm to generate compliance monitors; Section 4 presents the results of applying the algorithm to rights and obligations extracted from the HIPAA Privacy Rule [15]; Section 5 discusses limitations and future work, with our conclusion in Section 6.

2. Related Work

This section discusses approaches that model regulations, requirements and scenarios, that perform consistency and model checking, and approaches to requirements monitoring. It is important to stress that requirements are limited to the scope of software systems, whereas stakeholder rights and obligations (as expressed in law) govern the broader scope of business processes. Although requirements-based methods generally assume a degree of system control that cannot be assumed in a complete compliance

framework, these methods are highly relevant to the specification of stakeholder rights and obligations as requirements.

Much work has focused on modeling regulations in artificial intelligence [16, 22, 23, 24]. Sherman developed Prolog models from the Income Tax Act of Canada [24]. Sergot et al. have conducted similar case studies, developing logic programs from the British Nationality Act [22] and the Indian Central Service Pension rules [23]. These logic programs abstract rule elements as predicates to evaluate regulations. Our general approach has been to further decompose these predicates into semantic models [4, 5, 7]. We leverage this decomposition to align rights and obligations along shared events (see Section 3.3.2). Alternatively, Kerrigan and Law propose a system that provides question-answering assistance with environmental regulations modeled in first-order logic [16]. In addition to requirements monitoring, we believe our monitors can be used to answer compliance-related design questions based on risk.

In requirements engineering, relevant approaches include those to model policies [5], regulations [12] and stakeholder goals [17]. The notation developed from our previous work modeling privacy policies in healthcare and finance [5] is used to specify semantic models in the algorithm proposed in this paper. In addition, Giorgini et al. present Secure Tropos (ST), a formal framework for modeling security requirements applied to Italian privacy regulations [12]. ST distinguishes between rights or permissions (at-most) and obligations (at-least) in the context of delegation. ST employs Datalog to perform model checking and find inconsistencies. Our work complements their framework by providing new insight into how rights and obligations are conditioned on shared events. Landtsheer et al. show how to map KAOS goal models into Software Cost Reduction (SCR) specifications amenable to event-based model checking [17]. Goals are prescriptive actions of intent whose satisfaction may require agent cooperation. Like goals, our work with rights and obligations also express stakeholder intent yet within the expressed confines of regulations and normative theory.

Sutcliffe et al. and Maiden describe a partially automated method to generate scenarios using use cases and object system models [18, 26]. Maiden defines object system models as patterns of requirements that include attributes for agents, actions, objects, and pre-conditions, among others. Our approach differs in that our compliance monitors derive event sequences from regulations whereas use cases are generally elicited from stakeholders. Furthermore, we intend our FSMs to be used in verifying requirements in post-deployed systems,

whereas Maiden's scenarios were intended to be used in requirements validation.

We distinguish aspects of our approach from those in consistency and model checking [1, 3, 8, 13]. Heitmeyer et al. propose consistency checking as a formal method to statically identify ambiguities in requirements specifications [13]. We perform static checks on semantic models to identify ambiguities and further balance rights with obligations and call this process consistency checking. However, model checking seeks to assert formal properties across model states and has been applied to requirements specifications to: check safety properties using temporal logic [1]; reduce the number of model states using abstraction [3]; and derive FSM from design flow graphs (DFG) to check consistency between requirements and design [8]. A future goal of our work includes applying model checking techniques to the generated FSMs.

In requirements monitoring, various approaches use different techniques to specify and deploy requirements monitors. Among these, we recognize the need to: provide a generalized interface to query assumptions at runtime [9]; model relationships between events to distinguish between expected and recorded system behavior [25]; and align requirements monitoring with design methodologies [21]. Our approach seeks to accommodate these needs by generating FSMs for event-based compliance monitoring and to evaluate design decisions based on risk and compliance costs. On the other hand, Peters and Parnas discuss design issues for requirements monitors in real-time systems such as real-time notification under discrete-time sampling, sample quantization to measure error, and non-determinism [20]. However, the U.S. federal regulations we analyzed in healthcare did not exhibit these issues.

Fickas et al. describe a case study in which ephemeral requirements are modeled as finite state automata in the Promela language and monitored using a web service called Emu [11]. Ephemeral or personal requirements are difficult to monitor because the system environment is beyond the scope of reasonable control [11]. For example, advancing from one state to another may require an indefinite human response. The same limitation exists in compliance monitoring since not all regulations are implementable within the scope of software systems. However, due to diversity in regulated industries, it is infeasible for the law to prioritize implementing regulations based on factors such as available technology, business value or costs. Therefore, we believe all regulations should appear in the scope of compliance monitors, so as organizations evolve, they can evaluate their individual non-compliance risks against these factors.

3. Generating Compliance Monitors

The algorithm to generate compliance monitors (see Figure 1) accepts semantic models of regulations [7] as input and produces finite state machines (FSMs) as output in two phases: (1) a consistency checking phase to identify ambiguities and balance rights with obligations; and (2) a generation phase to produce the states-event and transition tables comprising the FSMs.

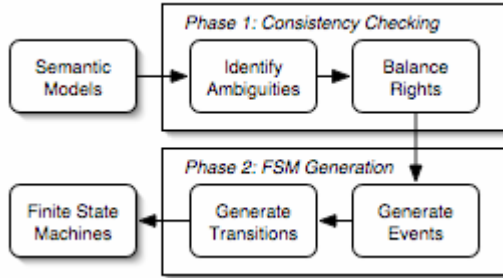


Figure 1: Algorithm Process Overview

We now describe in detail the inputs and outputs to the algorithm in Section 3.1 before proceeding to discuss the first and second phases in Sections 3.2 and 3.3, respectively.

3.1. The Input and Output

We now discuss the algorithm’s inputs and outputs. All examples employed below are derived from the HIPAA Privacy Rule [15].

3.1.1. Input. Requirements and software engineers provide semantic models of regulations as inputs to the algorithm. These semantic models describe rights, obligations or constraints: a *right* is an action a stakeholder is permitted to perform; an *obligation* is an action a stakeholder is required to perform; and a *constraint* describes either a state-of-being or past event that is a pre-condition to a right or obligation. Each right and obligation is defined by a pair (M, E) for a semantic model M of a right R_x or obligation O_x for an index x and a logical expression E comprised of individual semantic models of constraints $C_1 \dots C_n$ using logical operators (and, or, not).

We have developed a methodology that engineers can use to extract these models from policy and regulation text [7]. The methodology, which was developed from a pilot study [6] and a summative case study [7], employs a process called Semantic Parameterization to derive the models from restricted natural language statements [4, 5]. To illustrate, consider the restricted statements below, extracted from the HIPAA Privacy Rule, for the obligation pairs $(O_{4.10}, C_1 \wedge C_2)$ and $(O_{4.11}, C_1 \wedge \neg C_2)$ in which a covered entity (CE) provides protected health information (PHI) to the individual:

- C_1 : The individual requests to access PHI in a format.
- C_2 : The requested format is readily available.
- $O_{4.10}$: The CE must provide the individual with access to PHI in the requested format.
- $O_{4.11}$: The CE must provide the individual with access to PHI in a readable hard copy format.

The obligation $O_{4.10}$ requires the CE to provide access to PHI in the requested format if the individual requests the format (C_1) and the format is available (C_2). Otherwise, the obligation $O_{4.11}$ requires the CE to provide access to PHI in a readable hard-copy format.

Semantic models are expressed in the KTL notation [5], which uses two relations: (1) the alpha relation $\alpha(x, y)$, where $x\{y\}$ reads “ x has y ” and $y:x$ reads “ y of x ” and (2) the delta relation $\delta(x, y)$, where x/y and $y=x$ reads “ x is y .” Symbols in semantic models are restricted to one part-of-speech from nouns, adjectives, verbs and adverbs; articles and prepositions are not allowed. Symbols preceded by an exclamation point are negated, while symbols preceded by a question mark are query variables. Using unification [2], models can be used to query other models [5] — a procedure that is invoked throughout the algorithm. We use the term *pattern* to refer to model components that are formally defined by a query.

The model for obligation $O_{4.10}$ appears in Figure 2, expressed in the KTL notation. In Figure 2, symbols taken from the obligation statement appear in **bold**; all other symbols comprise the *meta-model* that is provided by the user or acquired from a template.

```

1 activity [ obligation ] {
2   subject = CE
3   action = provide
4   object = access {
5     subject = individual
6     action = access
7     object = PHI {
8       format [ requested ]
9     }
10  }
11  target = individual
12 }
  
```

Figure 2: Example Semantic Model

The algorithm discussed herein uses the activity pattern that requires the attributes *subject*, *action* and *object* [4, 5]. An instance of the activity pattern appears in Figure 2 for the *activity* (Line 1) and the *access* (Line 4) with the subject (Lines 2 and 5), action (Lines 3 and 6) and object (Lines 4 and 7). For a model M matching the activity pattern with values assigned to the subject s , action a , and object o in the activity, we define the function $T : M \rightarrow \langle s, a, o \rangle$ that maps M to a SAO-triple comprised of those values. For example,

$T(O_{4.10}) = \langle CE, provide, access \rangle$. The function T is implemented using static queries [5].

3.1.2. Output. The algorithm produces finite state machines as output in which each state indexes the rights and obligations for which stakeholders are accountable. To “reach a state” means to assign a right or obligation to a stakeholder; otherwise, the rights or obligations are considered to be unassigned. From each state, a stakeholder who is assigned a right may invoke that right and a stakeholder who is assigned an obligation must achieve or maintain that obligation. It is considered a violation of a right or obligation if the stakeholder cannot invoke an assigned right or cannot achieve or maintain an assigned obligation. An *event* is the performance of the entitled or obliged activity. This definition further elucidates the need for rights to be balanced with obligations.

Each state is connected by one or more transitions and each transition coincides with an event or state-of-being. The pre-conditions to rights and obligations are in-transitions to states, whereas the act of invoking a right or achieving and maintaining an obligation is an out-transition from states. For the obligation pair $(O_{4.10}, C_1 \wedge C_2)$ from the earlier example in Section 3.1.1, we derive the following events:

- E₁:** The individual requests to access PHI in a format.
- E₂:** The requested format is readily available.
- E₃:** The CE provides the individual with access to the PHI in the requested format.

The constraint C_1 maps to the event E_1 , the constraint C_2 maps to the event E_2 and the act of achieving the obligation $O_{4.10}$ maps to the event E_3 . Since E_1 and E_2 were derived from the constraints, they become in-transitions to the state that indexes $O_{4.10}$. The achievement E_3 becomes an out-transition to that state.

The state and transitions are illustrated in Figure 3.




Figure 3: Example State with Transitions

The conjunction $C_1 \wedge C_2$ maps to $E_1 \wedge E_2$ and appears as a bridge between the corresponding in-transitions. In all such figures, the transitions for events produced by invoking rights are illustrated using dotted lines whereas the transitions for events produced by achieving or maintaining obligations are illustrated using solid lines.

We now describe the individual steps in the first and second phases of the algorithm in Sections 3.2 and 3.3, respectively.

3.2. Phase 1: Checking Model Consistency

In the first phase of the compliance algorithm, semantic models are checked for ambiguities and transformed to balance rights with obligations. The procedure to perform these steps and their contribution to the second phase of the algorithm is discussed in each sub-section that follows.

3.2.1. Identify Ambiguities. Semantic models use the activity pattern [4, 5], which has three co-requisite attributes: subject, action and object. These attributes must be specified in each activity because they are required to generate events in the second phase. We automatically detect these ambiguities by applying a query algorithm based on unification [2] that proceeds as follows: for each symbol x in a model M , if x is a type of *activity* then $\alpha(x, y)$ and $\delta(y, z)$ are true for some symbol z and $y \in \{subject, action, object\}$. The algorithm compares each symbol in the model for these ambiguities and the user must resolve any ambiguities before proceeding to the next step.

In addition to the symbol *activity*, other nouns are types of activities such as a *request, denial, review, agreement*, etc. For these nouns, the action is implied by the noun (e.g., the action *agree* is implied by the noun *agreement*) and the user maintains a list of these nouns for use in the algorithm. Based on our prior work analyzing policies in healthcare and finance [4, 5], it is reasonable to expect many of these nouns are generalizable across domains. For example, in Figure 2 (above) if the object access (Line 4) were specified without the object PHI (Line 7), the algorithm would detect this ambiguity during this step and require the user to complete the specification.

3.2.2. Balance Rights with Obligations. Rights and obligations are balanced by identifying their implied rights and obligations. Implied rights or obligations may not be directly expressed but they are necessary to satisfy expressed rights and obligations. For this reason, they are logically derived from the expressed rights and obligations and we provide patterns to do so.

Balancing is automated for four cases including where rights or obligations are implied by (1) delegations, (2) direct provisions, (3) indirect provisions, and (4) an act where a stakeholder is expressly not obliged, called an *anti-obligation* [7]. Each case uses a *transformation* comprised of a unique query to match the input model and identify relevant values that are in turn mapped into an output model describing the implied right or obligation. For

example, consider the delegation right $R_{6.3}$ balanced by implied obligation $O_{R-6.3}$:

R_{6.3}: The CE may require an individual to request in writing that the CE amend their PHI.

O_{R-6.3}: The individual must request that the CE amend their PHI in writing.

In Figure 4, the right $R_{6.3}$ (Lines 1–15) and the implied obligation $O_{R-6.3}$ (nested in Lines 4–14) is extracted as a separate obligation (Lines 16–25). In general, the transformation uses a unique query to recognize the actions *permit* and *require* as delegation verbs in which the object of the delegation is always the implied right or obligation, respectively. Consequently, for the action *require* (Line 3) the activity (Lines 4–14) is identified as an implied obligation (Line 16–25).

```

1  activity [ right ] {
2    subject = CE
3    action = require
4    object = activity {
5      subject = individual
6      action = request
7      object = activity {
8        subject = CE
9        action = amend
10       object = PHI : individual
11     }
12    instrument = writing
13    target = CE
14  }
15 }

16 activity [ obligation ] {
17   subject = individual
18   action = request
19   object = activity {
20     subject = CE
21     action = amend
22     object = PHI : individual
23   }
24   instrument = writing
25 }
```

Figure 4: Right Balanced with an Obligation

Direct and indirect provisions are also balanced using transformations that rely on a unique query to identify and automatically resolve these cases [7].

In the fourth case, *anti-obligations* describe actions that stakeholders are not required to perform. In this case, the stakeholder’s implied right is to choose whether or not to perform that action. Anti-obligation models are expressed using a negated *obligation* symbol and balanced by replacing the negated *obligation* symbol with a *right* symbol. These symbols appear in square brackets after the *activity* symbol at the head of each semantic model for anti-obligations.

We define the balancing function $B : M \rightarrow M^*$ that maps a model M to a set of models M^* that is equal to:

$\{M\}$ if M is not unbalanced for cases 1–4; $\{M\} \cup \{M_1 \dots M_n\}$ if M is unbalanced for cases 1–3 and the models $M_1 \dots M_n$ are accumulated from applying the transformations for each case; or $\{M'\}$ if M is unbalanced for case 4 in which the model M' (a right) replaces the model M (an anti-obligation). In this step, we apply the function B to the set of input models to balance rights and obligations as necessary.

Balancing rights and obligations requires special handling for constraints in the rule pair (M, E) for each case 1–4. For delegations and indirect provisions, the balanced right or obligation will produce a new rule pair $(M', (e))$ where the new right or obligation M' is conditioned on the invocation of the original delegation M that occurs in the event e . For direct provisions and anti-obligations, the new rule pair (M', E) uses the same constraint expression E , because these cases model the same rule but from a different stakeholder perspective [7].

In the second phase, we see how events generated from implied rights and obligations correspond to the pre-conditions of other rights and obligations. Balancing rights and obligations ensures these dependent events in pre-conditions are accounted for.

3.3. Phase 2: State Machine Generation

In the algorithm’s second phase, two tables are generated: (1) the state-event table is generated by querying the semantic models from the first phase; and (2) the transition table is generated by iterating constraints and entries in the state-event table. Both steps are discussed in detail below.

3.3.1. Generate States and Events. In the first step, we populate the state-event table by querying semantic models. Entries in the state-event table have four fields, including a unique index for the state or event and a SAO-triple with subject, action and object.

To populate the table from rights, obligations and their constraints, recall from Section 3 we introduced the function $T(M)$ to generate the SAO-triple from a semantic model M . In this step, we define the recurrence relation $T(o)$ for the object $o \in T(M)$, whenever the object o is a type of activity. We ensured $T(o)$ is well-defined by disambiguating activities in Section 3.2.1. In addition, we introduce a similar function $T'(M)$ to extract one of two possible SAO-triples conveying the regulation’s authority over the stakeholder: $T'(M) = \langle Rule, permit, T(M) \rangle$ for a model M of a right; and $T'(M) = \langle Rule, require, T(M) \rangle$ for a model M of an obligation. In both cases, the subject of the triple is the regulation, identified by the *Rule*, whose authority is described by the action, either *permit* or *require*. Consider the example obligation model $O_{6.3}$ in Figure 5, below.

```

1  activity [ obligation ] {
2    subject = CE
3    action = provide
4    object = denial [ written ] {
5      subject = CE
6      action = deny
7      object = request {
8        subject = Individual
9        action = request
10     object = amendment {
11       subject = CE
12       action = amend
13       object = PHI
14     }
15   }
16 }
17 target = Individual
18 }

```

Figure 5: Example Recurrence for SAO-triple

The model is an obligation (Line 1) that requires “the CE provide the individual with a written denial to their request for amendment to PHI.” Consequently, the function $T(M) = \langle Rule, require, T(M) \rangle$ and $T(M) = \langle CE, provide, T(denial) \rangle$. Note how the object is an activity (*denial* on Line 4) thus leading to the subsequent recurrence $T(denial)$ in $T(M)$. For now we ignore attributes other than those involved in the SAO-triple such as the *target* in Line 13. Applying functions $T(M)$ and $T(M)$ yields the entries in Table 1.

Table 1: Example State-Event Table

Index	Subject	Action	Object
$O_{6.3}$	Rule	require	E_1
E_1	CE	provide	E_2
E_2	CE	deny	E_3
E_3	Individual	request	E_4
E_4	CE	amend	PHI

In the state-event tables, states are entries where the subject is the *Rule* and all other entries are events. Note a state is either a right or an obligation depending on the value in the action field, either *permit* or *require*, respectively. Successive uses of the same subject, action and object fields will reuse the first index to that triple. Table entries for constraints are produced using only the function $T(M)$ and recurrence when applicable. For example, the right $O_{6.3}$ has the constraint “the CE denies an individual’s request to amend PHI” in which the function T applied to the model yields an event equivalent to event E_2 .

3.3.2. Generate Transitions. In the second step, we populate the transition table by generating transitions using events from the first step in phase two. The transition table has four fields: the set number shared by constraints in a conjunction; the source state from which the transition leads out; the event used to

generate the transition; and the target state to which the transition leads in.

Each right and obligation state has the following transitions: in-transitions generated for events that were derived from pre-conditions; out-transitions generated from the event in the object field for states in the state table; and, if the state is an obligation, transitions for the negation of the event in the object field of the state table. For obligations, the negation of the event always leads to a non-compliant state (NC) for violating the obligation. For rights, the target state of this transition is unspecified. Continuing with the example from Section 3.3.1, we generate the transitions for obligation $O_{6.3}$ in Table 2.

Table 2: Example Transition Table

Set	Source	Event	Target
1		E_2	$O_{6.3}$
2	$O_{6.3}$	E_1	
3	$O_{6.3}$	$\neg E_1$	$NC_{6.3}$

The in-transition (first row) to $O_{6.3}$ corresponds with the constraint on $O_{6.3}$ and the out-transition (second row) corresponds with the object from the state $O_{6.3} = \langle Rule, require, E_1 \rangle$ in Table 1. The transition to the non-compliant state (third row) must eventually be conjoined with a time-out event or deadline to complete this monitor. The graphic illustration of this monitor appears in Figure 6.




Figure 6: Example Compliance Monitor

After the state-event and transition tables have been generated, one can derive a combined compliance monitor by pairing in- and out-transitions to connect states. The combined compliance monitor more effectively illustrates the interactions between rights and obligations. We present such a graphic in Section 4 as an application of the algorithm.

4. Results from HIPAA Privacy Rule

In a previous case study [7], we derived semantic models from rights, obligations and constraints that were extracted from the Privacy Rule [15] – a U.S. federal regulation for the HIPAA [14]. The Rule governs use and disclosure of patient healthcare information. Based on our discussions with CSOs, CISOs and CPOs, companies prioritize compliance with those regulations most likely to interface with the

public and consumers. For this reason, we applied our algorithm to §164.520 – §164.526 in Subpart E of the Rule. We present results from §164.524 titled “Access of individuals to protected health information.”

The analysis of §164.524 in the case study yielded a total of 20 rights, 26 obligations and 67 constraints. From these, the following rights and obligations are most relevant to generate the largest combined compliance monitor using the algorithm. The following acronyms are used: covered entity (CE), licensed healthcare professional (LHP), and protected health information (PHI).

- R_{4.1}**: The individual has a right to request access to their PHI.
- R_{4.3}**: The CE may deny an individual access to their PHI. (**C₁**)
- R_{4.5}**: The individual may have a denial of requested access reviewed by an LHP. (**C₂**)
- O_{4.1}**: The CE must permit an individual access to their PHI. (**C₃**)
- O_{4.2}**: The CE must deny an individual access to their PHI. (**C₄**)
- O_{4.3}**: The CE must permit an individual to request access to their PHI.
- O_{4.5}**: The CE must inform the individual that requested access is permitted. (**C₅**)
- O_{4.7}**: The CE must inform the individual that the requested access was denied. (**C₂**)
- O_{4.16}**: The CE must designate an LHP to review a denial of requested access. (**C₆**)
- O_{4.18}**: The LHP must recommend that the CE permit or deny the individual access to PHI. (**C₇ ∧ C₈**)
- O_{4.19}**: The CE must inform the individual of the recommendation of the LHP. (**C₃ ∨ C₄**)

Each right and obligation above is annotated with the logical expression of constraints (in parenthesis) from the following list:

- C₁**: The individual requests access to their PHI.
- C₂**: The CE denies requested access to PHI.
- C₃**: The LHP recommends the CE permit access.
- C₄**: The LHP recommends the CE deny access.
- C₅**: The CE permits the requested access to PHI.
- C₆**: The individual requires an LHP review a denial.
- C₇**: The CE designates the LHP to review a denial.
- C₈**: The LHP reviews the denial of access.

For the purpose of this illustration, we highlight only those events that form transitions between states and we ignore constraints that only describe state-of-being as they contribute no such events. As a result, the following state-event and transition tables are simply incomplete under the law but sufficient as an exemplar.

In phase one (consistency checking), step one, the algorithm identified several ambiguities in the original semantic models for rights, obligations and constraints. To resolve these ambiguities, the user was required to specify 37 subjects, 35 actions and 32 objects to disambiguate activities in the semantic models.

In step two, the algorithm balanced two rights and one obligation. The rights $R_{4.1}$ and $R_{4.5}$ and the obligation $O_{4.3}$ were balanced with new models $O_{R-4.1}$, $O_{R-4.5}$ and $O_{O-4.3}$, respectively. Since $O_{R-4.1} \approx O_{4.3}$ and $R_{4.1} \approx R_{O-4.3}$, the right $R_{4.1}$ balanced directly with $O_{4.3}$, which means $O_{R-4.1}$ and $R_{O-4.3}$ were not new contributions. However, the right $R_{4.5}$ only balances with obligation $O_{R-4.5}$ requiring the LHP to review denials of requested access, so $O_{R-4.5}$ is a new contribution.

In phase two (FSM generation), step one, the states and events were generated (see Table 3). There were 17 instances where events previously entered into the table were reused.

Table 3: State-Event Table for HIPAA §164.524

<i>Index</i>	<i>Subject</i>	<i>Action</i>	<i>Object</i>
R _{4.1}	Rule	permit	E ₁
E ₁	Individual	request	E ₂
E ₂	Individual	access	PHI
R _{4.3}	Rule	permit	E ₃
E ₃	CE	deny	E ₂
R _{4.5}	Rule	permit	E ₄
E ₄	Individual	require	E ₅
E ₅	LHP	review	E ₃
O _{R-4.5}	Rule	require	E ₅
O _{4.1}	Rule	require	E ₆
E ₆	CE	permit	E ₂
E ₇	LHP	recommend	E ₆
O _{4.2}	Rule	require	E ₃
E ₈	LHP	recommend	E ₃
O _{4.3}	Rule	require	E ₉
E ₉	CE	permit	E ₁
O _{4.5}	Rule	require	E ₁₀
E ₁₀	CE	inform	E ₆
O _{4.7}	Rule	require	E ₁₁
E ₁₁	CE	inform	E ₃
O _{4.16}	Rule	require	E ₁₂
E ₁₂	CE	designate	LHP
O _{4.18}	Rule	require	E ₇
O _{4.18}	Rule	require	E ₈
O _{4.19}	Rule	require	E ₁₃
E ₁₃	CE	inform	E ₇
O _{4.19}	Rule	require	E ₁₄
E ₁₄	CE	inform	E ₈

The generated transitions from step two appear in Table 4. The in-transitions were generated from constraints (Sets 1–11). The out-transitions were generated from the object value of states (Sets 12–25), the alternate transitions from rights (Sets 26–30) and the transitions from states to non-compliant states (Sets 31–40) all from Table 3.

We illustrate the combined compliance monitor in Figure 7. The transitions from obligations to non-compliant states and the alternate transitions for rights that do not align with existing states have been

omitted. Unspecified states appear as clouds with events E_{10} , E_{11} , E_{13} and E_{14} leading to such states.

Table 4: Transition Table for HIPAA §164.524

Set	Source	Event	Target
1		E_1	$R_{4.3}$
2		E_3	$R_{4.5}$
3		E_4	$O_{R-4.5}$
4		E_3	$O_{4.7}$
5		E_3	$O_{4.16}$
6		E_7	$O_{4.1}$
7		E_7	$O_{4.19}$
8		E_8	$O_{4.2}$
9		E_8	$O_{4.19}$
10		E_6	$O_{4.5}$
11		E_{12}	$O_{4.18}$
11		E_5	$O_{4.18}$
12	$R_{4.1}$	E_1	
13	$R_{4.3}$	E_3	
14	$R_{4.5}$	E_4	
15	$O_{R-4.5}$	E_5	
16	$O_{4.1}$	E_6	
17	$O_{4.2}$	E_3	
18	$O_{4.3}$	E_9	
19	$O_{4.5}$	E_{10}	
20	$O_{4.7}$	E_{11}	
21	$O_{4.16}$	E_{12}	
22	$O_{4.18}$	E_7	
23	$O_{4.18}$	E_8	
24	$O_{4.19}$	E_{13}	
25	$O_{4.19}$	E_{14}	
26	$R_{4.1}$	$\neg E_1$	
27	$R_{4.3}$	$\neg E_6$	
28	$R_{4.3}$	$\neg E_3$	
29	$R_{4.5}$	$\neg E_4$	
30	$O_{R-4.5}$	$\neg E_5$	
31	$O_{4.1}$	$\neg E_3$	$NC_{4.1}$
32	$O_{4.2}$	$\neg E_6$	$NC_{4.2}$
33	$O_{4.3}$	$\neg E_9$	$NC_{4.3}$
34	$O_{4.5}$	$\neg E_{10}$	$NC_{4.5}$
35	$O_{4.7}$	$\neg E_{11}$	$NC_{4.7}$
36	$O_{4.16}$	$\neg E_{12}$	$NC_{4.16}$
37	$O_{4.18}$	$\neg E_7$	$NC_{4.18}$
38	$O_{4.18}$	$\neg E_8$	$NC_{4.18}$
39	$O_{4.19}$	$\neg E_{13}$	$NC_{4.19}$
40	$O_{4.19}$	$\neg E_{14}$	$NC_{4.19}$

Figure 7 makes it easier to recognize important aspects of the combined compliance monitor. For example, the unconditional rights and obligations such as $O_{4.3}$ have no in-transitions. Rights or obligations that immediately follow unconditional obligations, like right $R_{4.1}$, are consequently unconditional, unless the stakeholder violates the preceding obligation. In addition, loops on states for obligations require stakeholders to maintain that state. For example, based on obligation $O_{4.18}$, if the LHP determines the CE

should not provide access (via E_8 , bottom center), then the CE must deny access (via E_3 and the loop at $O_{4.2}$).

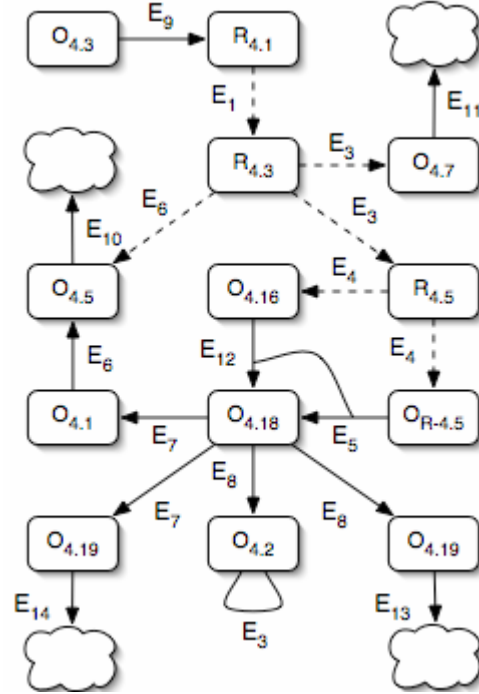


Figure 7: Combined Compliance Monitor

Rights that provide stakeholders choices are also easier to visualize. For example, in state $R_{4.3}$, the CE has a choice: (1) they can deny the requested access via E_3 , in which case they must (a) inform the individual via $O_{4.7}$ and (b) provide the right to review via $R_{4.5}$; or (2) they can permit the requested access, in which case they must inform the individual via $O_{4.5}$.

Rights and obligations assigned through delegation are clearly shown. For example, the obligation requiring the CE to permit the individual to request access (the path $O_{4.3} \rightarrow E_9 \rightarrow R_{4.1} \rightarrow E_1$) is distinct from the right permitting the individual to require an LHP to review a denial (the path $R_{4.5} \rightarrow E_4 \rightarrow O_{R-4.5} \rightarrow E_5$). However, the graphic does not clearly distinguish between rights to obligate other stakeholders and obligations that follow from invoking a stakeholder right, such as the path $R_{4.3} \rightarrow E_6 \rightarrow O_{4.5} \rightarrow E_{10}$.

5. Discussion and Future Work

Before discussing why the compliance algorithm works and therefore which compliance scenarios are likely to benefit most from this algorithm, we quickly address the current limitations.

In phase two, step one in Section 3.3.1, the algorithm must determine if two events are equivalent in order to reuse shared events. In this study, we only used the SAO-triple to compare events and identify duplicates; however, two different events can have the

same SAO-triple. For example, two similar requests to two different recipients or a repeated event with disjoint temporal constraints could both have the same *subject*, *action*, and *object*. We propose addressing this problem by comparing the semantic model sub-components used to generate the events instead of the SAO-triple, since they will have the necessary information to distinguish these events – including any temporal constraints.

In phase two, step two in Section 3.3.2, the set number is sufficient to assign transitions to logical conjunctions or disjunctions. However, we encountered the need to support exclusive-or on the out-transitions for obligations (see $O_{4.18}$ in Section 4) as a convenience to stakeholders. The consequences of exclusive-or on out-transitions impacts how transitions to non-compliant states are generated. For example, for the events A and B , the out-transitions for an obligation in the expression $(A \wedge \neg B) \vee (\neg A \wedge B)$ only require a transition to a non-compliant state for $(\neg A \wedge \neg B)$ and not for $\neg A$ independent of $\neg B$.

5.1. Why it works?

Why the algorithm consistently generates FSMs from semantic models of rights and obligations is not at first obvious. We can answer this question in more detail by examining the underlying principal in the SAO-triple and appealing to the structure of the recurrence tree. Recall from Section 3.3.2 the function $T(M)$ for a semantic model M and the recurrence $T(o)$ for $o \in T(M)$. Figure 8 illustrates the recurrence tree generated from events selected from the results in Section 4. The recurrence tree is rooted at the base triple $\langle \text{Individual}, \text{access}, \text{PHI} \rangle$. Arrows point from parent to child triples; children have the parent's index in the recurrence of the object field.

The recurrence only occurs in an SAO-triple when the action is a transitive-verb whose object is another activity; this is a unique case in semantic models [5]. In environments where stakeholders react to the actions of other stakeholders, these recurrences capture the object of that reaction: another activity. In the recurrence tree, stakeholder reactions map to triples where the stakeholder is the subject of that triple and the triple appears as a child node of the event to which they are reacting.

Regulations in these environments define which of these (re)actions are required or permitted. In the recurrence tree, regulations map to triples that appear as children of the actions they govern – these children are also leaves in the tree. The algorithm literally unravels stakeholder reactions and governing regulations to generate the pool of events that concern stakeholders. The algorithm binds regulations to states, recognizing that the event of performing the governed

action is an event that maps to an out-transition. In the recurrence tree, these events appear as the parents of regulation triples. Furthermore, regulated actions are pre-conditioned on other events that map to in-transitions on a state. When the regulation governs a reaction to an event produced by another regulation, these regulations become aligned in a compound compliance monitor. Consequently, the algorithm is most effective in highly regulated environments with complex stakeholder interactions – environments where tools such as this are needed most because the complexity of stakeholder interactions is greater and the cost of non-compliance may be severe.

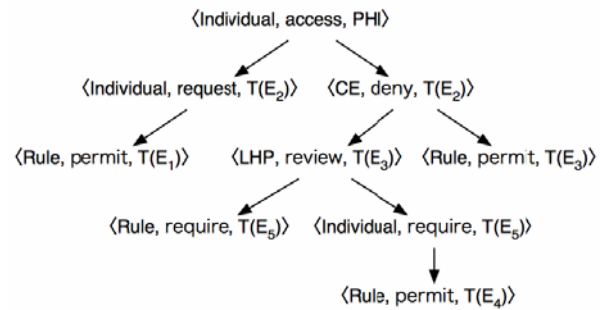


Figure 8: Recurrence Tree for SAO-triple.

5.2. Risk and Compliance

Future work includes integrating the compliance monitors produced by the algorithm into runtime systems. In addition, there is the need to evaluate risk and compliance associated with the decision to implement a right or obligation in systems. *Risk* is the probability that a regulation will be violated by a business or system process. For each right or obligation, calculating risk requires knowing the frequencies of events that pre-condition, satisfy or violate the right or obligation. Furthermore, risk must also factor in the penalty or cost of violating a right or obligation. For rights or obligations with a high penalty, frequent violation or with frequent satisfaction of pre-conditions, there is a higher priority to implement system processes to prevent violation and monitor compliance at runtime.

For example, in Figure 7 in Section 4, the individual is given the right to request access to PHI via right $R_{4.1}$ and, if denied that access, they receive the right to have an LHP review the denial via right $R_{4.5}$. The review involves a complex set of stakeholder interactions between the obligations of the LHP and CE and the rights of the individual. If individuals rarely exercise right $R_{4.5}$, the cost of implementing these obligations in systems may be incommensurate with the frequency of violations due to human error in the business process. Since the compliance monitors can be aligned with

both business and system processes, stakeholders can use these monitors to design and develop software systems commensurate with risk and compliance costs.

6. Conclusion

We present an algorithm for generating finite state machines (FSMs) from stakeholder rights and obligations that are extracted from policies and regulations. The algorithm has been applied and validated to date within the context of U.S. health regulation (HIPAA). The FSMs can be used to evaluate design decisions in terms of risk and monitor compliance in runtime systems. The user only provides semantic models to the algorithm, after which the algorithm proceeds in two phases to: (1) perform consistency checking to identify ambiguities and balance rights with obligations; and (2) generate the FSM in state-event and transition tables.

Acknowledgements

We thank the members of ThePrivacyPlace.org for their helpful comments. This work was funded by NSF ITR Grant #032-5269.

References

- [1] J. Atlee, J. Gannon, "State-based Model Checking of Event-driven System Requirements." *Conf. Soft. for Critical Systems*, New Orleans, LA, pp. 16-28, 1991.
- [2] F. Baader, J.H. Siekmann, Unification Theory. *Handbook of Logic in AI and Logic Programming*, Oxford University Press, New York, NY, pp. 41-125, 1994.
- [3] R. Bharadwaj, C.L. Heitmeyer, "Model Checking Complete Requirements Specifications Using Abstraction." *Auto. Soft. Engr.*, 6(1), pp. 37-68, 1999.
- [4] T.D. Breaux, A.I. Antón, "Deriving Semantic Models from Privacy Policies." *IEEE Workshop on Policies for Distributed Sys. & Networks*, Sweden, pp. 67-76, 2005.
- [5] T.D. Breaux, A.I. Antón, "Analyzing Goal Semantics for Rights, Permissions, and Obligations." *IEEE Req'ts. Engr. Conf.*, Paris, France, pp. 177-186, 2005.
- [6] T.D. Breaux, A.I. Antón, "Mining Rule Semantics to Understand Legislative Compliance." *ACM Workshop on Privacy in Electronic Society*, USA, pp. 51-54, 2005.
- [7] T.D. Breaux, M.W. Vail, A.I. Antón, "Towards Regulatory Compliance: Extracting Rights and Obligations to Align Requirements with Regulations" In submission *IEEE Int'l Conf. Reqs. Engr, NCSU CSC Tech. Report TR-2006-06*, 2006.
- [8] M. Chechik, J. Gannon, "Automatic Analysis of Consistency Between Requirements and Design." *IEEE Trans. Soft. Eng.*, 27(7), pp. 651-672, 2001.
- [9] D. Cohen, M.S. Feather, K. Narayanaswamy, S.F. Fickas, "Automatic Monitoring of Software Requirements." *IEEE Int'l Conf. Soft. Eng.*, pp. 602-603, 1997.
- [10] Ernst & Young, Global Information Security Survey 2005: Report on the Widening Gap, 2005.
- [11] S.F. Fickas, T. Beauchamp, N.A.R. Mamy, "Monitoring Requirements: A Case Study." *IEEE Int'l Conf. Auto. Soft. Eng.* Edinburgh, UK, pp. 299-304, 2002.
- [12] P. Giorgini, F. Massacci, J. Mylopoulos, N. Zannone. "Modeling Security Requirements Through Ownership, Permission and Delegation." *IEEE 13th Req'ts. Eng.. Conf.*, France, pp. 167-176, 2005.
- [13] C.L. Heitmeyer, R.D. Jeffords, B.G. Labaw. "Automated Consistency Checking of Requirements Specifications," *ACM Trans. Soft. Eng. Methods*, 5(3), pp. 231-261, 1996.
- [14] Health Insurance Portability and Accountability Act, USC H.R. 3103-168, April 2000.
- [15] "Standards for Privacy of Individually Identifiable Health Information." 45 CFR Part 160, Part 164 Subpart E. In Federal Register, vol. 68, no. 34, February 20, 2003, pp. 8334 – 8381
- [16] S. Kerrigan, K.H. Law, "Logic-based Regulation Compliance-Assistance." *Int'l Conf. AI and Law*, pp. 126-135, 2003.
- [17] R. de Landtsheer, E. Letier, A. van Lamsweerde, "Deriving Tabular Event-based Specifications from Goal-Oriented Requirements Models." *IEEE Req'ts. Eng. Conf.*, Monterrey, CA, pp. 200-210, 2003.
- [18] N.A.M. Maiden, "CREWS-SAVRE: Scenarios for Acquiring and Validating Requirements," *Auto. Soft. Eng.* 5(4), pp. 419-446, 1998.
- [19] Medical Privacy - National Standards to Protect the Privacy of Personal Health Information. Office for Civil Rights, US Department of Health and Human Services. 2000. <http://www.hhs.gov/ocr/hipaa/finalreg.html>.
- [20] D.K. Peters, D.L. Parnas, "Requirements-based Monitors for Real-time Systems." *IEEE Trans. Soft. Eng.*, 28(2), pp. 146-158, 2002.
- [21] W.N. Robinson, "A Requirements Monitoring Framework for Enterprise Systems." *Req'ts. Eng. Journal*, 11(1), pp. 17-41, 2005.
- [22] M.J. Sergot, F. Sadri, R.A. Kowalski, F. Kriwaczek, P. Hammond, H.T. Cory, "The British Nationality Act as a Logic Program," *Comm. of the ACM*, 29(5), pp. 370-386, 1986.
- [23] M.J. Sergot, A.S. Kamble, K.K. Bajaj, "Indian Central Civil Service Pension Rules: A Case Study in Logic Programming." *Int'l Conf. AI & Law*, pp. 118-127, 1991.
- [24] D. Sherman, "A Prolog Model of the Income Tax Act of Canada." *Int'l Conf. AI & Law*, pp. 127-136, 1987.
- [25] G. Spanoudakis, K. Mahbub, "Requirements Monitoring for Service-based Systems: Towards a Framework Based on Event Calculus." *IEEE Int'l Conf. Auto. Soft. Eng.*, Linz, Austria, pp. 379-384, 2004.
- [26] A.G. Sutcliffe, N.A.M. Maiden, S. Minocha, D. Manuel, "Supporting Scenario-based Requirements Engineering," *IEEE Trans. Soft. Eng.*, 24(12), pp. 1072-1088, 1998.