The Agile Project: Reconciling Agility and Legal Accountability

Alexander Boer[1] and Tom van Engers[1]

[1] Leibniz Center for Law, University of Amsterdam,
Kloveniersburgwal 48 1000AB Amsterdam, The Netherlands
{aboer,vanengers}@uva.nl

Abstract. While every IT professional knows from experience that handling change is much harder and more expensive than building IT systems from scratch, only little scientific attention has been paid to prepare for change. This issue is of great importance since usually 80-90% of the total cost of ownership (TCO) of IT systems are maintenance costs. In our networked society changes may result from different sources. The recently started AGILE project addresses the legal dimension of management of organizational change processes. This paper introduces the AGILE project, and discusses the relation between sources of law and the business processes and services of the administrative organization.

1 Introduction

ICT has penetrated most organizations. It is an essential asset for public administration, and is becoming one for the judiciary. Administrative organizations, including law enforcement organizations, dealing with taxation, immigration, etc. have been using ICT to improve the efficiency of case handling (primarily aimed at reducing handling time and operational costs) and the quality of services, a development also known as eGovernment. ICT is also used to make law enforcement more effective (see e.g. [6]). Also in the judiciary we see developments aiming at improving efficiency and effectiveness (see e.g. [4]).

In the judiciary the application of Artificial Intelligence techniques is still very limited. Application of case management and ePublishing have become common and found their way to courts and law firms, but application of AI based applications, e.g. natural language processing for making sense of blogs with legal comments, or applying automated classification of email and intelligent routing and possibly preparing elements that could be used for solving the legal issues at hand, are still to be developed and implemented.

A wider range of Artificial Intelligence techniques have found their way into organizations. In this paper we focus on those applications for several reasons:

1. Similar solutions will be applied in the judiciary in the near future;
2. These applications are legal decision support applications and the knowledge in those systems is at least to a large extent legal knowledge being either declarative or procedural;
3. These applications have big impact on organization design (rule management) and an increasing impact on legal drafting.

These legal knowledge representation applications are interesting from a technical perspective but even more from a design perspective. Especially the separation of different forms of knowledge and the explicit relations to the sources of law can help to significantly reduce the TCO (see also [6,4]). This reduction of TCO is mainly possible because of the increase in traceability to the sources of law. Knowledge representation promises to reduce the time needed to establish the impact of change caused by changes in the sources of law.

Currently the interest in legal knowledge representation in public administrations is gradually increasing but also changing in nature. The increasing legal convergence and legal complexity, an increasing pace of organizational change in public administration drives this growing interest, and increased use of IT and web services. In this paper we will describe the shift from the utility of fielding computer systems built using a knowledge engineering approach to the potential utility of knowledge representation for comparative and for maintenance purposes, and for increasing the efficiency of the organizational change process itself.

Compared to the standards set by knowledge engineering research, fielded systems in public administration and elsewhere that use explicit knowledge representation to support decision making processes are technically and theoretically straightforward. The required transparency and the great challenge real world knowledge representation poses for the people implementing such systems, act as a natural limit to the complexity of these decision support systems. The required functionality rarely by itself justifies state-of-the-art legal knowledge representation. As we will argue in section 2 and 3, however, the knowledge management problems involved in managing the consequences of organizational change processes triggered by changes in the law, for business processes, services, databases, fielded applications, forms and documents, internal education, etc, make a considerably better case for some state-of-the-art concepts in legal knowledge representation.

The recently started AGILE project, introduced in section 4, addresses management of institutional change, sometimes driven by changing legislation and sometimes by environmental factors. This paper introduces the AGILE project and discusses the relation between sources of law and the business processes and services of the administrative organization. This account is a first, tentative step towards a design method that should help organizations to adapt to new or changing legislation.

2 Background

Inside public administrations, and on the interfaces between them, ICT and Internet have a large impact. Some decision making processes are nowadays assisted by computer applications, and others are more or less autonomously performed by the computer.

At the same time, service-oriented architectures are becoming the prominent paradigm for building enterprise information systems, also in administrative agencies. Service-orientation leads to new network arrangements between administrative agencies for sharing data, etc. This development in itself leads to attention for the adaptability and accountability issues that arise (cf. [5]).

A major category of those services are of course those performed for the main client categories of the organization, often generally the citizen. The implementation of these services on the Internet is what is often understood in the narrow sense by e-Government.

The services in question are in an administrative setting often implementations of public legal acts, performed by public legal personalities, based in formal legislation. Legislation gives administrative organizations public personality, defines what the core functions of public organizations are, and what services they provide. It guides how the organization subdivides itself into administrative units, how it organizes business processes inside the organization, and eventually how the functions of the organization are realized by civil servants and computer systems.

Business process design and design of specialized computer systems are both usually based on explicit models in modeling languages, like the Unified Modeling Language (UML) or, more recently in the context of the Semantic Web, the Web Ontology Language (OWL), of what the business process or computer application should achieve. These models are supposedly used as a specification of the objectives of an organizational change process or application development process. When legislation changes, these models are updated, and the organization's structures and computer programs have to be changed to conform to the models.

In the past these changes were conceived of as temporary interruptions of long periods of everything staying the same. This was certainly the case when the adaptation of existing systems was still considered a frightening prospect: things did change but the changes where carefully orchestrated to not impact existing procedures, network arrangements with other organizations, and computer systems. But as the perceived capacity of organizations to organize change processes increases, and the number of fielded computer applications increases, so does the pace of change in legislation directly affecting existing computer applications.

Tax legislation is for instance changed every year, leading to continuous adaptation of relevant computer applications for next year and the years after that, while the legislation of the present year and previous years is still being applied. In the business process design literature, awareness of this phenomenon has led to a new conception of the organization as an entity that is constantly in the process

of changing: the organization is constantly conceptualizing and comparing what it is and what it is becoming.

Attention for knowledge representation of sources of law is very often triggered by such administrative change processes driven by new legislation and other sources of law, such as case law and internal written policies.

The change processes triggered by the legal system are increasingly expensive, especially if they involve changes in ICT infrastructure. Knowledge representation is seen as a means to potentially reduce costs and increase efficiency through increased control over the knowledge dimension of the change process. Our past work for the Dutch Tax and Customs Administration (DTCA; cf. for instance [3, 6]) was for instance clearly related to the huge change process triggered by the complete overhaul of the Dutch income tax law in 2001. The Juridisch Loket (cf. [7]) project on pro bono legal assistance, and the DURP project on spatial planning (cf. [2]) were for instance also driven by an overhaul of legislation.

The law does not, however only play a top-down role, as knowledge in decision making and a yardstick for performance. Also the feedback to the legislator plays an increasingly important role (in the DTCA case notably; cf. generally [6]). There are certain traditional forms of feedback, like lessons learned about effectiveness, efficiency, and enforceability, and information about relevant changes in the environment of the organization which threaten enforceability, effectiveness, or efficiency. In addition the increasing formalization of network arrangements between involved actors spurred by the possibilities of ICT tends to lead to requests to the legislator to give such agreements and technical arrangements a basis in law, so that the partners on which the organization depends are bound to provide service as the organization prefers it.

Whether the legislator can play a constructive role in organizing public administration depends on the quality of this feedback. The nature of this feedback moreover strongly depends on the agility of the organization and it's infrastructure; The less prepared the organization is for change, the more problems it will encounter in implementing changes in law effectively, and it will surely be inclined to demand legal formalization of current network arrangements in order to buttress it's own current business processes.

Direct feedback from the responsible organizations about shortcomings of proposed legislation, as uncovered during modeling or simulation, played an important role in our work for the DTCA (cf. generally [6]).

The AGILE project, described in this paper addresses a more general problem conception: that of balancing the requirements of accountability to the law with the agility of the organization in a changing environment. But before explaining the aims of and approach taken in Agile we will first describe the general problem of deciding how to adapt to changes in the law and the interpretation problem.

3 Changes in the legal system

Adapting to changes in the law is a challenge for the Public Administrations that then have to adapt their decision support systems that are commonly used for massive case handling. Already in the eighties and nineties of the previous century architectures and design methods have been proposed that enhance legal effectiveness (see e.g. [6]). Following a knowledge engineering tradition these approaches have focused on the knowledge representation problems and the technical aspects that come with automated (legal) reasoning, rather than focusing on the process characteristics that affect the sources of law those systems are build upon. Needless to say that the way systems are build have great impact on their maintainability and adaptability but this is only part of the puzzle.

The authors of this paper, having been involved involved in many large and complex systems design projects, have noticed that many of the insights have still had little impact on the actual way systems are designed and produced. Unfortunately relationships between the specifications and their sources, or between the specifications and the components that realize them, or simple things like versioning and distinctions between design, test, accepted and production status are in many cases still lacking. This might explain the high exploitation cost of many of our governmental IT systems.

In this paper however we assume that despite the backlog that still exists in design practice, the technical architecture and methodological issues are relatively small problems and almost solved. The remaining major problem therefore is the interpretation problem. Before we can translate the sources of law that are expressed in natural language in a computational formalism, we should decide on the intended behavior or meaning of the constitutive elements of the norms that are expressed in the legal sources. As we will argue interpretation of the constitutive elements of the norms is rather problematic. How serious the problem is can be showed by looking at so-called 'open norms'

Concepts such as 'undue hardship' or 'generally accepted conduct' obviously require an operationalization by legal experts first before the norms that use these concepts can be transformed into a computational formalism. Alternatively one could decide to exclude such norms from the automated system that is to make decisions automatically, and leave the interpretation to civil servants, but that would imply that for every individual case we should find a specific interpretation.

The clear advantage of such approach would be that the civil servant could take into account all relevant contextual information. So potentially this could lead to a more justifiable decision. However on the other side this may lead to unfair-ness and inequality before the law. Another clear disadvantage is that stakeholders cannot make reasonable predictions about the outcome of the decision.

If one wants to include the 'open norm' into an automated system these norms have to be translated into 'closed norms'. For closed norms the interpretation, i.e. the semantics are completely clear. One naively could argue that the legislator should avoid the use of open norms for the purpose of easy implementation, but

this would neglect the fact that the law, being a steering instrument for society, will always need to be adapted to new circumstances. In order to prevent frequent changes of the law legislators sometimes prefer broader or open concepts. Open norms are also sometimes used to allow for different interpretations and thus circumvent potential political difficulties in the parliament.

However it is not just the legislator that has to decide on the desired interpretations of norms and the consequent choice in using closed norms or open ones, accepting that in the latter case the administrations will have larger freedom in interpreting them. We should realize ourselves that the interpretations discussed so far are ad-hoc interpretations. But since one can always defeat decisions up to the highest court, court decisions should also be taken into account since they reflect further interpretations. The problem for the administrations is that it may take a while before such decisions are made and acting on the impact these court decisions should have on the handling of other cases is not always easy.

If we look more closely at the interpretation problem we can observe that the hard problem is actual the mapping between the brute reality (the real world) and the institutional reality (the law), which is a qualification problem and depends on the aims and preferences of the person that makes that mapping. Legislators, policy makers in administrative organizations, and judges have a particular goal in mind when choosing an interpretation, but their plans may be wrong, i.e. their understanding of the world or of the effects of their interpretation doesn't lead to that desired goal.

While some attempts have been made to simulate effects of regulations, the application of such simulations is still very limited. Unfortunately only a few judges are aware of the relationships between the individual case decision at hand and the potential effects on massive client handling by the administrations, and rarely the published case decision contains indications about how to handle similar cases.

Legal quality in general would benefit if judges would consult the administrations about the (undesired) effects of their draft case decisions and make this clear in their decisions. This process could be formed in analogy to the process that legislators use to ask advice to both administrative organizations and the courts about potential operational effects of draft legislation.

Interpretation of norms is a challenge for legislators, the public administrations and the courts, as the above shows. If we want to improve legal effectiveness and take into account the need for adaptivity we should not limit ourselves to just the technical aspects of the supporting ICT-systems. 'Rule management' as it is called today, should particularly address the organization of the processes that should be aligned in order to achieve an agile solution.

In the next section we will shortly explain the Agile project in which we address these problems.

4 The AGILE Project

In the AGILE project (acronym for Advanced Governance of Information services through Legal Engineering) we aim at developing a design method, distributed service architecture and supporting tools that enable organizations - administrative and otherwise - to orchestrate their legal information services in a networked environment. The AGILE project started in the second half of 2008 and will last for four years.

At issue is the adaptability of ICT infrastructure, of business processes, and of data and knowledge within the organization, given changing legal demands and constraints.

4.1 Complex Adaptive Systems

Based on complex adaptive systems (CAS) theory, the project will develop a service modeling and design method that should help organizations to adapt to new or changing legislation. The proposed method should take the resilience of existing systems, dependencies on the environment, and the unpredictability of change processes explicitly into account (cf. generally [5]). The objective of this part of the project is to improve the adaptability of ICT infrastructure, of business processes, and of knowledge in the organization.

The essence of CAS theory is the study of systems built of individual agents that are capable of adapting as they interact with each other and with an environment, in order to understand how the individual affects system-level responses.

The underlying premise is that simply determining future needs and requirements is not the right approach, due to the inherent unpredictability of a complex environment and the fact that there are already many working (social and information) systems in place which cannot and should not be ignored.

4.2 Legal Concepts

The models to be developed in the project are intended for validation, and – in combination with an agent planning framework – for simulation of potential implementations. Direct deployment of services – which is possible with the selected Semantic Web technology – in a newly designed infrastructure is - as one should expect in a project definition that embraces gradual and continuous change - not explicitly intended.

Of specific relevance to a world dominated by written declarations and decisions, databases, web services, and changing sources of law is an account of formal acts, and of the act of providing evidence for a legally relevant proposition. In this account the concepts developed in the MetaLex standardization effort for instance play an important role.

The design methodology to be developed in AGILE will cover modeling concepts found in legal knowledge representation such as realization, reference, representation (from MetaLex), applicability, and constitutiveness (from our work in [1]).

The best-understood legal concept is undoubtedly the norm: there are certain things the organization is not allowed to do. This aspect of law is however also of ephemeral interest for AGILE. In the AGILE project it is rather the projection of organizational reality on abstract legal reality that is of interest.

4.3 Pilot Studies

Results from the research tracks discussed will be tested in the context of two actual business cases. One at the Dutch Immigration and Naturalisation Service (IND) and one at the Dutch Tax and Customs Administration (DTCA).

In both organizations, timely and efficient adaptation to changing legislation, case law, and patterns of behaviour accommodating or evading law in the relevant environment is seen as a main organizational objective.

5 Discussion

To demonstrate the potential utility of legal knowledge representation tech-niques for increasing the manageability of the organizational change process, sev-eral pilots are planned within organizations that do have to deal with such regular changes. In the IND the most obvious impetus is from changes in migrant and refugee streams, which always lead to new legal issues. In the DTCA change is predictable as a clock: taxation appears to change as often as the legislator believes the DTCA can handle it, and sometimes more often.

A major obstacle for assessing the utility of the design method developed in AGILE is of course the time frame: whether the investment in knowledge representation (the TCO, so to speak) is worthwhile can only be judged in relation to its lifespan. The pilots should in this respect be considered formative evaluations.

Finally, the critical reader may wonder to which extent the proposed interpretation of the relation between law and public administration is generally applicable. It is certainly the case that there are considerable differences between approaches to law making that need consideration.

How the legal system deals with complexity is in the end the less important factor; The external factors that increase complexity of change processes – legal convergence, increasing dependence on network arrangements with other actors, increasingly costly infrastructure – are certainly present all over Europe and, perhaps to a lesser extent, the rest of the developed world.

If we want to make administrative organizations more agile we shouldn't however limit ourselves to developing ICT-tools en methods. It is equally important to

strengthen the links between the legislator, the policy departments of the administrations, and the courts in order to develop shared ad-hoc interpretations. The same holds for those situations in which post-hoc adjustments have to be made. Also in those post-hoc cases the legal system would benefit from taking a broader perspective. Both cases would benefit from an architecture that clearly distinguishes the different interpretation problems and allows for testing, i.e. simulating effects of such decisions. With the Agile project we hope to set a first step towards such a environment.

AGILE is a Jacquard project funded by the Netherlands Organization for Scientific Research (NWO). In the AGILE project, The Leibniz Center for Law of the University of Amsterdam cooperates with the Technical University of Delft, which has experience in the application of CAS theory to organizations. The IND and two companies, O&I and BeInformed, provide matching effort to the project. We would also like to thank the DTCA for its cooperation in the project.

References

[1] T. van Engers. Legal engineering: A structural approach to improving legal quality. In A. Macintosh, R. Ellis, and T. Allen, editors, Applications and Innovations in Intelli-gent Systems XIII, proceedings of AI-2005, pages 3–10. Springer, 2005.

[2] Hilhorst, R.A., Van Engers, T.M., 2009, E-dossiers at the Dutch Council of State, in Proceedings of the 12th International Conference on AI and Law (ICAIL), ACM, ISBN 1-60558-597-0.

[3] M. Janssen. Adaptability and accountability of information architectures in interorganizational networks. In ICEGOV ’07: Proceedings of the 1st international confe-rence on Theory and practice of electronic governance, pages 57–64, New York, NY, USA, 2007. ACM.

[4] Boer, T. van Engers, and R. Winkels. Using Ontologies for Comparing and Harmoniz-ing Legislation. In Proceedings of the International Conference on Artificial Intelli-gence and Law (ICAIL), Edinburgh (UK), 2003. ACM Press.

[5] T. van Engers, R. Winkels, A. Boer, and E. de Maat. Knowledge management and the dutch legal aid service counter. In J. J. Schreinemakers and T. van Engers, editors, Advances in Knowledge Management, volume IV, Würzburg, 2006. Ergon Verlag.

[6] Boer, T. van Engers, R. Peters, and R. Winkels. Separating law from geography in gis-based egovernment services. Artificial Intelligence & Law, 15(1):49–76, 2007.

[7] Boer. Legal Theory, Sources of Law, & the Semantic Web. Frontiers in Artificial Intelligence and Applications 195. IOS Press, 2009. To appear.