

# Towards a Framework for Law-Compliant Software Requirements

Alberto Siena  
FBK - Irst  
via Sommarive 18 - Trento, Italy  
siena@fbk.eu

Anna Perini  
FBK - Irst  
via Sommarive 18 - Trento, Italy  
perini@fbk.eu

John Mylopoulos  
University of Trento  
via Sommarive 14 - Trento, Italy  
jm@cs.toronto.edu

Angelo Susi  
FBK - Irst  
via Sommarive 18 - Trento, Italy  
susi@fbk.eu

## Abstract

*During the requirements elicitation phase, analysts have often to take into consideration laws and regulations enacted by different levels of government. The purpose of this paper is twofold. First, a systematic process is outlined which, given a problem and a collection of legal prescriptions, generates a set of requirements that address the problem while complying with the prescriptions. Second, the conceptual framework is outlined, which characterises the process by providing both legal concepts proposed in theoretical studies in the legal domain and concepts from goal-oriented requirements engineering. The issues and challenges of the proposed framework are also evaluated, with regard to expected results.*

## 1. Software and the law

Nowadays, as information systems become more pervasive, laws are having an increasing impact both on their functionalities and on the people that use them. Laws and regulations enacted by different levels of government are continuously addressed by analysts and developers during the software development activities. However, its systematic processing is essentially excluded from the mainstream software development processes and tools, and left to the ingenuity of individual designers. It has been estimated that in the Healthcare domain, organisations have spent \$17.6 billion over a number of years to align their systems and procedures with a single law, the Health Insurance Portability and Accountability Act (HIPAA), introduced in 1996<sup>1</sup>.

<sup>1</sup>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/nalreg.html>

In the Business domain, it was estimated that organisations would spend \$5.8 billion in one year alone (2005) to ensure compliance of their reporting and risk management procedures with the Sarbanes-Oxley Act (SOX, for short)<sup>2</sup>.

In recent years, some research has been undertaken to investigate this aspect of the software development. Antón and Breux [1] have developed a systematic process for extracting rights and obligations (and auxiliary concepts such as actors and constraints) from legal text thereby generating a formal model of a law. This makes an important step forward in dealing with the complexity and syntactic ambiguity of legal sources. About legal concepts, the LRI-Core [2] is a layered ontology of law, rooted in a foundational ontology that can be instantiated into domain ontologies. It is founded on the thesis that law is driven by common world concepts and words, and as such the ontology contains concepts such as agent, action, organisation, and so on, together with legal concepts. Somehow, this idea about laws is implicitly contained in some works that attain requirements modelling. Darimont and Lemoine use KAOS as a modelling language for representing objectives extracted from regulation texts [3]. Such an approach is based on the analogy between regulation documents and requirements documents. Partially similar are the techniques adopted by Ghanavati et al. [4], who use GRL to model goals and actions prescribed by laws. This work develops on the intuition of using the same modelling framework for both the regulations and the organisation, and this allows to establish traceability links between the law and the requirements. The Normative *i\** framework [8] allows for modelling laws inside an intentional framework and produces effective additions to the requirements system.

With respect to these approaches, we stress the impor-

<sup>2</sup>Online News published in DMReview.com, November 15, 2004

tance of a systematic process, to achieve the result of ensuring compliance by construction.

## 2. Impact on requirements specification

The first step of the software development process is to understand and specify the requirements of the system-to-be. Goal-oriented Requirements Engineering (GORE) frameworks rely on the concept of goal to represent the needs and objectives of the stakeholders, and derive the system requirements from the goal model of the domain. In contrast, laws are not expression of the needs of the stakeholders, and often prescribe behaviours contrary to the wish of the stakeholders, thus destructively interacting with the role of goals. Searching, identifying, tracing and solving conflict situations between laws and goals become then necessary steps to reconcile goal-based requirements elicitation processes with the presence of laws. The resulting requirements model is at the same time a subset of the acceptable alternatives, with respect to what the law permits (Fig. 1).

To prevent this situation to happen, we propose to consolidate the architecture of the requirements specification process, such that compliance is ensured by construction. In order to do this, we rely on the intuition that laws modify the space of alternatives by introducing invariants. In each alternative the same legal invariants produce different effects that are part of the cost/benefit analysis of the alternative. For example, what if a law prescribes the liability of a certain class of subjects because of privacy issues? In a centralised organisational structure, the central office has to deal with the privacy law, and the peripheral offices act as clients. Viceversa, if the organisational structure is distributed, every office has to face the privacy law liability. In any case, liability for data protection is invariant in the space of alternatives.

Laws are complex artefacts that attempt to prescribe how the world should be; i.e., it describes the properties of the desired world, out of all the possible alternative worlds. As such, the law can be seen as a collection of constraints on the domain under analysis. More formally, given a domain  $D$ , a law  $L$  is a subset of  $D$ , such that  $\forall P' \subset L$  then  $comp(P') = true$ , where  $P'$  is a behaviour and  $comp(x)$  is a function that returns *true* if  $x$  is compliant, and *false* if  $x$  not compliant. A strategy  $S$  is as well a subset of  $D$ , such that  $\forall P \subset S$  then  $goal(P)$ , where  $P$  is a behaviour and  $goal(x)$  is a function that returns *true* if  $x$  is an admissible strategic solution, and *false* if it is not admissible.

Three cases exist:

$$S \cap L = \emptyset \quad (1)$$

In this case, the goals of the organisation completely fall outside the boundaries of the law; this means that in no way it is possible to comply with the law. This is a suspicious case, as shown below.




Fig. 1: Any behaviour that is at the same time wanted by the stakeholders and prescribed by the law is a strategic, law-compliant behaviour

$$S \cap L = C \mid C \subset S \quad (2)$$

In this case, the possible, alternative strategic choices have potentially dangerous differences, because moving from one to another could lead the organisation from being compliant to being non-compliant.

$$S \cap L = C \mid C = S \quad (3)$$

In this case, no matter what alternative the organisation will choose, it will always be law compliant. This is the optimal solution.

## 3. A legal-strategic framework

We have depicted the possible cases of compliance through the use of sets. However, if the boundaries of  $L$  are defined by legal prescriptions, the boundaries of  $S$  are defined by stakeholders' objectives. So intuitively, these two sets are hardly comparable because they have (a) different languages, (b) different concepts, (c) different interests, and (d) different scope. But the language of laws is strongly influenced by the matter the law is regulating [5]. Broadly speaking, the matter regulated by laws are human societies, comprised by interacting actors with different goals and behaviours, so the legal concepts should also be able to describe such kind of elements.

In requirements engineering, goal-orientation is a paradigm that use models of actors and their goals to describe the organisation. For example,  $i^*$  [10] (but other frameworks have similar characteristics) is a goal-oriented requirements engineering framework able to capture the *why* of the system-to-be; it provides a modelling framework tailored to model the domain as composed of heterogeneous actors with different goals. Actors depend on each other to undertake their tasks and achieve these goals.  $i^*$  addresses two aspects of the domain: the **strategic dependencies** among actors - i.e., the system-wide strategic model based on the matching between the depender, which is the actor who "wants" something and the dependee, who has the "ability" to do something; and the **strategic rationale** of the actors - i.e., a description of how each actor pursues its objectives, expressed in terms of intentional elements

such as goals, tasks, resources and softgoals, linked by task decomposition links, means-end links, and the contribution links.

A conceptual language to capture legal prescriptions is the fundamental hohfeldian legal taxonomy - a taxonomy grounded on 8 elementary concepts classified by Hohfeld [6] as privilege, claim, power, immunity, and their correlatives no-claim, duty, liability, disability. **Privilege** is the entitlement for a person to discretionally perform an action, regardless of the will of others who may not claim him to perform that action; for example, giving a tip at the restaurant is a liberty, and the waiter can't claim it. **Claim** is the entitlement for a person to have something done, and to legally pretend it; for example, if John has the right to exclusively use of his land, others have a corresponding duty of non-interference. **Power** is the (legal) capability to produce changes in the legal system; examples of legal powers include the power to contract and the power to marry. **Immunity** is the right of being kept untouched from other performing an action; for example, one may be immune from prosecution as a result of signing a contract. Two rights are **correlatives** [6] if the right of a person A implies that there exists another person B (A's counter-party), who has the correlative right. For example, **duty** and claim are correlatives, because if someone has a claim - let say, to access some data - then somebody else will have the duty of providing that data; similarly privilege-noclaim, power-liability, immunity-disability are correlatives. The concept of *correlativeness* implies that rights have a **relational nature**. In fact, they involve two subjects: the owner of the right and the one, against whom the right is held - the *counterparty*. The objects of rights are "actions" [7]. Two types of actions exist: *behavioural* and *productive*. **Behavioural actions** are described by the actual behaviour performed by actors ("A does x"); **productive actions** attain the results produced by the behaviour of the actors ("A brings it about that x") [7].

The capability of law to produce effects stays in its *prescriptive* strength, which ultimately brings to deontics. Viceversa, the rights-based legal taxonomy presented above is essentially *descriptive* and successfully captures the *effects* produced by the legal prescription. As such, a right-based conceptual model acts as the join-point between the legal domain ( $L$ ) and the intentional domain ( $S$ ). Specifically, we map the hohfeldian taxonomy into the  $i^*$  goal-oriented meta-model, describing its intentional properties. As an example, focusing on the actors of the domain and their behavior, we can notice that:

- The *subjects* addressed by the law  $L$  are stakeholders of the intentional domain  $S$ .
- The *actions* prescribed by the law  $L$  turn into goals and tasks (in the intentional domain  $S$ ) that can be respec-

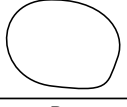
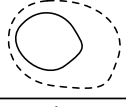

	Domain characterisation	Law modelling	Goal modelling
Activities undertaken	Preliminary domain exploration	Laws identification; rights identification and modelling	Model refinement; intentional modelling
Input		Knowledge of stakeholders and activities nature	Model of the legal superset
Output	Knowledge of stakeholders and activities nature	Model of the legal superset	Model of the legal-strategic set
			
	D	L	$S \cap L$

Table 1: The compliance-by-construction process. The purpose of the process is to refine the boundaries of the domain up to the identification of  $S \cap L$ .

tively achieved and performed by the stakeholders in  $S$ .

- The existence of rights articulate the way stakeholders can operationalise their goals or decompose their tasks.

#### 4. Results: a refinement process

How can this conceptual language be used to improve the elicitation process? We imagine to be in charge of eliciting requirements for a large organisation, which needs to set up a system that deal with sensitive data. We build the model of Fig. 1 by means of a refinement process, which aims at cutting the space of alternatives from the unfeasible (non-compliant) ones. The process is comprised by 3 steps, namely *Domain definition*, *Law modelling* and *Goal modelling*, and is summarised in Table 1.

**Domain characterisation.** What are the laws that have to be taken into consideration (if any)? Not every law has to be taken into account for a given domain. So the first step is to give a preliminary characterisation of the domain boundaries (as in Table 1) in terms of stakeholders and their basic behaviour. This acts as the invariant model to be used for further modelling and will allow to identify the relevant laws. For example, acquiring the knowledge that the main stakeholder is a hospital, will allow to select a medical data protection law rather than a trading regulation law. This activity allows to exclude from the domain those laws that are not applicable.

**Law modelling.** Hypothesising that the main actor is a US based hospital, we have to face the US Health Insurance Portability and Accountability Act (HIPAA). For example,

article §164.314 prescribes: *A CE is not in compliance [...] if the covered entity knew of a pattern of an activity [...] of the business associate that constituted a material breach [...] of the business associate's obligation under the contract [...], unless the CE took reasonable steps to cure the breach [...], and, if such steps were unsuccessful - terminated the contract or reported the problem to the Secretary.*

The fundamental legal taxonomy exposed in section 3 constitutes the conceptual language for law modelling. By means of that language, we can model the legal prescriptions and define the legal set  $L$ , and then refine it with intentional elements of  $S$ . As in [9], to create a model of the law 4 questions are relevant: who is the subject of the law? What is the object of the law? Which are the counter-parties? And what is the legal right being created? Answering to these questions produces a model that is the reference model for the *comp* function. For example, in this fragment we can identify 3 subjects: the Covered Entity (CE), the Business Associate (BA) and the Secretary. To each of them we can associate some goals or tasks of the legal model. The Secretary has to receive the report of non-compliant behaviour. The BA has to cure the breach. The CE has to restore the compliance state in one of the 3 possible ways: cure the breach - if the BA consents to do it; terminate the contract, if possible; or report the problem to the Secretary. In this fragment we can identify the *claim* of the CE over the BA and the correlative *duty* of the BA; the *liability* of the CE towards the Secretary; and the *immunity* of the CE towards the Secretary, under proper conditions. So ultimately, we have actors, goals or tasks, and the constraints or opportunities that are given to actors by the prescriptions.

What we have done now has been to define the boundaries of the set  $L$ . It means that the elements of the model are marked as invariant for the subsequent activity, and no requirements model will be acceptable if the structure of  $L$  will not be present or the goals and tasks in  $L$  will not be achievable.

**Goal modelling.** Typically, goal-oriented methodologies are guided by the top goals of the stakeholders. In activity of law modelling, we have defined a higher level set of requisites, which shall guide the decisions of the actors. Goal modelling becomes now a refinement step of the law pattern built in step 2 through the modelling of intentional elements. This way, the rules of modelling through refinement have two effects. First, they can implicitly exclude the case of Eq. (1), because the top goals of the stakeholders are incompatible with the law; for example, any decision to make patient's personal data available via Internet is unfeasible by law. Second, they can ensure to be in the subset  $C$  of Eq. (2) from the space of alternatives, because adding these alternatives to the model is structurally prevented. The result is a realisation of an organisational structure in line with the normative prescriptions - i.e., the case of Eq. (3).

## 5. Challenges

In this paper, we have presented the guiding rules and a conceptual framework for a systematic approach for deriving compliant-by-construction requirements. In the cases the law plays an important, distorting role on the effectiveness of the requirements elicitation phase, the direction indicated by the presented approach goes in the direction of reducing the complexity of the problem.

The success of the described process strongly depends on the representation capabilities of the adopted modelling language. Currently, the  $i^*$ -based language used here does not allow to express legal invariants, so that the subset of (3) can't be formally guaranteed. As a further consequence, the definitions of the sets  $D$ ,  $L$  and  $S$  are indeed weak and their boundaries fuzzy. This problem mainly concerns the nature of the law itself, which is ambiguous and subject to interpretation.

## References

- [1] T. D. Breaux, M. W. Vail, and A. I. Anton. Towards regulatory compliance: Extracting rights and obligations to align requirements with regulations. In *RE'06: Proceedings of the 14th IEEE International Requirements Engineering Conference (RE'06)*, pages 49–58, Washington, DC, USA, September 2006. IEEE Society Press.
- [2] J. Breuker, A. Valente, and R. Winkels. Legal ontologies in knowledge engineering and information management. *Artif. Intell. Law*, 12(4):241–277, 2004.
- [3] R. Darimont and M. Lemoine. Goal-oriented analysis of regulations. In *ReMo2V*, 2006.
- [4] S. Ghanavati, D. Amyot, and L. Peyton. Towards a framework for tracking legal compliance in healthcare. In J. Krogstie, A. L. Opdahl, and G. Sindre, editors, *CAiSE*, volume 4495 of *Lecture Notes in Computer Science*, pages 218–232. Springer, 2007.
- [5] R. Hoekstra and J. Breuker. Commonsense causal explanation in a legal domain. *Artificial Intelligence and Law*, 15(3):281–299, 2007.
- [6] W. N. Hohfeld. *Fundamental Legal Conceptions as Applied in Judicial Reasoning*. Yale Law Journal 23(1), 1913.
- [7] G. Sartor. Fundamental legal concepts: A formal and teleological characterisation. *Artificial Intelligence and Law*, 14(1-2):101–142, April 2006.
- [8] A. Siena, N. A. M. Maiden, J. Lockerbie, K. Karlsen, A. Perini, and A. Susi. Exploring the effectiveness of normative  $i^*$  modelling: Results from a case study on food chain traceability. In *20th International Conference on Advanced Information Systems Engineering (CAiSE'08)*, pages 182–196, 2008.
- [9] A. Siena, J. Mylopoulos, A. Perini, and A. Susi. From laws to requirements. In *1st International Workshop on Requirements Engineering and Law (Relaw'08)*, 2008.
- [10] E. S.-K. Yu. *Modelling strategic relationships for process reengineering*. PhD thesis, University of Toronto, Ontario, Canada, 1996.