Chapter 1

Triply-Articulated Modelling of the Anticipatory Enterprise

P. Boxer
Boxer Research Ltd <p.boxer@brl.com>
Prof. B. Cohen
City University, London <b.cohen@city.ac.uk>

1.1. Risk and System Orders

The engineering of products and services is exposed to three kinds of risk:
performance risk: that a component might not work as specified by the supplier;
composition risk: that hitherto independent components might be not be able to interoperate in such a way as to provide an intended product or service; and
implementation risk: that the intended product or service might not satisfy the client's demand when used in its actual context-of-use, even though it meets the specification agreed with the supplier.
Performance risk involves exposure to errors of execution, the supplier being unable to sustain the performance of its constituent capabilities, and is typically evaluated by behavioural analysis of models of those capabilities, using simulation, analytical techniques or exhaustive testing. These forms of analysis are computationally effective for first-order systems.
Composition risk involves exposure to errors of planning, the supplier being unable to ensure the validity of its approach to composing constituent capabilities in order to deliver the product or service that, it believes, the client expects. In order to evaluate its composition risk systematically, the supplier would require a model of its own organisation of composition and have the ability to change this organisation without damaging its integrity. This self-organising, or autopoietic [Maturana 1980] property is characteristic of second-order systems and distinguishes agent-based architectures from their simpler, first order, software counterparts. Second-order systems pose analytical problems related to those of higher-order logics (involving statements in which quantifiers range over predicates), which are tractable only when the closed-world assumption is valid, as it is in most of the knowledge representation schemes

currently being used in agent technology [Sowa 1999]. Unfortunately, that assumption is explicitly denied in open systems.
Implementation risk involves exposure to errors of intention, the supplier being unable to guarantee that products or services will satisfy the client's demand when deployed in its context-of-use. Here, the client is an *actor*: an *anticipatory system* [Rosen 1991] who derives her demands for services from her formulation of herself as their context-of-use. The dissatified client experiences a *value deficit*: the gap between what she wants and what the service delivers. The greater this gap, the more *asymmetric* becomes the demand with respect to what the supplier can deliver. (We assume that value deficits exist for the *enterprise* as actor as much as they do for the *individual*.) This risk is particular to the 'relational' response to asymmetric demand (the 'positional' response being to ignore the value deficit associated with it [Porter 1980]). As an embodied actor, the client's model of her own organisation of demand orients her behaviour leading her to specify services that might satisfy her demands and to anticipate, possibly erroneously, their satisfaction by those services. This apparently teleogical, or *anticipatory*, property is characteristic of third-order systems, which exhibit *closed loops of entailment* [Rosen 1991]. The logical problem posed by these third-order systems is related to that of composing statements made in different modalities, identified by C. S. Peirce, in relation to his Gamma graphs [Peirce 1958], as an outstanding difficulty that future generations would have to tackle. Most knowledge representation languages are based on Sowa's Conceptual Graphs [Sowa 1999] which, deriving directly from Peirce's graphs, do not solve this problem.

1.2. Asymmetric Demand

The actor's experience of her own needs, and of the things in the world that relate to these needs, are parts of her *semantic formation*, which changes as she learns more about her own needs, as well as about what is available in her world. Her experience of herself as context-of-use will always be, to some extent, asymmetric, leaving 'something to be desired'; that is, that she will never know everything about her own needs. When suppliers assume that demand is symmetric — independent of the client's semantic formation — the client becomes a *customer* and competitive intensity tends to increase as more suppliers enter markets through the effects of globalisation. Further, both the *economies of scale*, induced by commoditisation, and the *economies of scope*, induced by providing for variety of supplied services, tend to increase with *digitisation*, the effects of computational and networking infrastructures on the economics of suppliers' abilities to respond. Under both of these influences, demand asymmetry tends to increase precisely because suppliers' success in satisfying symmetric demand leads clients to expect ever more value for money with respect to their specific contexts-of-use. As value deficits grow and demand becomes increasingly asymmetric, strategy is under pressure to become relational. This requires a shift in the place where power is held in the enterprise, from the centre, in the form of a positional strategy for the enterprise as a whole, towards the edge, where there is relational knowledge of clients' contexts-of-use [Alberts 2003]. As a

result, the focus of the enterprise has to move towards *managing the risks* associated with deploying multiple strategies, each particular to a client relationship.
We have encountered these phenomena in fields as diverse as the EC *AgentCities* Programme, Healthcare systems and military strategy studies, where both clients for and suppliers of services have exhibited similar consternation at the intractibility of the problems they entail. We believe that these fields, and many others, share four categories of long-term challenge:

- Automation (Service Management):
- Interoperability (Service-Service Communication):
- Coordination (Service Orchestration): and
- Knowledge Acquisition (Interfaces between worlds).

'Interfacing' between the worlds of service users and service suppliers could be effective only when the same models could be assumed for all actor-observers involved in the composition of services. Since this closed-world assumption does not hold in a truly open environment, the processes of orchestration and interfacing cannot be fully automated. They require the active participation of the actors involved in such transactions, including the users whose demands those suppliers are seeking to satisfy, supported by computational tools in which their respective models are respresented, composed and analysed for consistency. This paper presents such a tool, called PAN, and illustrates its use in the field of Health Care.

1.3. Collaborative Composition

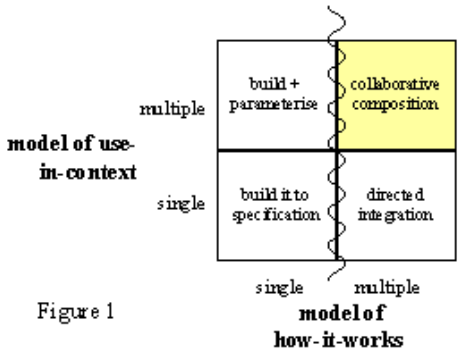

Figure 1

The shift to a relational strategy separates the organisation of the services from the organisation of demand. Thus, in fig. 1, the nature of the services, which are more or less over-determined by the causal nature of the processes from which they are constructed (the model of how-it-works), are distinguished from the nature of the demand, which introduces constraints relating to the client's particular context-of-use (the model of use-in-context). In the left column, there is a single model of how the service works, either built for a particular use-in-context (bottom-left), or parameterised so that it can be adapted to multiple forms of use (top-left). On the right, however, there is no single model, but multiple components, each with its own model, composable with other models, to form a *System of Systems* (Maier 2000). In the single use-in-context, a designer can impose, externally, a composition on these components (bottom-right), in which the component systems maintain an ability to operate independently, but their normal operational mode is subordinated to the usage requirements. This *directed* approach is typically used in large, complex System-of-Systems projects. However, the presumption of a symmetrical relation to demand, arising from the *a priori* construction of the service, cannot be maintained in the top-

right quadrant. There composition has to be a dynamic response to the user's asymmetric demand and we need to define components with a *granularity* and *stratification* that can support the requisite variety of uses-in-context. Granularity refers to the level of abstraction at which component services are defined and stratification to the way the different layers of composition of component services relate to the end use. *Orchestration*, the act of determining which components get to participate in the collaborative process of composition, is the form of composition demanded by an actor of an agent-based architecture. The more asymmetric the demand, therefore, the more granularity and stratification have to be elicited from the nature of the context-of-use in which the demand arises, and the more composition has to be collaborative. PAN is designed to support this collaborative process.

1.4. A Topical Exemplar: The Electronic Health Record

Asymmetric demand is intrinsic to health care. Patient-related information is generated from within a particular clinician's practice, and its communication among healthcare practitioners, particularly in the provision of shared care under multiple clinics, practices and practitioners, depends on a shared understanding of the patient's context-of-use — in this case, the patient as context to their 'condition'. Further, clinical approaches to treating the patient's condition within the patient-as-context differ from clinic to clinic, precisely because of their differing relations to this context-of-use. The Electronic Health Record (EHR) is an attempt to provide a supporting infrastructure governed by a universal ontology specific to healthcare. It is the subject of international standardisation activity that has been ongoing for over a decade. Although WHO, G8 and all national governments consider such standards to be significant factors in the improvement of public health, they have had limited uptake. For example, European 'demonstrator' projects were unable to share information even though they conformed to the European standard (CEN ENV 13606), because each 'adapted' it in ad hoc, and mutually incompatible, ways.

In fig. 2, which has the same structure as fig. 1 but is specific to the Healthcare domain, the bottom-left quadrant contains those situations in which everyone in the community agrees both on the relevance of the standard to the situation in which it is used, and on the capabilities to which the standard is to apply. In the bottom-right quadrant, a simpler approach can be admitted, but only as long as the situations/contexts to which the standard applies can be separately identified. In the top-left quadrant, differences in the way services are supplied necessitate the definition of distinct abstract domain models of ever-increasing complexity.

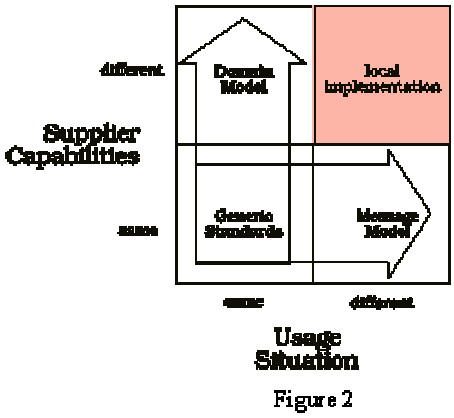

Figure 2

The progressive differentiation of care specialisms (i.e. supplier capabilities), necessitates the development of specific domain models insofar as the generic standards prove ineffective. And the challenge of asymmetric demand, represented here by the movement towards increasingly differentiated usage situations, arises as the satisfaction of basic health care needs gives rise to demand for treatment of increasingly complex and chronic conditions. Where there are well-defined treatment strategies (e.g. CPGs) for these conditions, situation-based models, typically based on standard message formats, are used to relate them to the patient's context-of-use. But in the top-right quadrant, the particular nature of the patient's context-of-use becomes inescapable and the need for relational strategies in response to the patient's value deficit — their particular needs for treatment — becomes overwhelming. This quadrant presents a severe challenge to the healthcare standards endeavour. The pragmatic argument is that 80% of what-is-going-on can be governed by standards, the rest being insignificant, but this ignores the sheer complexity and political sensitivity of the task of finding out which 80% can ignore asymmetric demand, and can, therefore, be standardised. Far from being insignificant, the top-right quadrant houses the clinician's crucial role in formulating the appropriate response to the particular patient and, as a result, deciding what can be ignored.

Thus, in our approach, we reverse the usual standards procedure and seek, first, to understand the nature of the clinician's practice within which the information generated by a clinic is embedded. This practice is the context within which negotiation must be conducted to determine what can safely be made generic, standardised and promulgated within a universal ontology (the symmetric part), and what must remain particular to the demands of the patient's condition (the asymmetric part), without jeopardising the consistency of the whole.

1.5. PAN's Triple Articulation

PAN enables the relational enterprise to manage the three kinds of risk to which it is exposed, and to make appropriately informed strategic decisions, by facilitating the representation, composition and analysis of three models, or *articulation*, that are inherent in the way the actor-observer constructs her world. For the purposes of illustration, we will relate this triple articulation to clinicians' practice.

The **existential** articulation models behaviour: the processes that the clinician deploys, the events in which she participates and the structures within which the processes she uses are coordinated.

The **deontic** articulation models the organisation of mutual obligations and constraints on behaviour: the outcomes that the clinician undertakes to deliver, the synchronisations to which these outcomes are subject, and the programme structures within which they are organised.

The **referential** articulation models the organisation of demand: the drivers that drive the clinician's patient's condition, expressed as situations that the patient presents to the clinician and as the patient qua context-of-use in which those situations are embedded.

These models are distinct, the clinician's practice being formed through the particular way she brings them into relation with each other, forming a composite articulation that is particular to the patient's context-of-use. PAN supports this process of composition by checking for consistency and identifying gaps in the composite articulation which represent a lack of interoperability among its constituent parts.
The actor-observer is here construed not just as having a point of view, but as embodying a particular, problematic form of relation to demand, in which anticipation does not merely specify that which is to be satisfied but directly affects the nature of the composition. The triple articulation is a model of the actor-observer's model of her own semantic formation — in this case, the clinician's formation of her practice. Here, performance risk arises from errors of execution described in terms of the existential articulation; composition risk from errors of planning in the relation between the deontic and existential articulations; and implementation risk from errors of intention in the relation between the deontic articulation, on one hand, and the existential and referential articulations on the other. Clearly, performance risk can be evaluated by objective modelling of the existential articulation alone, and therefore benefits little from PAN. But the other two forms of risk reflect errors in the actor-observer's articulation that can be observed and rectified only through the actor-observer's own reflections. PAN enables the actor-observer to mitigate these risks by identifying them and presenting them for her consideration, thereby enabling her to propose modifications and/or refinements to her articulations, and to anticipate their effects in relation to demand.

1.6. Using PAN

Interaction with PAN may be based on any method of relational modelling. The examples in fig. 3 describe:
(i) the organisation of demand, referred to here as a value ladder (for a patient, it takes the form of a referral pathway);
(ii) organisational hierarchy;
(iii) horizontal relationships across the organisation; and
(iv) the processes and infrastructures of the clinic.
The schema elicited by relational modelling is mapped into the triple articulation by a *meta-knowledge base* (MKB) written in the PAN notation. An MKB populated by a model formulated by an actor-observer, becomes an *observer's knowledge base* (OKB). More than one actor-observer's knowledge is likely to be relevant in the formation of the clinician's practice, not only the clinician's, but those of other administrators and clinicians, as well as that of the patient herself. Thus PAN needs to support the identification of objects from different OKBs, as well as being able to add additional objects and relations as these differing OKBs are composed with each other. The triply articulated model produced by this process contains all the

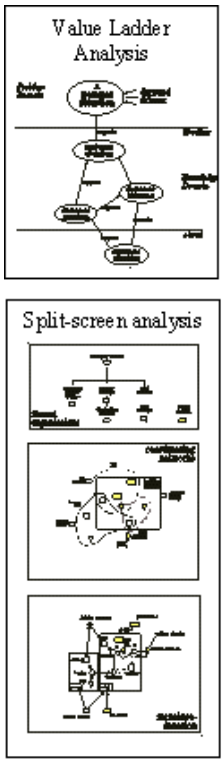

figure 3

knowledge from its constituent OKBs. Note that it is not the clinical enterprise itself that is being modelled but observers' models of the formation of the clinic.

The **composition** of a triple articulation may take a number of different forms, depending on which articulation is *privileged* and in which *sequence* the articulations are composed. The privileged articulation defines the domain in relation to which the composite articulation is formed, each of the objects that is not mapped to some object in this domain being deemed to be absent from the world denoted by the composite and pruned away. The resulting *composite knowledge base* (CKB) is defined in terms of objects that together comprise an ontology embedded in the particular semantic formation of a particular (singular or composite) actor-observer. It can be projected as a collection of matrices, arranged in strata that successively connect the worlds of supply and demand. The simplest such stratification, shown in fig. 4., arises because of the way the articulations are organised in relation to the referential, showing the particular relation of the enterprise to its pragmatic context. The lower strata represent products and services that are incorporated into the products and services of the higher-level strata, that finally arrive at their use within clients' contexts-of-use. Crucial in this analysis is the way in which the ontology of the lower strata is aligned to that of the upper strata by the particular form of relation to the context-of-use.

In the case of the clinic, the strata of this projection would contain:

6. The patient's context-of-use, describing the ultimate context in which a need arises;
5. The usage within an organisation of patient situations (the referral pathway);
4. The collaborative composition of a capabilities in the form of a system of systems, which can provide a composite service that satisfies the patient situation;
3. The capabilities of the healthcare system;
2. The sub-systems that themselves have to be composed to form the system capability
1. The components used by the sub-systems.

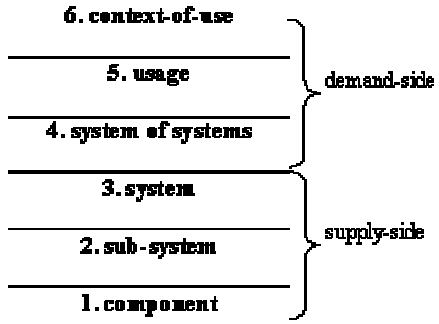

Figure 4

Although there is no limit to the number of tiers that can be identified in such an analysis, they can always be projected into six distinct strata because of the way they are defined in relation to the formulation of demand, three of them being demand-side, governed by the clinician's practice in relation to the patient's condition, and three supply-side, governed by the way the healthcare system supports the clinician.

Finally, regardless of the form of composition chosen, each stratum, considered as a *simplicial complex*, can be subjected to **landscape analysis**, an extension of Q-analysis [Atkin 1974], to identify the structural gaps within each stratum and display them graphically in a 3D histogram. These gaps indicate the absence of direct relationships between objects within each stratum, reflecting the particular ways in which they are unable to interoperate and, hence, the particular forms of risk facing the (clinical) enterprise as a whole.

1.7. The Clinician's Support Platform

The clinician governs the demand-side of the stratification in fig. 4 through the particular formation arising from her clinical assessment. At its heart are the clinician's judgements about what constitutes the patient's episode defining condition (EDC), which best characterises the patient's need for future treatment episodes (tier 6 above), and the treatment defining process (TDP), which is both a judgement about how to apply the treatment in the patient's case (tier 5), and a judgement about how that treatment should be constituted (tier 4). The clinic is formed through the way the clinician defines the TDPs in relation to patients' EDCs for the particular mix of patient conditions referred to the clinic (themselves the result of other clinicians' anticipations of how the clinic will provide care). The clinician must be able to manage the supply-side of the clinic in a way that is embedded in her own practice, using a repository of all the information that she needs to support her approach.

This 'support platform' (fig. 5) must not merely represent those aspects of the clinician's practice for which there are records (i.e. its processes) but must be organised in a way that is specific to the formation of the clinician's practice, enabling her to exercise 'data-pull' on the way the platform supports her, rather than being subject to 'data-push' based on someone else's (e.g. a standard) model of her practice. However, since some of the clinician's data originates from external sources whose organisation cannot be assumed to be consistent with her semantic formation, mappings must be defined that translate (*shred*) the formation of those sources to one consistent with that of the clinician's practice (see fig. 6).

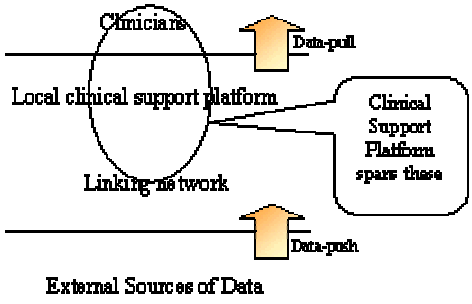

Figure 5

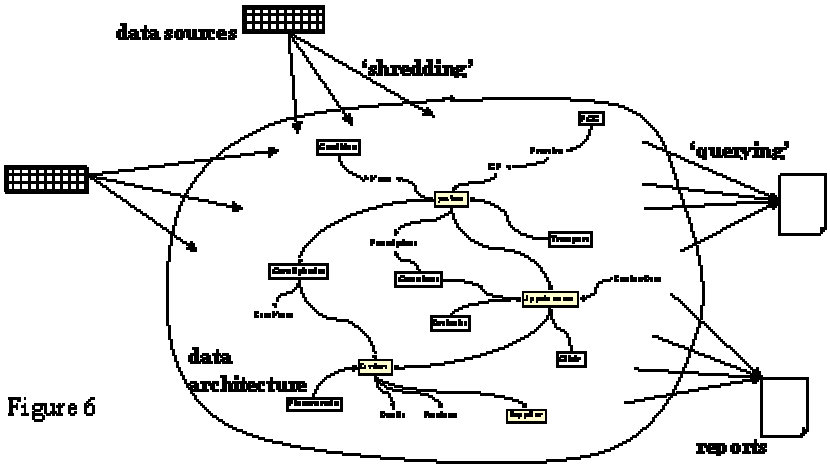

Figure 6

The ontology of the data elements within the platform itself (its data architecture) must be capable of supporting this mapping in order to satisfy the reporting requirements of the clinician (querying). It follows from the top-right quadrant in fig.1 that the way this platform reports cannot be standardised, even though the challenge in designing it must be to make use of standard data to the extent that it does not jeopardise its ability to support the clinician.

From this perspective, we infer that the generic approach to the electronic health record (EHR) is necessarily limited in its scope because it fails to provide the clinician with the means of managing the composition and implementation risks in relation to her patients. Projections of a health record can be computed on support platforms only in response to the practices of the clinicians it supports. The transfer of information between the support platforms of two different clinicians is therefore possible only if consistency can be constructed between the semantic formations in which their support platforms are embedded which, in turn, requires the composition of models of the semantic formation of the relevant clinicians' practices. This composition may reveal inconsistencies that must be rectified before any meaningful use of data from one platform may be used in support of the practice of the other, and a clinician may have to establish such compositional consistency with several others' practices. As a result, every strategic development by a clinician of her practice, and therefore of her support platform, may invalidate consistency established between her support platform and those of other clinicians. Equally, if the clinician wishes to retain previously established consistency, then she must accept restrictions on the strategic development of her practice. A generic EHR is an attempt to define an organisation of information with which all clinicians' support platforms are consistent. Clinicians within any healthcare enterprise committing to a generic EHR standard must therefore accept restrictions on the future strategic development of their practices. Clinicians are reluctant to accept such restrictions precisely because of the asymmetric nature of their patients' conditions and their need to form a care response that is particular to the needs of each patient. This reluctance is compounded when such restrictions are further intended to standardise their responses in the interests of reducing costs. PAN therefore provides the clinician with the means of deciding to what extent a standardised approach may be used and hence the extent to which she is able to orchestrate services from standard healthcare service components. On the other hand, given that the clinician has chosen to pursue a relational strategy in response to patients who insist that the asymmetric nature of their demand be addressed, it enables her to manage the risks of her practice in relation to her patients.

1.8. Conclusion

The Healthcare exemplar has demonstrated how a clinical support platform must necessarily support the particular semantic formation of the clinician in the way she composes treatments for her patients, and in relation to which she must orchestrate the supporting healthcare infrastructure so that it is aligned to the needs of her patients.

We suggest that in every open system, asymmetric demand will pose similar problems to the composition of services, whose solution will require the use of a tool, such as PAN, for the elicitation and composition of semantic formations. These semantic formations then provide the framework within which to manage the risks that arise in seeking to satisfy those demands.

Notes

PAN is the property of Boxer Research Ltd. It consists of a suite of software that:
1. parses the notation in which MKBs and UKBs are written and manages the articulations that they denote;
2. supplies the user with a graphical user interface in which the knowledge bases may be edited and appropriate composition operators invoked;
3. checks for, and reports violations of, internal consistency in the individual and composite knowledge bases so constructed;
4. outputs selected matrices from the stratifed compositions,
5. performs landscape analysis on selected matrices and displays the results as 3D surfaces.
The PAN applications that have been conducted so far are, by their nature, both elaborate and extremely sensitive. It is therefore impossible to display real results in the public domain. On the other hand, made-up examples small enough to include in these pages would appear so trivial as to defeat their own object. We therefore invite any readers who wish to learn more about PAN to contact Boxer Research Ltd. (www.brl.com) directly.

References

Alberts, D.S. & Hayes, R.E., 2003, *Power to the Edge. Command and Control Research Programme*, www.dodccrp.org.

Atkin, R.H., 1974, *Mathematical Structure in Human Affairs*, Heinemann (London).

Maier, Mark W., 2000, *Architecting Principles for Systems of Systems*, www.infoed.com/Open/PAPERS/systems.htm

Maturana, H. R. & Varela, F., V, 1980, *Autopoiesis and Cognition: The Realization of the Living* (Boston Studies in the Philosophy of Science, V. 42), D. Reidel

Peirce, C.S., 1958, *Collected Papers, Vol. 4* Harvard University Press

Porter, M.E., 1980, *Competitive Strategy: Techniques for Analyzing Industries and Competitors*. The Free Press (New York).

Rosen, R., 1991, *Life Itself*, Columbia University Press.

Rosen, R., 2000, *Drawing the Boundary between Subject and Object: Comments on the Mind-Brain Problem*, in *Essays on Life Itself*, Columbia University Press.

Sowa, J. & Dietz, D. 1999, Knowledge Representation: Logical, Philosophical, and Computational Foundations, Brooks Cole.