From Laws to Requirements

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

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

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

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

Abstract

Legal prescriptions are increasingly impacting on information systems and on organisations that must comply with them in order to avoid to be prosecuted or fined. Addressing law compliance in early phases of the requirements analysis helps in improving the alignment of information systems with the law. In this paper, we point out ontological differences between legal concepts and requirements and set the basis for a systematic process able to support decision making about requirements for law compliant systems.

1. Introduction

Laws and regulations are having an increasing impact on legacy and future software systems that must comply or face penalties. 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[1]. 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)[2].

But what exactly does "compliance" mean for a software system, and how does a company achieve it through a systematic, tool-supported process? How does one derive requirements from a law? The purpose of this paper is to begin to address these questions. A key problem that underlies these questions is that the concepts in terms of which a law is expressed are fundamentally different from those in terms of which requirements are defined. Laws are about rights and obligations, privileges and liabilities [7]. Requirements, on the other hand, are about stakeholders and their goals. To derive requirements from law then, amounts to establishing a systematic process for transforming legal concepts into stakeholder goals so that if the goals are fulfilled through a particular system design, then the law is upheld.

We are not alone in this quest. Anton and Breaux [1] have developed a systematic process for extracting rights and obligations (and auxiliary concepts such as actors and contraints) from legal text thereby generating a formal model of a law. So, our starting point is precisely a formal model of a law, expressed in an extended metamodel that includes legal concepts beyond rights and obligations, but we focus on how those laws are transformed into requirements.

This is a position paper, intended to motivate and sketch a research agenda for going from laws to requirements in a systematic way. The rest of the paper is structured as follows: section 2 summarises the related works; section 3 presents our research baseline; section 4 describes the problems by means of examples; section 5 introduces our proposal for a solution; section 6 discusses future work and expected results; finally, section 7 concludes.

2. Related work

By systematically extracting rights and obligations from legal texts, the work of Breaux and Anton [1] makes an important step forward in dealing with the complexity and syntactic ambiguity of legal sources, thus constituting the starting point for our analysis. However, their set of concepts is

[1] 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

[2] Online News published in DMReview.com, November 15, 2004

tailored for the purpose of automatic data extraction from documents; so, for example, obligations and rights are defined as *statements*, and right are defined as something that stakeholders are *permitted* to perform. This is akin to deontic permission rather than to the notion of "legal right" (as will be discussed in section 3). In contrast, our work attempts to clarify the difference between requirements and legal concepts and bridge the gap between them through a systematic process.

About legal concepts, the LRI-Core [3] is a layered ontology of law, rooted in a foundational ontology that can be instantiated into domain ontologies. It is built on top of the consideration 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 [4]. Such an approach is based on the analogy between regulation documents and requirements documents. Partially similar are the techniques adopted by Ghanavati et al. [5], who use *i** 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. SecureTropos [6] is a security-enhanced version of the Tropos methodology, which introduces the concepts of services ownership and delegation. In order to ensure access control, strategic dependencies are refined with the conditions of permission and commitment.

With respect to these works, we look at the whole process of extracting requirements from law, and try to ground it on a the proper meta-model to systematically deriving law-compliant requirements.

3. Research baseline

We rely on the assumption that the *why* of the choices about an information system is successfully captured by the analysis of the goals of stakeholders. *i** [10] supports this assumption by providing 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.

The intentional paradigm for requirements modelling, comprised by actors, goals and strategic dependencies, is hardly applicable to the legal domain, which is made by prescriptions and deontic necessities. We adopt the fundamental legal taxonomy grounded on 8 elementary concepts classified by Hohfeld [7] 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 noninterference. **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** [7] if the right of a person A implies that there exists another person B (A's counterparty), 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" [8]. 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 that are produced by the behaviour of the actors ("A brings it about that x") [8].

The described legal ontology is substantially different from the *i** ontology. *i** is about *intentions*: actors *want* something (goals) and depend on each other to achieve their goals, and from these intentions the overall system configuration emerges. On the contrary, laws constrain stakeholders by means of legal prescriptions, thus affecting their behaviour and social relationships. The above proposed set of legal concepts provides a way to represent those constraints. Compared to requirements, legal concepts are at a higher level, in the sense that there are many combinations of actor intentions (and corresponding behaviours) that can fulfill a legal constraint. It is part of the requirements analysis activity to understand their effects in a given domain, and to ensure that the settings of the domain is compliant with the

law. The Normative *i** framework [9] allows for modelling laws inside an intentional framework and produces effective additions to the requirements system [9]. However, it still lacks of a systematic debate of more sophisticated legal concepts and their effects on alternative requirements, as discussed below.

4. Laws and requirements: an example

HIPAA regulation §164.314(a)(1)(ii) prescribes that: "*a Covered Entity (CE) is not in compliance [...] if the CE knew of a pattern of an activity or practice of the business associate that constituted a material breach or violation of the business associates obligation under the contract or other arrangement*". Without this law, every entity covered here had the possibility to interact with contractors, ignoring their behaviour. But the law raises a problem of compliance, and to solve this problem, it is important to know what is the law actually prescribing, what can be done to comply with the law in terms of possible alternatives, and which alternative is better for our goals. So ideally, we would need a model of the law that, once applied, allows for the analysis of the alternatives to select the best one. In order to do this, we try to identify a conceptual model of the law that can act as a guiding criterion for the requirements representation. This conceptual model is based on the fundamental concepts described in section 3, namely the *object of a right*, the *subjects of a right*, including the *counterparties*, and the *nature of a right*.

Object of the right. What is the action actually specified by the law? Sometimes, the law prescribes exactly the "behavioural" action to perform. Other times, it prescribes a "productive" action [8], a state of the world - a goal - that is supposed to be ensured by the CE to be compliant with the law. Article §164.314(a)(1)(ii) specifies a condition of un-compliance. We claim that this condition is not the right information to extract from the law. To avoid this condition from being true, a CE adopts the high level goal of ensuring that its contractors comply with article §164.314(a)(1)(ii). This top goal can be decomposed into different alternative sub-goals and implemented with different processes, but each process, if correctly engineered to achieve the top goal, will lead to actual compliance.

Subjects of the right. Laws have different scope than requirements and have different (and higher) level of abstraction. Laws typically identify the addressed subjects in terms of their common characteristics or behaviour, thus resembling strictly *agent types*[3]. On the other hand, requirements models describe the intentional settings of the domain; as such, they typically refer to a particular organisation and to a set of *roles* defined in terms of their organisational functions. This subtle difference has implication on the requirements modelling activity. Agents can be discretionally assigned to play roles; but if an agent or a role has the characteristics given by the law, it is necessarily the subject defined by the law. So for example, changing the goals assigned to a role can change a stakeholder from being a CE into not having to comply with HIPAA.

Counterparties. It is worth mentioning that, given the relational nature of rights, the same considerations hold for counterparties. Some kinds of legal prescriptions can be represented by means of deontic impositions to perform certain actions. We refer in this case to the concept of *erga omnes* right [8]. But the nature of legal rights often involves other subjects as the holder of certain counter-rights. In this case, the subject will have to interact with its counterparty

Nature of the right. With respect to the identified counterparties (or *erga omnes*), a CE is holder of legal rights. Rights of different nature cause alternative requirements to be generated, according to the stakeholders' goals. To give an example, let consider how a *duty* could call for a different solution than a *claim* or a *power*.

Duty. Duties are the most intuitive legal rights and many of the considerations explained above (see "Object of the right") can be repeated here. On the one hand, a duty constrains the behaviour of the addressee; on the other hand, it allows free choice on how to "realise" the duty. For example, article §164.314(b)(2)(ii) says that a CE *must [...] ensure that the adequate separation [...] is supported by reasonable and appropriate security measures;*

Claim. Article §164.314(a)(2)(i) says that *the contract between a covered entity and a business associate must provide that the business associate will [...] (C) Report to the covered entity any security incident of which it becomes aware*. If the business associate has intention to keep closed the information about security incidents, a CE has the legal claim of pretending this from the counterparty. It is up to the CE to decide about how to use this claim: for example, it could decide to assign the responsibility Ensure access to security incidents information in contracts to the commercial office, or to create a specific role such as Associates security officer in the organisation.

Power. Article §164.314(a)(1)(ii)(A) later confers to the CE the power to terminate the contract with the counterparty if it knows that the counterparty behaves in infringement of the law. The power that the CE has is not a requirement for the system-to-be, but it might generate requirements: a goal such as Contracts with business associates in violation of law be terminated could be assigned as a responsibility to an internal office of a company, generating this way a requirement about its implementation.

Summarising, rights of different nature generate different and alternative organisational configurations in terms of the responsibilities that are assigned to the stakeholder. Each alternative generates a possibly different set of re-

[3] See http://www.loa-cnr.it/mostro/files/MostroDel5.pdf

quirements for the system-to-be, and the choice of the best one is expected to be a result of the requirements analysis. In any case, if every stakeholder fulfills its goals as modelled upon the law, then compliance is ensured. These different allocations of responsibilities, not the law itself, are the actual requirements for the system-to-be.

5. Intentional compliance

Organisations have strategies to pursue their institutional objectives. Requirements modelling aims at capturing the needs that concern the information system, by understanding the overall organisational settings that support the strategies. Laws break those strategies, in that they make prescriptions that have (a) different language, (b) different concepts, (c) different interests, (d) different scope than the strategies. Organisations face law prescriptions by trying to adapt their strategies and comply without compromising their objectives.

The contribution of goal orientation to requirements engineering consists in the capability to understand the *why* - beside the *what* and *how* - of the system-to-be. Following this paradigm, we raise the question of what does it mean capturing the *why* of the choices about an information system when laws and regulations are involved.

As discussed in previous section, legal prescriptions can't be univocally transformed into requirements: legal prescriptions generate alternative possibilities to be compliant. This means that legal compliance is a matter of *decision making* that involves the goals of the stakeholders. An allocation of goals - a strategy - is in compliance with a law if some condition holds for stating that the strategy is inside the boundaries defined by the law. Supporting the decision making means providing the right level of *abstraction*, such that the condition for compliance can be evaluated.

The representation of the knowledge for human comprehension should be adequate to the representation of both the legal concepts, together with the alternatives they create, and the intentions of the stakeholders, together with their preferences. If this is possible, we derive the following property: we define ***intentional compliance*** *the design-time distribution of responsibilities, such that if every actor fulfills its goals, then the compliance is ensured.* Modelling the intentional compliance means modelling the effects of the law together with the strategy of the organisation. The result is a realisation of an organisational structure *intentionally* in line with the legal prescriptions.

Intentional compliance can play a crucial role in guiding the development of the system, and keeping it compliant through all the phases of the development, so that the running system will also result compliant. We define this condition of compliance for the running system and the actual processes it supports as ***actual compliance*** or *dynamic compliance*. If the system or the supported processes are not ensured to be actually compliant, a violation may occur and they need to be re-aligned. We define this condition of reinforced compliance ***strong compliance*** or *late compliance*. The later a compliance violation is detected, the more expensive is to re-align the system. Requirements models characterised by intentional compliance should allow for early detection of violations, avoiding the need of imposing strong compliance later, and thus reducing costs.

6. Discussion

To ensure intentional compliance we have to fill the gap between intentional and legal knowledge representation, and here is where the described taxonomy is useful. Law is not a world apart: law's language and concepts are strongly influenced by the subject that it is intended to regulate [2]. Accordingly, since the goal-orientation paradigm is tailored to describe the stakeholder needs, we may expect to find in the law concepts akin to goals. The Hohfeldian taxonomy considers high-level concepts (such as claim, power, immunity) as primitive. Due to this abstraction level of the taxonomy, we can intuitively observe a coincidence between - for example - the concept of subject and the concept of actor, and between the concept of productive action and the concept of goal. We plan to identify the join-points between Hohfeld and the i^* meta-model, which is able to define the behaviour of the stakeholders. Then, build an extended i^* meta-model, which incorporates the legal concepts. Such a meta-model should be able to both represent legal prescriptions and support selection among alternatives. In this framework modelling law is the first step of a process, which continues in the modelling of the intentional part of the domain - i.e., the stakeholders with their goals - and in the derivation of the requirements. Intentional compliance is then a property of a model of the domain, built after the joint legal-intentional meta-model.

This challenge calls for an answer to important questions. The adopted taxonomy offers high-level concepts; but the atomic elements of laws are deontic statements that prescribe what is forbidden or permitted. Is there any chance to support the translation from a formal model of the prescriptive texts to the rights that they generate? Another question concerns the *Intentional compliance* - we defined it as a property of a model built on the basis of a legal-intentional meta-model. What are the conditions for this property to hold, and which reasoning mechanism could verify it? Finally, after a legal model has been created, it is desirable that modelling the intentional domain will never violate the legal model. Is it possible to elaborate rules of intentional compliance, such that the intentional model is necessarily compliant? Hopefully, the answers to these questions will be the key actions of a fully supported pro-

cess for law-compliant requirements modelling and analysis.

7. Conclusion

In this position paper, we have pointed out the ontological difference between legal concepts and requirements. As alternative ways of being compliant exist, the choice among the alternatives needs to be supported while inferring the requirements for a compliant system. We claim that a systematic process can support such decision, but it should be rooted in a formal model that embraces all the legal and intentional concepts, and is founded on the notion of intentional compliance.

References

[1] T. D. Breaux, M. W. Vail, and A. I. Antón. Towards regulatory compliance: Extracting rights and obligations to align requirements with regulations. In *14th IEEE International Requirements Engineering Conference (RE'06)*, pages 49–58, Washington, DC, USA, September 2006. IEEE Society Press.

[2] J. Breuker and R. Hoekstra. Core concepts of law: taking common-sense seriously. In *In Proceedings of Formal Ontologies in Information Systems (FOIS-2004. IOS-Press*, pages 210–221, 2004.

[3] J. Breuker, A. Valente, and R. Winkels. Legal ontologies in knowledge engineering and information management. *Artifificial Intelligence and Law*, 12(4):241–277, 2004.

[4] R. Darimont and M. Lemoine. Goal-oriented analysis of regulations. In *International Workshop on Regulations Modelling and their Verification & Validation*, June 2006.

[5] 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, *19th International Conference on Advanced Information Systems Engineering (CAiSE'07)*, pages 218–232. Springer, 2007.

[6] P. Giorgini, F. Massacci, J. Mylopoulos, and N. Zannone. Requirements engineering meets trust management: Model, methodology, and reasoning. In *Proceedings of the 2nd International Conference on Trust Management (iTrust 2004)*, 2004.

[7] W. N. Hohfeld. *Fundamental Legal Conceptions as Applied in Judicial Reasoning*. Yale Law Journal 23(1), 1913.

[8] G. Sartor. Fundamental legal concepts: A formal and teleological characterisation. *Artificial Intelligence and Law*, 14(1-2):101–142, 2006.

[9] 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, Montpellier, France, June 2008. Springer.

[10] E. Yu and J. Mylopoulos. Understanding "why" in software process modelling. In *16th International Conference on Software Engineering*, pages 159–168. IEEE Computer Society Press, 1994.