AgentWorks Prefeasibility Project Report

B. Cohen[1], P. Boxer[2], A Martino[3]

Introduction

The partners in the project are Boxer Research Ltd. (BRL) and City University, in the UK, and Aegis, in France. They have been investigating the methodological, theoretical and computational issues involved in integrating two different, but related, approaches to the representation of knowledge in enterprises: BRL's *Articulated Enterprise Modeling* and Aegis' *Distributed Agent-based Architecture*.

The former is a technique for knowledge elicitation and intervention in enterprises that has been developed and practiced for 15 years by BRL. It is designed to assist the enterprise to understand the complex relationships between the demand situations that it seeks to address, as it elaborates them into requirements for products and services; and the capabilities that it has at its disposal, as it coordinates them into the products and services that its business units present to the marketplace. The composition of these two 'stratified articulations' of the enterprise's knowledge of itself is the source of the enterprise's 'business rules': those obligations and responsibilities shared by its actors in pursuit of its intentions. Empirically derived metrics have been defined over these representations that indicate the presence of 'holes' between the enterprise's intentions and its practices, the 'repair' of which may require the reformulation of both articulations.

The latter is a computational paradigm derived from the principles of object-oriented, distributed and artificial intelligence systems. Agents are autonomous software devices that represent the actors of enterprises by populating a distributed network, their interoperation being mediated by representations of their actors' knowledge. The FIPA consortium, which defines the *de facto* standard architectural primitives and knowledge representation notations, commissioned Aegis to develop the first FIPA-compliant 'demonstrator' system. In Aegis's development environment, software implementing the agents and their communications protocols may be compiled from knowledge representations formulated as Conceptual Graphs (CG) [Sowa]. These representations rely on definitions, in first order logic, of the ontological concepts and relationships in terms of which the semantics of actors' statements may be constructed and reasoned about.

The project has addressed issues of:
structure: how the structural principles of the two approaches -- stratified articulations and nested conceptual graphs, respectively -- are related;
theory: in what mathematical framework might the semantics of stratified articulations be cast, and how that framework is related to the first order logic of conceptual graphs;
praxis: comparing and contrasting ways that the notations and tools of the two approaches are embedded in their practices of knowledge elicitation and organisational development;
tools: how the tools used by the two approaches might be composed and, if necessary, extended to provide more comprehensive and powerful support for the praxis.

This report presents explanations of the two approaches, identifies progress made on the issues investigated, and suggests directions for future work together with the benefits that might be expected from its successful completion.

[1] School of Informatics, City University, London
[2] Boxer Research Ltd, Chiswick, London
[3] Department of Mathematics, City University, London

Elicitation and Representation of Enterprise Knowledge

In **Articulated Enterprise Modeling**, the enterprise's knowledge of itself is elicited in terms of three 'stratified articulations', as follows:

Existential Articulation

The only means available to an enterprise for inducing effects in its world are its *capabilities*. To describe an individual capability requires the representation of:

- a model of (the part of) the world in which it has its effects, i.e. those 'states of affairs' whose observation can distinguish among the behaviours produced by the use of a capability;
- the names of the 'processes' through which it can produce its effects, where each process effects a transformation between distinguishable states of affairs; and
- a description of all the possible behaviours it can produce, i.e. all the sequences of processes in which it can engage and all the states of affairs that it can influence by engaging in these sequences[4].

These capabilities can therefore be thought of as *coordinations* of behaviour. An enterprise usually arranges for its capabilities to inter-operate in various ways so that, in concert, they achieve effects that the enterprise considers to be of value to it. Such *coordinations* of capabilities may be represented by identifying them with the effects that they collectively achieve. Thus, an enterprise will identify certain coordinations of capabilities with *competencies* that are typical of the *workgroups* used by the enterprise, and these competencies may themselves be coordinated within identified *task groups* that are charged with the achievement of certain *outcomes*. Such *stratification* of coordination continues through *business units*, which coordinate task groups so as to deliver *products* and *services*. All of the strata up to task groups may be considered to be 'inside' the enterprise, but the strata above that level, starting with business units, become observable 'outside' the enterprise. Thus the *business units*, producing products and services, are themselves coordinated by *market propositions*, collections of which are themselves coordinated in order to address the *customer situations* to which the enterprise chooses to direct its attention[5].

Each stratum of coordination may itself be multi-layered and *underdetermined*, in that it admits more than one way to achieve the effects required of it. But by the same token, the stratified articulation omits some coordinations of capabilities that are possible but do not feature in the enterprise's model of itself. Such a stratified description, which is called an *existential articulation*, may therefore be said to describe the enterprise's *know-how*: what is referred to in BDI modeling as its *beliefs* about *what it is possible to do*.

The lowest level of coordination in the existential articulation is represented by its capabilities, which denote its ability to engage in certain coordinations of behaviour (C^1). Every higher stratum describes sets of alternative behaviours expressed in terms of coordinations of lower levels of behaviour (coordinations of capabilities - C^2).

Referential Articulation

The existential articulation depends implicitly on the value that an enterprise places on the effects it can achieve. Changes to that valuation might lead the enterprise to reconsider its organisation of coordination, or even the capabilities on which it is based. However, the existential articulation gives no indication of the origins of that valuation. To describe that aspect of the enterprise requires the representation of:

- a collection of *drivers* that it perceives as generating, and being punctuated by,
- those *demands* to which it chooses to respond; and

---

[4] Since this part may be infinite in extent, it must be presented in the form of a finite model from which the behaviours may be inferred, i.e. in which the behaviours comprise the model's *closure*. Many formal notations, including set theory, first order logic, Petri Nets and various process algebras, have been used for this purpose [Sow]. Algebraically, any such model is a graph [Mar] whose vertices denote states of affairs and whose edges denote processes.

[5] This meta-ontology defines the schema for a relational model which, when populated by specific data about an enterprise ontology, may be processed by BRL's DECON tool (see below). A similar schema is defined for the referential articulation. The outputs of the DECON tool are analysed by the PAN tool to calculate the values of certain topological metrics that are used to reveal the existence of 'holes' in the enterprise ontology (see under 'metrics' below). [Box]

- a description of all those possible configurations of demand which it can support, i.e. the configurations of drivers which its behaviours can 'satisfy'.

An enterprise typically elaborates demand in the form of *anticipations of satisfaction* i.e. that which capabilities, and coordinations of those capabilities, would satisfy. This analysis typically proceeds by identifying successive levels (stratified, as before) at which demands and drivers can be brought into relation with each other as anticipations. Thus *problems* are particular anticipations of configurations of demands and drivers, *value propositions* are anticipations of configurations of problems, and *value profiles* are anticipations of configurations of value propositions. These strata denote the *identity* of the enterprise's relation to demand. Further elaboration generates strata that identify *service requirements* that are anticipations of configurations of value profiles, *integrating processes* as anticipations of configurations of service requirements, and *capability outcomes* as anticipations of configurations of integrating processes. These latter strata denote the *viability* of the enterprise's anticipated relation to demand satisfaction, from which its *intentions* to achieve particular forms of outcome in the existential articulation may be derived. Again, each stratum may be multi-layered and underdetermined, in that it identifies more than one way in which satisfaction may be anticipated. This stratified articulation will also omit some ways of anticipating satisfaction that are conceivable but do not feature in the enterprise's model of itself. Such a stratified description, which is called a *referential articulation*, may therefore be said to describe the enterprise's relation to demand: what in BDI modeling is referred to as its *desires* and *intentions* (see figure 1).

The surface level of anticipation in a referential articulation is represented by the problems that are anticipations of particular configurations of demands and drivers (A^1). Each deeper stratum describes sets of alternative anticipations of satisfaction expressed in terms of configurations of anticipation of satisfaction at shallower levels (anticipations of problems - A^2).

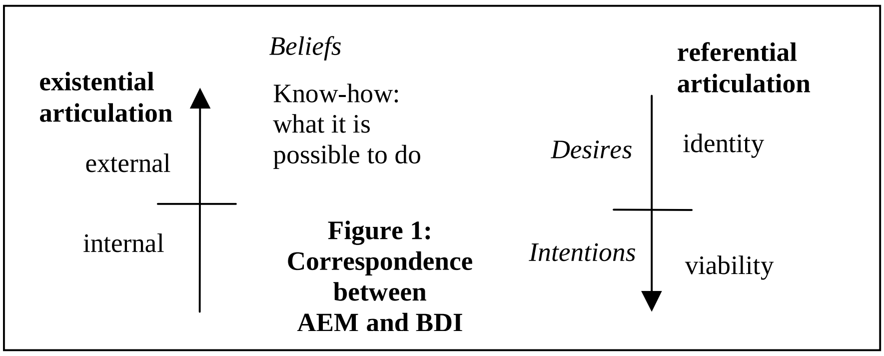

Figure 1: Correspondence between AEM and BDI

Deontic Articulation

The existential articulation provides a description of the *behavioral ontology* of an enterprise, or set of enterprises. The referential articulation provides a description of the *demand ontology* – the organisation of demand in relation to an enterprise. By bringing these two together, one should be able to construct a common *enterprise ontology*, but one is likely to be frustrated by the absence of direct correspondences among the different anticipations of satisfaction and coordinations of behaviour. For example, one might find that some elaboration in a referential elaboration leads to a concept that is not realisable by any coordination in the existential articulation. This indicates the existence of a 'hole of the second kind' [Mar] in the existential articulation. Similarly, holes of the second kind in the referential articulation would indicate that something realisable does not appear in any elaboration of demand. The aim of Articulated Enterprise Modeling is precisely to highlight these 'holes' and to propose solutions to them. These proposed solutions are of three kinds:

(i) they may introduce new conceptual elements (at the lowest existential and shallowest referential levels). This would be tantamount to introducing new knowledge about what it is possible to do, and/or about what is demanded.

(ii) They may introduce into the enterprise ontology new coordinations/anticipations over them (C^2 and A^2) by creating new knowledge that the enterprise has about itself.

(iii) They may introduce constraints into how the existing enterprise ontology is deployed in order to compensate for the existence of the holes. These introductions of constraints then become coordinations of the coordinations of capabilities in the existential articulation (C^3), which are also anticipations of the anticipations of problems in the referential articulation (A^3).

The role of proposing this third kind of solution is essentially *deontic* in nature, and itself leads to a stratified knowledge about how to 'manage' the enterprise ontology. It is what characterises the knowledge of an *actor* as the capacity to make choices about how the enterprise's ontology is to be used to create value - to distinguish the notion of an actor at this C^3 level from an actor as a capability (C^1), and/or as the ability to coordinate capabilities (C^2).

The **DECON** tool, that supports BRL's practice of Articulated Enterprise Modelling, provides for two 'knowledge bases': one 'user-defined' (UKB) and one 'expert-defined' (EKB). The EKB contains partial models, imputed by an enterprise's actors, of the 'semantic domains' in which models of their enterprise can be expressed. It imposes syntactic and semantic constraints that define 'legal forms of expression' in those domains. DECON provides for the definition of both 'simple' and 'complex' object classes and the types of relation in which they may participate in the UKB ontologies. Simple objects are not further 'decomposed', and complex objects are defined by stratified compositions, axiomatically restricted as necessary, of simple or complex objects. The UKB is derived from statements made by the actors about the enterprise ontology. Each such statement should, of course, be an instance of a legal form of expression as defined in the EKB.

In **Conceptual Graphs**, ontologies are also stratified, but this stratification reflects the logical need to distinguish between statements that refer to things in the world and 'meta-statements' that refer to the status of such statements (and so on, for successive meta-levels). For example, the statement: 'This statement is false' is paradoxical as it stands since, whether it is true or false, it contradicts itself. This kind of paradox can be avoided by allowing such observation to be made only as meta-statements, as in: 'The statement made earlier was false', which does not refer to itself. This scheme has the distinct logical advantage of accommodating arbitrary numbers of meta-levels while retaining the analytical power of first-order logic. This stratification into meta-levels, called *contexts* in CGs, can be used to represent the knowledge in the existential articulation. One could, for example, establish an inner context in which a logical system described the basic behaviours of the enterprise's capabilities, whose coordination could be described by an enclosing context that defined a logical system for communicative behaviour. Each successive stratum of coordination (C^2) would then be the subject of a new enclosing meta-level context.
In CGs, contexts can also be used to 'situate' statements, in time, in space and with reference to the subject, as in Situation Theory [Bar]. This distinction between *subject-referenced* and *object-referenced* contexts is needed to present appropriate views of the enterprise's *knowledge of itself* (see the remarks under *Mathematical and Computational Issues*, below).
CGs of this nature are typically interpreted as describing the *beliefs*, *desires* and *intentions* of the actors whose knowledge they represent.

In the **BDI** approach to agent-based systems, these *modalities* are investigated by applying suitably defined modal logics to the CG representations of agents' knowledge. The semantics of these modal operators is defined in terms of behaviour graphs, whose edges represent actions that might be performed by certain agents and whose vertices represent states of affairs that pertain at moments in time in the world in which those agent act. This kind of semantic domain is common to BDI semantics, to conceptual graphs and to the existential articulation.
The strata of the referential articulation could also be represented in stratified CGs, the innermost context representing primitive demands and drivers and successive meta-levels representing elaborations of anticipations of satisfaction (A^2). But such a representation would not admit an interpretation in behavioral graphs, and therefore could not be subjected to modal analysis in the BDI framework.

Composition Risk

In BDI systems, composition of agent models is required to formalise *co-operation* (in teams with a common goal), or *social conventions* (where mutual beliefs or joint intentions can be recognised), or *group know-how* (where a multi-agent system is to be treated as a single agent). These are currently the subject of a great deal of research because modal treatments of individual agents do not compose easily.
In articulated enterprise models, these issues of composition cannot be ignored.
The enterprise is always a composition of actors-in-roles, so that the question of composition has to be addressed not only for each articulation (see the discussions of existential and referential composition in

[Mar]), but also for the composition of existential and referential. Such a composition into an enterprise ontology, and its subsequent consequences for a deontic articulation, cannot be simply an embedding of the former into the latter. The holes that the existential and referential articulations induce in each other indicate ontological mismatches, proposals for the repair of which would require adding or modifying concepts, relations or even contexts in each articulation: new implementations of an A^2, or new anticipations for a C^2. In terms of their stratifications, we can see that:

- there is an articulation of demand (identity) that must be elaborated in terms of a specification of outcomes that must be achieved (viability);
- these outcome specifications must be formulated in terms of the behaviours in which the enterprise must participate; and
- these behaviours must be supportable by the capabilities and competencies of the enterprise's actors and agents (including those known, in the computing literature, as the 'legacy system').

These different forms of compositional act relate to the *risks* that the enterprise runs at the deontic level in trying to achieve a desired, composite behaviour through the formation of a particular business:

1. *Will the business concept capture the demand*, or how can one determine whether any composition of the demands being addressed by the enterprise may be elaborated in viability term?
2. *Will the planned design of the business deliver the business concept*, or is there a conceivable coordination of competencies and capabilities that might support these forms of viability?
3. *Will the legacy environment be capable of supporting the business design*, or can the enterprise's current infrastructure support these conceivable coordinations? In other words, can the beliefs embedded in its current existential articulation, and supported by its legacy systems, be composed to deliver the desired forms of composite behaviour?

Mathematical and Computational Issues

During the pre-feasibility study, a great deal was learned about the relationships between the two approaches. A sound mathematical foundation, based on graph theory and category theory, was constructed for Articulated Enterprise Modeling [Mar]. Coincidentally, a major textbook on Conceptual Graphs was published [Sow] and the Aegis demonstrator system was displayed at the FIPA conference in London. But more importantly, both sides perceived gaps in their mathematical and computational foundations and agreed on directions for developments that might fill those gaps, and thereby unify the two approaches, methodologically, analytically and computationally.

A combination of the two forms of representation could have major repercussions, providing a computational framework in which *the enterprise's knowledge of itself* might be represented in a standard form that would support:

- logical manipulation, i.e. formal reasoning;
- ontological extension, i.e. increasing granularity and modifying stratification;
- calculation and interpretation of topological metrics, especially those concerned with composition risk;
- presentation of different 'views' of the enterprise's knowledge, suitable for different purposes of the enterprise, both as the *subject* and as the *object* of analysis;
- compilation of an existential articulation (subject to the constraints in the deontic) to implementations, where appropriate, of agents in a distributed architecture;
- compilation of the articulated enterprise model to an 'enterprise KBS', comprising suitable 'data platforms' on which management information systems and decision support systems might be executed.
- the use of 'passive agents' to collect data from the interactions of actors with the enterprise's legacy environment, and from their 'conversations' with each other, to support a new praxis of knowledge elicitation and intervention.

In order to progress the development of this powerful framework, the following issues have to be resolved:

Metrics

In Articulated Enterprise Modeling, the measurement of *granularity* and *eccentricity*, using graph-theoretic techniques derived from Q-Analysis [Atk], provides important insights into the presence and nature of ‘holes’ which, in turn, indicate the magnitude of the composition risks facing the enterprise. No such metrics are defined over CGs.[6]

It is an open question as to whether the nested contexts of CGs admit these kinds of topological metrics and, if so, what interpretations they might support and how these interpretations might affect the practice of knowledge elicitation and the implementation of distributed agent-based systems.

Logics

Conceptual Graphs support powerful computational reasoning thanks to their sound foundation in first order logic. They also support modal reasoning thanks to their interpretation as BDI systems.

The DECON tool is also a first order logic system (essentially a knowledge base represented in the Horn-clause logic of Prolog), but it has been used only for limited forms of reasoning over doubly articulated enterprise models.

There is obviously scope for extending DECON to encompass Sowa’s ontological hierarchy, providing terms in the knowledge base with suitable logical axioms and thereby providing much greater support for the semantic analysis of the existential and referential articulations.

It is an open question, however, as to whether any CG-based tool could have sufficient expressive power to represent the deontic articulation.

Ontologies

The ontological content of an articulated enterprise model is contained in its ‘naming dags’ [Mar]. There is clearly some relationship between the structure of these graphs and the nested contexts of CGs. However, since the naming dags are not related to any underlying ‘universal’ ontology by axiomatic definition, there is, as yet, no way to unify the two representations.

Some evidence from computer science suggests that the ‘composition’ of articulated models, both within the existential or referential articulations and between them, requires the category theoretic operation of ‘superposition’ [Fia], which is not equivalent to the logical operation of ‘unification’ in first order logic. This suggests that the operations on ontological structures supported by the CG framework are inadequate for the domain of articulated enterprise modeling. Their extension in this direction would have significant effects on many other applications of ontology and semiotics.

Trade-offs

Case histories in articulated enterprise modeling have demonstrated that the repair of ‘holes’ detected by the application of topological metrics might introduce unacceptable levels of ‘deontic complexity’, that is, a proliferation of ‘business rules’ that can assume Byzantine proportions. On the other hand, changes to the underlying enterprise ontology in order to introduce low deontic complexity might introduce coordinations that require unacceptably high ‘traffic’ among the responsible agents. There are, at present, no mechanisms in either of the approaches that could evaluate this trade-off.

References

[Sow] John F. Sowa, *Knowledge Representation: Logical, Philosophical and Computational Representations*, Brooks/Cole, 2000

[Mar] Armando Martino, *Enterprise Ontology*, Agentworks Project Report, Jan 2000

[Box] Philip Boxer and Bernard Cohen, *Analysing the ‘lack’ of Demand Organisation*, in Daniel M. Dubois (ed) *Computing Anticipatory Systems*, American Institute of Physics, 1998, pp157-181

[Atk] W. Atkin, *Q-Analysis: a hard language for the soft sciences*, Futures December 1978, pp492-499

[Fia] J.L.Fiadeiro, A.Lopes and T.Maibaum, *Synthesizing Interconnections* , in *Proc. IFIP TC 2 Working Conference on Algorithmic Languages and Calculi*, D.Smith and J.P.Finance (eds), Chapman Hall.

[Bar] Jon Barwise and John Perry, *Situations and Attitudes*, CUP, 1998.

[6] These metrics are applied to the strata of an Articulated Enterprise Model by the PAN tool, which complements DECON in BRL’s practice.