Service Supply and Asymmetric Demand

P. Boxer and B. Cohen
Boxer Research Ltd.

Abstract

People expect the behaviours of large, complex systems of systems to be embedded in their personal contexts-of-use. But suppliers of such systems often persuade their users to settle for something less than their particular needs because that something is better than nothing. Digitisation is turning this whole approach to designing complex systems of systems inside out. New approaches are required to understanding the nature of these complex systems which address risk, meaning and composition in ways that place the user of such systems at the centre of their formation instead of at their periphery. This paper outlines these approaches and the demands they place on the processes of collaborative composition.

Background

The Eureka E!2187 AgentWorks project was set up to establish the feasibility of developing engineering tools for the design and implementation of Multi-Agent Based Systems. The tool developed by Boxer Research Ltd formed part of this project – a *Composition Agent*. This tool enabled managers to model complex Enterprise Architectures and the changing requirements on them both for how business processes could be aligned to increasing varieties of demand, and for how the managers responsible for responding to those demands could be supported by information environments that enabled them to manage that alignment. At the heart of this project was a problem of composition:

> ‘ ... *the generic problem of mapping a user request with a set of services that can be composed to deliver a new service and satisfy the user demand*.’ [1]

There is continuing controversy over how this composition should be undertaken in relation to open environments[2]:

> *Technology development by itself is not enough to bring about the commonly articulated visions of vast information GRIDs populated by intelligent, dynamically interacting systems working on behalf of their users. In particular, this new generation of on-line systems must deal with the challenges raised by deployment in a truly open environment. In particular from our experience in AgentCities and other developments we believe there are four categories of long-term challenges:*
>
> - *Automation (Service Management):*
> - *Interoperability (Service-Service Communication):*
> - *Coordination (Service Orchestration):*
> - *Knowledge Acquisition (Interfaces between worlds).*

Our position was that the complete automation of the processes of orchestration and interfacing in a truly open environment was impossible. Thus ‘interfacing’ between the worlds of the user and of service suppliers could only be effective when the same models could be assumed for all actor-observers involved in the composition of services. In a truly open environment, this ‘closed-world’ assumption does not hold.

In what follows, therefore, we argue that the solution to this generic problem in a truly open environment necessarily involves the active involvement of the *actors* involved in such transactions, viz: the users whose demands those suppliers are seeking to satisfy. Thus, although it cannot be solved by purely computational mechanisms (e.g. *agents*), a certain kind of computational tool, exemplified by the *Composition Agent*, can effectively support the actors’ collaboration in the processes of orchestration and collaboration. This paper aims to explore what the Composition Agent has to be capable of as a result, and looks at an example in Health Care where the need for such a solution is apparent.

[1] Pierre, D. email of 7 February 2003.
[2] OpenNet ‘Objectives and Overview’ version 1.5, 2 February 2003.

Introduction

Asymmetric Demand for Services

The clients for services are *actors*: anticipatory systems [ROS1] deriving their demands for services from their formulation of themselves as context to their use of those services, referred to in what follows as a *context-of-use*. In this way we can think of an enterprise as an actor as much as we can an individual person.

If the service delivered to a client does not satisfy the client's demand when deployed in its context-of-use, then the client will experience a *value deficit*. This is the gap between what the client wants, and what is actually 'delivered' by the service. The greater the value deficit, the more *asymmetric* becomes the demand on the supplier with respect to the supplier's service. The formulation of a particular relation to an actor's demands takes the form of a *semantic formation*. This formation will include her experience of her own needs as the context-of-use, so that it will change 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 to some extent be asymmetric, leaving 'something to be desired' as a reflection of the fact that she will never know everything about her needs.

This paper assumes that a value deficit exists for enterprises as actors as much as it does for individuals, given that an enterprise is always ultimately formed by individuals. The normal presumption by a supplier, however, is of demand symmetry: that the nature of the client's demand is the same as that anticipated by the supplier. Such demand symmetry arises when the demand can be defined by both parties in a way that is independent of the client's context-of-use, so that it can be generalised across clients. Under these circumstances we normally refer to the client as a 'customer'.

The supply of services where demand is assumed to be symmetrical tends to become more competitively intense as more suppliers enter markets through the effects of globalisation. Equally, both the economies of scale through commoditisation and the economies of scope through providing for variety of supplied services tend to increase with 'digitisation'. Digitisation is here understood to be the effects of computational and networking infrastructures on the economics of suppliers' abilities to respond. Under both of these influences (globalisation and digitisation), demand asymmetry tends to increase precisely because of suppliers' success in satisfying symmetrical demand, allowing clients to expect more and more value for money with respect to the specificity of their demands in their contexts-of-use.

The doctrines of competitive strategy [POR] emerged around the elaboration of 'positional' competitive advantage. Those who adopted a positional approach (for example, by developing 'standard' services) sought to increase the value of their business without having to address the client's context-of-use. As value deficits grow, however, and the client becomes increasingly insistent on the asymmetric nature of their demand, so 'strategy' is under pressure to become 'relational', formulated more in terms of the particular client's context-of-use. Thus increasing asymmetry of demand increases the pressure on suppliers to switch from 'positional' to 'relational' approaches to demand. This requires a shift, from power being held at the centre in the form of a positional strategy for the enterprise as a whole, towards power being held at the edge of the enterprise, where there is relational knowledge of the client's context-of-use [ALB]. As a result the focus of the enterprise has to move towards managing the *risks* associated with deploying multiple strategies that are particular to the client relationship.

Collaborative Composition

With this shift to a relational strategy, it becomes necessary to separate out the organisation of the services from the organisation of demand. Thus in Figure 1 below, the nature of the services themselves, 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 itself, which introduces constraints relating to the client's particular context-of-use (the model of use-in-context). We can consider the alternative approaches to the problem of composing services, therefore, in terms of how these two are brought into relation with each other.

In the left column, there is a single model of how the service works. This model may be built for a particular use-in-context (bottom-left), or it may be parameterised so that it can be adapted to multiple forms of use (top-left). Either way, what the service actually does remains unchanged. On the right, however, there is no single model, but instead multiple components, each with its own model, and each one able to be composed with other models. Composition defines a System of Systems (SoS)[3]. In the single use-in-context, a designer can impose a *directed* integration 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 requirements of the single use-in-context. This directed approach is the one normally used for making large complex SoS projects work, and works through imposing composition externally.

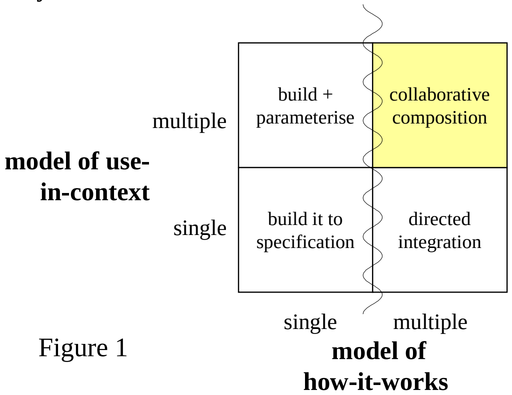

Figure 1

In all three instances so far, however, there is a presumption of a symmetrical relation to demand arising from the a priori construction of the service. This approach cannot be used top-right, whether because the composition needed to satisfy a particular use-in-context cannot be predicted by the designer, or because there is too great a variety of such uses to be able to predict which one a given user will want until the time of use. Either way, composition has to be a dynamic response to the user's demand. Here, then, the demand can be allowed to be asymmetric.

Under these circumstances, we need to be able to define a *granularity* and *stratification* of components that can support the variety of uses-in-context. This granularity is the particular way in which component services are defined, and the stratification is the way the different layers of composition of component services ultimately relate to the end use. We also need to adopt an approach to composition that is *collaborative,* it not being possible to impose composition externally, the component systems instead needing to be composed at the time of use in collaboration with the user's demand. Orchestration now involves determining which components get to participate in the collaborative process of composition. This 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 derived from the nature of the context-of-use in which the demand arises, and the more composition has to be collaborative. It is this collaborative process that the Composition Agent is designed to support, through establishing the suitable conditions of granularity and stratification for collaboration as well as defining the form of composition itself. Healthcare provides us with an example of asymmetric demand. So how does this demand for collaborative composition show itself there?

[3] An SoS can be understood to have the following characteristics [MAI]:
"1. Operational Independence of the Elements: If the system-of-systems is disassembled into its component systems the component systems must be able to usefully operate independently. The system-of-systems is composed of systems which are independent and useful in their own right.
2. Managerial Independence of the Elements: The component systems not only can operate independently, they do operate independently. The component systems are separately acquired and integrated but maintain a continuing operational existence independent of the system-of- systems.
3. Evolutionary Development: The system-of-systems does not appear fully formed. Its development and existence is evolutionary with functions and purposes added, removed, and modified with experience.
4. Emergent Behaviour: The system performs functions and carries out purposes that do not reside in any component system. These behaviours are emergent properties of the entire system-of-systems and cannot be localized to any component system. The principal purposes of the systems-of-systems are fulfilled by these behaviours.
5. Geographic Distribution: The geographic extent of the component systems is large. Large is a nebulous and relative concept as communication capabilities increase, but at a minimum it means that the components can readily exchange only information and not substantial quantities of mass or energy."

A Topical Exemplar: The Electronic Health Record (EHR)

The asymmetric nature of demand in health care is unavoidable. 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, making it difficult to establish universal measures of care performance.

Within this domain, the Electronic Health Record (EHR) is an attempt to provide a supporting infrastructure governed by a universal ontology specific to healthcare. The EHR is currently the subject of international standardisation activity that has been ongoing, in ISO, CEN, ASTMS and HL7, for over a decade. WHO, G8 and all national governments consider such standards to be significant factors in the improvement of public health. In considering the difficulties encountered by the implementation of this EHR, we can understand some of the particular challenges that give rise to the need for collaborative composition.

The standards bodies address these challenges by seeking "*to define generic characteristics of EHR information and to embody these in information models and message models that could be used by system developers as a standard interface between clinical systems*."[4] However, "*these standards have so far had limited uptake by industry*"[5] on the grounds that their information models are:

- "*incomplete, or have offered only partial coverage of the healthcare domain;*
- *unnecessarily complex;*
- *too generic, leaving the various implementations too much variability in how the models are applied in a given domain;*
- *flawed, with some classes and attributes not implementable as published;*
- *requiring expensive re-engineering of systems;*
- *containing features not required by the purchasers of clinical systems.*"[6]

Further, in those 'demonstrator' projects that chose to use a CEN EHR standard (ENV 13606) to enable interoperability between systems and enterprises, "*the adaptations made to ENV13606 have been rather ad hoc, so the exchange of EHR information between demonstrators has not been possible, thus defeating the object of such a standard*."[7]

Similar difficulties have plagued all other attempts to standardise the EHR, both *de re* and *de facto*. It is our contention that it is the failure to include the context-of-use as part of the way we understand the clinician's practice that results in these attempts being plagued in this way. But how are we to understand this challenge in terms of the clinician's practice?

The challenge of the clinician's practice

Standards work best when 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[8]. Consider Figure 2 below, using the same formulation as in Figure 1, but related to the Healthcare domain. Such situations appear in the bottom-left quadrant, whereas in the top-left quadrant differences in the way services are supplied necessitate the definition of distinct abstract domain models of ever-increasing complexity. 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.

Thus the progressive differentiation of care specialisms (the movement towards different 'supplier capabilities'), necessitate the development of *domain models* for the different specialisms as a result of the generic standards proving unable to deal with their differing specific requirements. And the challenge of asymmetric demand, represented here by the movement towards increasingly differentiated 'usage situations', arise 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

---

[4] Extracted from *Invitation to Contribute to work of the new CEN Task Force on Electronic Healthcare Record Communication*, CHIME, March 2002
[5] Op.cit.
[6] Op.cit.
[7] Op.cit.
[8] As in the exemplary standards for nuts and bolts, pioneered by Joseph Whitworth.

treatment strategies for these that can be made more widely accessible to patients, situation-based '*message models'* are used to relate them to the patient's context-of-use. But this still leaves us with the top-right quadrant, where the need for relational strategies, in response to the patient's 'value deficit' in the form of their particular needs for treatment, becomes overwhelming. Here the challenge of the particular nature of the patient's context-of-use becomes inescapable. But how significant is this quadrant?

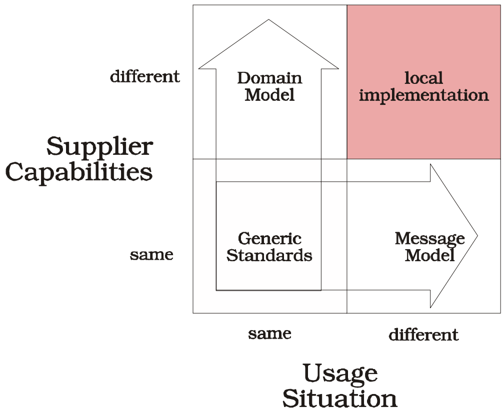

Figure 2

The top-right quadrant presents a severe challenge to the standards endeavour, as evidenced by the complaints above concerning the 'ad hoc' nature of local implementations of ENV 13606. The pragmatic argument is that the 80/20 principle can usually be applied, so that most of what is going on can be governed by standards, the rest being insignificant. The complaint about limited industrial uptake is symptomatic of the problem of asymmetric demand, combined with the sheer complexity and political sensitivity of the task of finding out which 80% can ignore asymmetric demand. This suggests that the top-right quadrant cannot be assumed to be insignificant, with its associated question of knowing *which* 80% to standardise! So how is the top-right quadrant to be approached?

Of course, the top-right quadrant defines the crucial nature of the role of the clinician in formulating the appropriate response to the particular patient *qua* context-of-use, 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;
- then, to identify which elements of this practice can be covered by standards without jeopardising the consistency of the whole.

From this there follows an alternative approach, which defines the clinician's practice within the context of which it is possible to negotiate what can safely be made generic, standardised and delegated within a universal ontology (the symmetric part), and what must remain particular to the demands of the patient's condition (the asymmetric part).

The Composition Agent

Risk and System Orders

In considering the implementation of a treatment strategy, we can distinguish three kinds of risk:

- *performance risk*: the possibility that a component product or service might not work as specified by the supplier.
- *composition risk*: the possibility that component, and hitherto independent, products or services might be not be able to interoperate in such a way as to provide an intended product or service as a whole.
- *implementation risk*: the possibility that the intended product or service might not satisfy the client's demand as expected when used in its actual context-of-use, even though it meets the product or service specification agreed with the supplier.

The first two of these risks are faced by any form of supplier in developing products and services, but the third is particular to the relational response to asymmetric demand (the positional response being to ignore the value deficit associated with it). We are particularly interested, therefore, in understanding what makes this third kind of risk distinctive, and how it affects the nature of the second risk, since it is this risk that most distinguishes the clinician's practice.

Performance risk involves exposure to *errors of execution*, when the supplier is unable to sustain the performance of its constituent capabilities. It is normally evaluated by behavioural analysis of models of those capabilities, using simulation, analytical techniques or exhaustive testing. These forms of analysis are typical of, and are computationally effective for, those capabilities describable in classical systems theory, i.e. **first-order systems**.

Composition risk involves exposure to *errors of planning*, when the supplier is unable to ensure the validity of its approach to composing the product or service as a whole 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 as a whole. It is this ability that distinguishes agent-based architectures from their simpler, first order, software counterparts. This self-organising, or *autopoietic* [MAT] property is characteristic of **second-order systems**.

These systems pose analytical problems related to those that arise in higher-order logics (i.e. those logics involving statements in which quantifiers range over predicates), and 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 [SOW]. Unfortunately, that assumption is explicitly denied, by definition, in *open systems.* Approximate solutions to second-order analytical problems may nevertheless be found by assuming that a single ontology can cover a supplier's specific domains of activity, thereby, in effect, restoring the closed world assumption[9]. This, however, prevents us from being able to deal with asymmetric demand.

Implementation risk involves exposure to *errors of intention*, when the supplier is unable to guarantee that products or services will satisfy the client's situation within its context-of-use. Thus the more asymmetric the demand, the more likely it is that the supplier will encounter implementation risk. The point here is that the client is an embodied actor, possessing a model of her own organisation of demand in relation to which her own and others' behaviour is oriented, through specifying services that might satisfy her demands and anticipating, possibly erroneously, their satisfaction by those services. Thus the client's model of demand has itself to be included in the process of service composition. This apparently teleogical[10] (or *anticipatory* [ROS1S]) property is characteristic of **third-order systems**, which exhibit what Rosen referred to as 'closed loops of entailment'. Again, we have no computationally effective analytical procedures for such systems[11].

[9] Taken to the extreme, if the fabled 'Standard Universal Ontology' [CYC] were to be constructed, this approach could be made to work in general. But there are not only strong philosophical grounds for doubting its existence, but also strong practical ones evidenced, for example, in the difficulties with Electronic Health Records that we discuss below. It is these difficulties that are addressed through the concept of *implementation risk*.
[10] As in Aristotle's 'final cause', not to be confused with the more simplistic 'goal-driven' systems.
[11] In fact, very little is known about the mathematical structure of anticipatory system models. Rosen hinted that they might not form a cartesian closed category and would therefore be neither expressible in a logic nor computationally tractable, but that is mere speculation. "It should be stressed that, by advocating the objectivity of

The logical problem these third-order systems pose is related to that of composing statements made in different modalities, which was identified by C. S. Peirce, in relation to his *Gamma graphs* [PEI], as an outstanding difficulty that future generations would have to tackle. Most knowledge representation languages are based on Sowa's Conceptual Graphs which, although they derive directly from Peirce's graphs, do not solve Peirce's problem of Gamma graphs. The Composition Agent tackles this problem through using *a triple articulation.*

Triple Articulation

We can understand the nature of the three risks in terms of a triple articulation of the clinician's practice. Triply articulated modelling furnishes three different types of model that are inherent to the way the clinician *qua* actor-observer constructs her world:

- 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 the patient *qua* context-of-use in which those situations are embedded.

These models are distinct, so that the clinician's practice is formed through the particular way she brings them into relation with each other. In doing this, she forms a composite articulation that is particular to the patient's context-of-use. The role of the Composition Agent is to support this process of composition by checking for consistency and identifying the gaps in the composite articulation – gaps representing a lack of inter-operability between the constituent parts of the articulation.

Thus the actor-observer is here being 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 itself. The triple articulation is in this sense a model of the actor-observer's model of her own formation – in this case the clinician's formation of her practice. In these terms, performance risk arises from errors of execution described in terms of the existential articulation, while composition risk arises from errors of planning in the relation between the deontic and existential articulations. Implementation risk, however, arises from errors of intention in the relation between deontic and existential and referential articulations.

From this is is apparent that performance risk can be evaluated by objective modelling of the existential articulation, and therefore benefits little from the Composition Agent. 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. The purpose of the Composition Agent is therefore to enable 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.

The Semantic Formation of the Actor-Observer

The three risks and their associated errors are characteristics of particular compositions of the actor-observer's triple articulation, which as a whole describe the semantic formation of the actor-observer. At the lowest level of this semantic formation, we find the *lexical formation* of 'data' – the way the atomic terms used by an actor-observer are formed. At the next level we have their *syntactic formation*, involving constraints on the ways in which the actor composes atomic terms into well-formed statements, i.e. the *syntagmatic structure* of the actor-observer's language. At the next *operational* level, the meaning of well-formed statements is defined. When these are predicates or imperatives, i.e. when they denote queries about, or changes to, the state of some system (as in

---

complex systems, systems with nonformalizable models and hence closed loops of entailment (impredicativities), I am advocating the objectivity of at least a limited kind of final causation. This is precisely what closes the causal loops. It simply describes something in terms of what it entails, rather than in terms of what entails it. This, it will be observed, need have nothing whatever to do with Telos, any more than, say Godel's incompleteness Theorem does." (ROS2, p95) programming languages), their meaning may be defined by *denotational semantics*[12], the state space of the system being a *semantic domain*. But can the semantic formations of actor-observers be defined denotationally in relation to such semantic domains?

For a supplier-actor, an appropriate semantic domain would denote the functional components, together with their combinations and relations, which the supplier-actor may dispose for a client. This is what is articulated in the relation between the existential and deontic articulations, which characterise composition risk. But for an actor-client, the semantic domain would have to denote the *pragmatics* of the client's context-of-use, as well as the characteristics of anything that she might use in support of those pragmatics. These pragmatics form a fourth level of semantic formation. The existence of asymmetric demand implies therefore that the semantic formation of an actor must include the *pragmatics* associated with her context-of-use.[13] It is the inclusion of this fourth pragmatic level that leads us to formulate a triple articulation through the inclusion of the referential articulation.

If such a (four-level) semantic formation is internally consistent, we can ask two further questions of it: firstly, is it sufficiently *expressive* of its domain; and secondly, is it *consistent with* the semantic formation of some other actor? It is the role of the Composition Agent to support this process of establishing consistency within and between semantic formations that include pragmatic aspects. To ensure expressiveness, this process has to be undertaken in collaboration with the appropriate actors. Given adequate expressiveness, implementation risk can then be detected as inconsistency between an actor-supplier's and the actor-client's semantic formation.

But what happens without such a composition agent? The actor-supplier starts with a specification of what he thinks the user wants and formalises it as a denotational semantics which is (a) adequate to his articulation of the user's requirement, (b) consistent with the functional dynamics of the semantic domain in which the denotational semantics is to be operationalised (i.e. of which it is adequately expressive), and (c) over-determining of the user, by limiting the user's pragmatics to what the designer has articulated [FLO].

Again we see how this avoids the top-right quadrant in Figures 1 and 2. The over-determination derives from the actor-supplier's assumption of a (locally) universal ontology that obviates the need to explicate the particular semantic formation that it supports, and any inconsistencies that it might have with that of the actor-client. As a result, we lose the ability to discuss the expressibility 'gap' between the actor-supplier's semantic formation and its domain and the semantic formation and domain of the actor-client, and with it any ability to discuss demand asymmetry.[14]

---

[12] cf Frege's "denotational imperative" that the meaning of any statement be definable only in terms of the meaning of its parts.

[13] In the world of the military and the new doctrines of network-centric warfare, this distinction is present in the separation of data fusion (the formation of an integrated picture from composing the outputs of multiple sensor systems) and situational awareness, which involves understanding how it is possible to compose the different fusions of data given the differing pragmatics of the different commanders *qua* interpreters.

[14] The extent to which this 'expressibility gap' has not been acknowledged within FIPA has limited its ability to deliver on its promise, and constrained it to being not much more than a dialect of CORBA.

Supporting the Clinician

Using the Composition Agent

Interaction with the Composition Agent[15] may be based on any method of relational modelling. The examples in figure 3 describe (working downwards) (i) the organisation of demand, referred to here as a value ladder, but for a patient taking the form of a referral pathway; (ii) organisational hierarchy; (iii) horizontal relationships across the organisation; and (iv) the processes and infrastructures of the clinic. Whatever form these modelling methods take, the method of relational modelling within which they are constructed is itself expressed as a meta-knowledge base (MKB) that also defines the relation of the relational modelling to the triple articulation. Thus the MKB enables a model formulated by an actor-observer to be 'read' as an observer's knowledge base (OKB).[16]

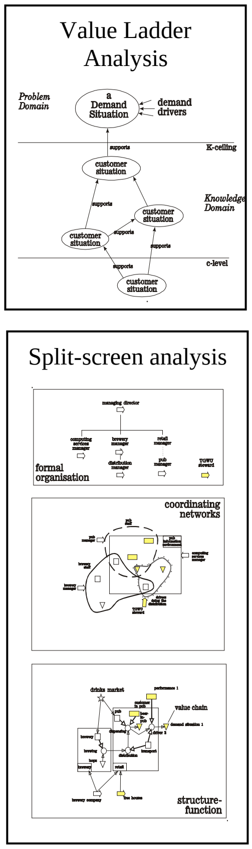

figure 3

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 OKB, but also those of other administrators and clinicians, as well as that of the patient herself. Thus the Composition Agent 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 then becomes the minimal graph that contains all the knowledge from its constituent OKBs.

What is being modelled here is not the clinical enterprise itself, so much as *observers' models* of the formation of the clinic. Thus, the *composition* of a triple articulation takes a number of forms, depending on two things:

- which articulation is to be *privileged*, and
- in which *sequence* should the articulations be composed.

In the resultant composition, the privileged articulation defines the domain in relation to which the composite articulation is formed, with 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.[17] The resultant composite knowledge base (CKB) is internally consistent, defined in terms of objects that together comprise an 'ontology' embedded in the particular semantic formation of a particular (singular or composite) actor-observer[18].

This CKB can itself be projected in the form of a number of matrices that are *stratified* in relation to each other, the simplest stratification of which is shown in Figure 4 below. Its stratification arises because of the way the articulations are organised in relation to the referential, showing the particular relation of the clinical enterprise to its pragmatic context. Thus 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 patients' contexts-of-use. Crucial in this analysis is the way in which the 'ontology'

---

[15] BRL's implementation of the Composition Agent is called PAN. PAN and its associated tools were developed as part of the Eureka AgentWorks Project E!2187.

[16] The visual modelling techniques shown here are particularly designed to distinguish components of the referential, existential and deontic articulations. Other modelling techniques, however, may identify objects that cannot be directly allocated in this way. The relational structures implicit in these objects must be unpacked (a process that requires the cooperation of the analyst and the actor-observer in defining the relevant MKB operators) into object types, the objects assigned to those types, and relations between objects within and between the articulations. The agreed unpacking is represented in the MKB by suitably defined complex object operators (written in PAN), along with rules for checking that the populated models are syntactically valid. The extended OKB that is generated by the application of these complex object operators forms the actor-observer's input to the Composition Agent.

[17] There are in fact six different forms of composition, each one reflecting a different relationship of the actor-observer to her semantic formation.

[18] In contrast to those knowledge representation approaches that adopt the 'closed world assumption' of a Universal Ontology.

of the lower level strata is aligned to the upper level strata by the particular form of relation to the context-of-use.

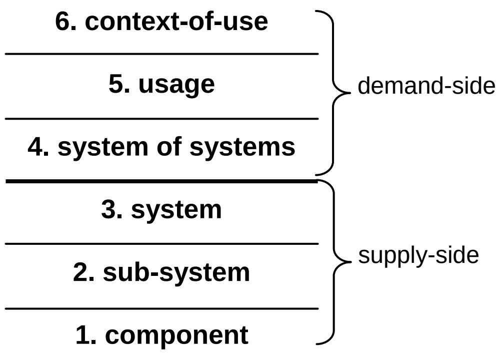

Figure 4

The contents of the strata within this projection can be summarised as follows:

6. The user'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 number of 'system' capabilities in the form of a system of systems, which can provide a composite service that satisfies the patient situation
3. The (healthcare) 'system' capabilities themselves.
2. The sub-systems that themselves have to be composed to form the 'system' capability
1. The components used by the sub-systems.

There is of course no limit to how many tiers can be separated out within each of these strata, but they can nevertheless 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 projection chosen, each stratum can be subjected to '*landscape analysis*'. This is a particular form of analysis which identifies the structural gaps within each stratum. These gaps identify the absence of direct relationships between objects within each stratum, reflecting the particular ways in which they are unable to inter-operate, and as a result the particular forms of risk facing the (clinical) enterprise as a whole. So how to we apply this to the practice of the clinician?

Supporting the Clinician

A clinician is a good example of an actor-observer. Even though the organisation of demand is not her own, her ethics require her to articulate that of her patient in a way that is particular to the patient – that addresses the asymmetric nature of the patient's condition in such a way that the patient can be appropriately treated. It is this ethic which enables the clinician to 'govern' the demand-side of the stratification in Figure 4 through the particular formation arising from her clinical assessment.

At the heart of this clinical assessment are two things. Firstly, there is the clinician's judgement about what constitutes the patient's *episode defining condition* (EDC). This is that condition, with its implied treatment, which best characterises the patient's need for future treatment episodes, and arises in tier 6 above. Secondly, there is the clinician's judgement about the particular form of treatment(s) that are appropriate to the patient's condition: the *treatment defining process* (TDP). This is both a judgement about how to apply the treatment in the patient's case (tier 5 above), and a judgement about how that treatment should be constituted (tier 4 above). Thus the clinician's clinic is formed through the way she defines these TDPs in relation to patients' EDCs for the particular mix of patient conditions referred to the clinic, these referrals being themselves the result of other clinicians' anticipations of how the clinic will provide care.

As a result, the 'demand-side' treatment of the patient leaves open the question of how the 'supply-side' of the clinic provides support to the clinician's particular practice. To do so it has not only to enable the clinician to keep track of her patients through holding definitions and relations that were particular to the formation of her treatments of them: i.e. patients, treatments, patient health records, EDCs, and referral pathways. It also has to enable her to track the supply-side use of facilities, products and services in support of those treatments.

In order to be able to manage the alignment of the supply- and demand-sides, the clinician needs information about the relation of each to the other. We call the ability to provide this information a clinician's 'support platform', because it enables her to manage the supply-side in a way that is embedded in her own practice, acting as a repository of all the information that she needs to support her approach to:

- comprehending her clients' condition, expressed as an EDC;
- coordinating the clinic's own processes and those of its suppliers, organised in relation to TDPs; and
- synchronising the outcomes of the constituent processes of those TDPs to satisfy the demands of the former EDCs.

Without this it is not possible either for the clinician or anyone else to hold the clinician accountable for the outputs of her clinic other than in terms which do not account for the asymmetric nature of the demands on her clinic.

Thus the clinician's support platform in Figure 5 below is not only itself a realisation of those aspects of the clinician's practice for which there are records – namely its processes – but is also organised in a way that is specific to the formation of the clinician's practice.[19] This enables the clinician to exercise 'data-pull' on the way the platform supports her, rather than being subject to 'data-push' based on someone else's model of her practice – typically one based on standards.

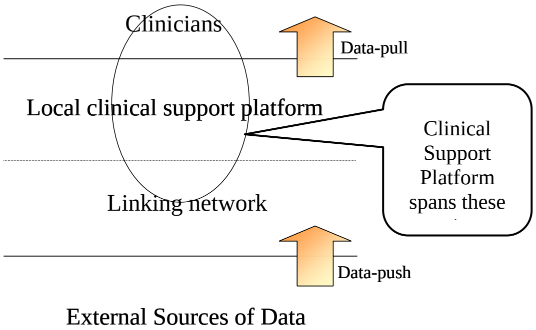

Figure 5

If we think of this 'data push' being from external as well as from local data sources, then the form taken by the 'data pull' must reflect the particular semantic formation of the clinician's practice. In this sense, the 'data pull' is being exerted from the top right quadrant of figures 1 and 2, and the way these data sources are organised cannot be assumed to be consistent with the semantic formation of the clinician. As a result, mappings must be defined ('shredding') that translate the formation of the data sources to one consistent with that of the clinician's practice (see Figure 6 below).

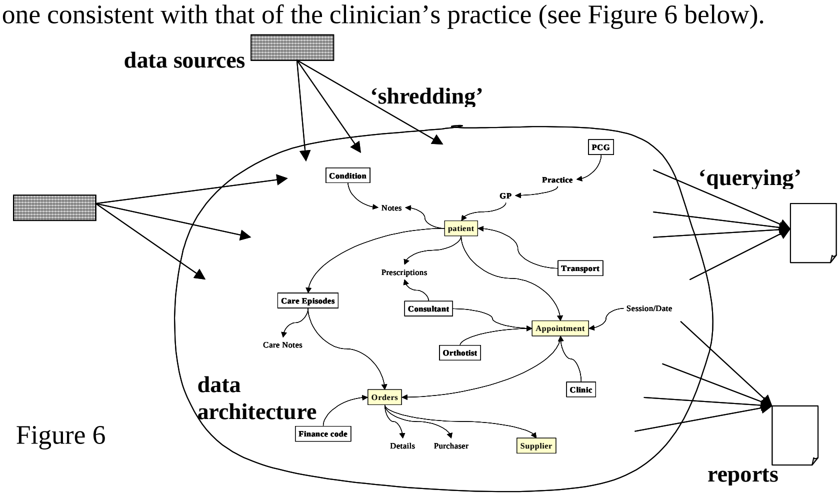

Figure 6

---

[19] Thus the knowledge support platform cannot itself be standardised. The data architecture and shredding processes must be configurable in support of the clinician's semantic formation, and the ontology of the data architecture must be consistent with that of both the supplying sources of data and that of the clinician. Once

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 Figure 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. This is the question of how much of the clinician's practice falls in the top-right quadrant in Figure 1.

The limitations of generic EHR

From this perspective, we deduce that the *generic* approach to the 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. Our reasoning is that projections of a 'health record' can only be computed on support platforms 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. This construction of consistency 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, which need not necessarily be compatible with each other.

As a result, every strategic development by a clinician of their practice, and therefore of their support platform, may invalidate consistency established between their support platform and those of other clinicians, and equally, if the clinician wishes to retain previously established consistency, then she must accept restrictions on the strategic development of her practice. In these terms, a *generic* EHR is an attempt to define an organisation of information with which all clinicians' support platforms are consistent, as a result of which the clinicians within any healthcare enterprise committing to a generic EHR standard must accept restrictions on the future strategic development of their practices.

Clinicians are reluctant to accept such restrictions on their practice, 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 – a reluctance compounded when such restrictions are further intended to standardise their responses in the interests of reducing costs. The Composition Agent therefore provides the clinician with the means of deciding on the one hand to what extent a standardised approach may be used.[20] This will determine the extent to which the clinician 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.

Conclusion

This 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 the Composition Agent, 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.

---

defined in these terms, however, it can then function as an organising overlay on the healthcare information environment, determining how data is mapped and composed.

[20] These issues can be seen in the NHS, where plans for data fusion (see footnote 8) around the Integrated Care Records Service [IPU] have not yet been distinguished from the requirements of the clinician's 'situational awareness' – working out how to interpret both records and the evidence presented directly by the patient in formulating both the patient's EDC and an appropriate TDP. It is worth noting that, in the case of the military, success in providing situational awareness in support of a commander's Command and Control has increased the demand from the commander for self-synchronisation – the commander's scope for local initiative in relation to asymmetric forms of threat within the context of an overall intent.

References

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

[CYC] http://www.cyc.com/cyc-2-1/cover.html

[FLO] Floyd, C. "Software Development as Reality Construction"

[IPU] "Consultation Draft, National Specification for Integrated Care Records Service", in *Delivering 21st Century Support for the NHS*, Information Policy Unit, Department of Health, July 2002, http://www.doh.gov.uk/ipu/develop/index.htm

[MAI] from "Architecting Principles for Systems of Systems", by Mark W. Maier, U. Alabama. http://www.infoed.com/Open/PAPERS/systems.htm

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

[OPE] OpenNet 'Objectives and Overview' version 1.5, 2 February 2003.

[PEI] Peirce, C.S., 1931-1958, *Collected Papers*, Vol. 4 (eds. P. Weiss, A. Burks, C. Hartshorne), Cambridge: Harvard University Press (CP).

[PIE] Pierre, D. email of 7 February 2003.

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

[ROS1] Rosen, R. *Life Itself*, Columbia University Press, 1991

[ROS2] Rosen, R. "Drawing the Boundary between Subject and Object: Comments on the Mind-Brain Problem", in *Essays on Life Itself*, Columbia University Press, 2000.

[SOW] Sowa, J. and Dietz, D. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*, Brooks Cole, 1999.