

AGENT-BASED COMPUTING

Nicholas R. Jennings

Dept. of Electronics and Computer Science,

University of Southampton, Southampton SO17 1BJ, UK.

`nrj@ecs.soton.ac.uk`

Abstract Agent-based computing represents an exciting new synthesis for both Artificial Intelligence and, more generally, Computer Science. It has the potential to improve the theory and the practice of modelling, designing and implementing complex computer systems. Yet, to date, there has been little systematic analysis of what makes the agent-based approach such an appealing and powerful computational model. To rectify this situation, this paper aims to tackle exactly this issue. The standpoint of this analysis is the role of agent-based software in solving complex, real-world problems. In particular, it will be argued that the development of robust and scalable software systems requires autonomous agents that can complete their objectives while situated in a dynamic and uncertain environment, that can engage in rich, high-level interactions, and that can operate within flexible organisational structures.

Keywords: autonomous agents, agent-oriented software engineering, complex systems

Introduction

Building high quality, industrial-strength software is difficult. Indeed, it has been argued that developing such software in domains like telecommunications, industrial control and business process management represents one of the most complex construction tasks humans undertake (both in terms of the number and the flexibility of the constituent components and in terms of their interconnections). Against this background, a wide range of software engineering paradigms have been devised. Each successive development either claims to make the engineering process easier or to extend the complexity of applications that can feasibly be built. Although evidence is emerging to support these claims, researchers continue to strive for more effective techniques. To this end, this paper will argue that analysing, designing and implementing complex software

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35602-0_35](https://doi.org/10.1007/978-0-387-35602-0_35)

systems as a collection of interacting, autonomous agents (i.e., as a *multi-agent system*) affords software engineers a number of significant advantages over contemporary methods. This is not to say that *agent-oriented software engineering* [Jennings 2000] represents a silver bullet [Brooks 1995] – there is no evidence to suggest it will represent an order of magnitude improvement in productivity. However, the increasing number of deployed applications [Jennings and Wooldridge 1998, Parunak 1999] bears testament to the potential advantages that accrue from such an approach.

In seeking to demonstrate the efficacy of agent-oriented techniques, the most compelling argument would be to quantitatively show how their adoption had improved the development process in a range of projects. However, such data is simply not available (as it is not for other contemporary software engineering approaches like patterns, application frameworks and component-ware). Given this fact, the best that can be achieved is a qualitative justification of why agent-oriented approaches are well suited to engineering complex, distributed software systems.

1. MANAGING COMPLEXITY IN SOFTWARE SYSTEMS

Industrial-strength software is complex: it has a large number of parts that have many interactions [Simon 1996]. Moreover this complexity is not accidental [Brooks 1995], it is an innate property of such systems. Given this situation, the role of software engineering is to provide structures and techniques that make it easier to handle complexity. Fortunately for designers, this complexity exhibits a number of important regularities [Simon 1996]:

- Complexity frequently takes the form of a hierarchy. That is, a system that is composed of inter-related sub-systems, each of which is in turn hierarchic in structure, until the lowest level of elementary sub-system is reached. The precise nature of these organisational relationships varies between sub-systems, however some generic forms (such as client-server, peer, team, etc.) can be identified. These relationships are not static: they often vary over time.
- The choice of which components in the system are primitive is relatively arbitrary and is defined by the observer's aims and objectives.
- Hierarchic systems evolve more quickly than non-hierarchic ones of comparable size (i.e., complex systems will evolve from simple

systems more rapidly if there are clearly identifiable *stable intermediate forms*, than if there are not).

- It is possible to distinguish between the interactions *among* sub-systems and those *within* sub-systems. The latter are both more frequent (typically at least an order of magnitude more) and more predictable than the former. This gives rise to the view that complex systems are *nearly decomposable*: sub-systems can be treated almost as if they are independent, but not quite since there are some interactions between them. Moreover, although many of these interactions can be predicted at design time, some cannot.

Drawing these insights together, it is possible to define a canonical view of a complex system (figure 1). The system's hierarchical nature is expressed through the "related to" links, components within a sub-system are connected through "frequent interaction" links, and interactions between components are expressed through "infrequent interaction" links.

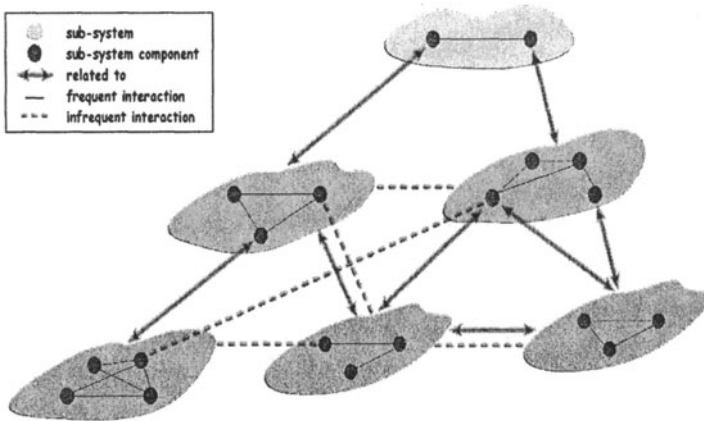


Figure 1. View of a Canonical Complex System

Given these observations, software engineers have devised a number of fundamental tools of the trade for helping to manage complexity [Booch 1994]:

- *Decomposition*: The most basic technique for tackling large problems is to divide them into smaller, more manageable chunks each of which can then be dealt with in relative isolation (note the nearly decomposable sub-systems in figure 1). Decomposition helps tackle complexity because it limits the designer's scope.

- *Abstraction*: The process of defining a simplified model of the system that emphasises some of the details or properties, while suppressing others. Again, this works because it limits the designer's scope of interest at a given time.
- *Organisation*¹: The process of defining and managing the inter-relationships between the various problem solving components (note the sub-system and interaction links of figure 1). The ability to specify and enact organisational relationships helps designers tackle complexity by: (i) enabling a number of basic components to be grouped together and treated as a higher-level unit of analysis and (ii) providing a means of describing the high-level relationships between various units.

2. THE CASE FOR AGENT-ORIENTED SOFTWARE ENGINEERING

The first step in arguing for an agent-oriented approach to software engineering involves identifying the key concepts of agent-based computing. The first such concept is that of an agent:

an agent is an encapsulated computer system that is situated in some environment, and that is capable of flexible, autonomous action in that environment in order to meet its design objectives [Wooldridge 1997]

There are a number of points about this definition that require elaboration. Agents are: (i) clearly identifiable problem solving entities with well-defined boundaries and interfaces; (ii) situated (embedded) in a particular environment over which they have partial control and observability – they receive inputs related to the state of their environment through sensors and they act on the environment through effectors; (iii) designed to fulfil a specific role – they have particular objectives to achieve; (iv) autonomous – they have control both over their internal state and over their own behaviour; (v) capable of exhibiting flexible problem solving behaviour in pursuit of their design objectives – being both reactive (able to respond in a timely fashion to changes that occur in their environment) and proactive (able to opportunistically adopt goals and take the initiative) [Wooldridge and Jennings 1995].

When adopting an agent-oriented view, it soon becomes apparent that most problems require or involve multiple agents: to represent the decentralised nature of the problem, the multiple loci of control, the multiple perspectives or the competing interests. Moreover, the agents will need to interact with one another: either to achieve their individual objectives or to manage the dependencies that ensue from being situ-

ated in a common environment. These interactions can vary from simple semantic interoperation (information passing), through traditional client-server type interactions, to rich social interactions (the ability to cooperate, coordinate and negotiate about a course of action). Whatever the nature of the social process, however, there are two points that qualitatively differentiate agent interactions from those that occur in other software engineering paradigms. Firstly, agent-oriented interactions generally occur through a high-level (declarative) agent communication language (often based on speech act theory [Mayfield et al. 1995]). Consequently, interactions are conducted at the *knowledge-level* [Newell 1982]: in terms of which goals should be followed, at what time and by whom (cf. method invocation or function calls that operate at a purely syntactic level). Secondly, as agents are flexible problem solvers, operating in an environment over which they have only partial control and observability, interactions need to be handled in a similarly flexible manner. Thus, agents need the computational apparatus to make context-dependent decisions about the nature and scope of their interactions and to initiate (and respond to) interactions that were not foreseen at design time.

In the majority of cases, agents act either on behalf of individuals/companies or as part of some wider initiative. Thus, there is typically some underpinning organisational context to agents' interactions. This context defines the nature of the relationship between the agents. For example, they may be peers working together in a team or one may be the manager of the others. To capture such links, agent systems have explicit constructs for modelling organisational relationships (e.g. manager, team member). In many cases, these relationships are subject to ongoing change: social interaction means existing relationships evolve (e.g. a team of peers may decide to move to a more hierarchical structure as their numbers grow) and new relations are created (e.g. a number of unrelated agents may decide to form a team to deliver a service that no one individual can offer). The temporal extent of these relationships can also vary enormously: from providing a service as a one-off, to a permanent bond. To cope with this variety and dynamic, agent researchers have devised protocols that enable organisational groupings to be formed and disbanded, specified mechanisms to ensure groupings act together in a coherent fashion and developed structures to characterise the macro behaviour of collectives [Jennings and Wooldridge 1998, Wooldridge and Jennings 1995].

Drawing these points together (figure 2), it can be seen that: (i) adopting an agent-oriented approach to software engineering means decomposing the problem into multiple, autonomous components that can

act and interact in flexible ways to achieve their set objectives; (ii) the key abstraction models that define the agent-oriented mindset are agents, interactions and organisations; and (iii) explicit structures and mechanisms are often used to describe and manage the complex and changing web of organisational relationships that exist between the agents.

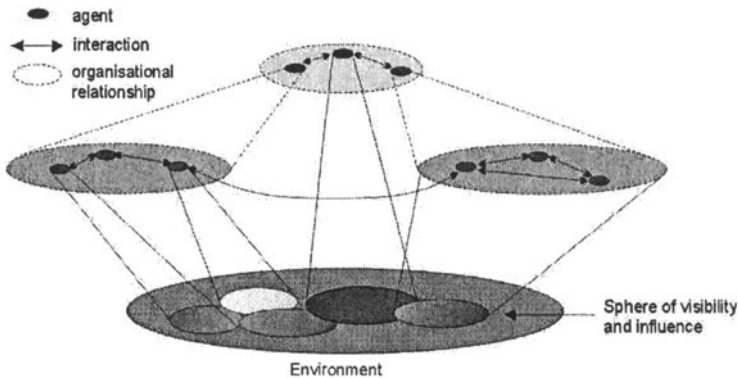


Figure 2. Canonical View of a Multi-agent System

2.1. The Software Engineering Credentials of the Agent-Oriented Approach

Here the argument in favour of an agent-oriented approach to software engineering is composed of the following steps:

- show that agent-oriented decompositions are an effective way of partitioning the problem space of a complex system;
- show that the key abstractions of the agent-oriented mindset are a natural means of modelling complex systems; and
- show that the agent-oriented philosophy for modelling and managing organisational relationships is appropriate for dealing with the dependencies and interactions that exist in complex systems.

When taken together, these steps form a complete mapping between the characteristics of a complex system and the key software engineering abstractions for handling complexity as they apply to agent-based systems.

The Merits of Agent-Oriented Decompositions.

Complex systems consist of a number of related sub-systems organised in a hierarchical fashion (figure 1). At any given level, sub-systems

work together to achieve the functionality of their parent system. Moreover, within a sub-system, the constituent components work together to deliver the overall functionality. Thus, the same basic model of interacting components, working together to achieve particular objectives occurs throughout the system. Given this fact, it is entirely natural to modularise the components in terms of the objectives they achieve². In other words, each component can be thought of as achieving one or more objectives. A second important observation is that complex systems have multiple loci of control: “real systems have no top” [Meyer 1988] pg 43. Applying this philosophy to objective-achieving decompositions means the individual components should localise and encapsulate their own control. Thus, entities should have their own thread of control (i.e. they should be active) and they should have control over their own actions (i.e. they should be autonomous).

For the active and autonomous components to fulfil both their individual and collective objectives, they need to interact (recall complex systems are only nearly decomposable). However the system’s inherent complexity means it is impossible to *a priori* know about all potential links: interactions will occur at unpredictable times, for unpredictable reasons, between unpredictable components. For this reason, it is futile to try and predict or analyse all the possibilities at design-time. It is more realistic to endow the components with the ability to make decisions about the nature and scope of their interactions at run-time. From this, it follows that components need the ability to initiate (and respond to) interactions in a flexible manner.

The policy of deferring to run-time decisions about component interactions facilitates the engineering of complex systems in two ways. Firstly, problems associated with the coupling of components are significantly reduced (by dealing with them in a flexible and declarative manner). Components are specifically designed to deal with unanticipated requests and can spontaneously generate requests for assistance if they find themselves in difficulty. Moreover because these interactions are enacted through a high-level agent communication language, coupling becomes a knowledge-level issue. At a stroke this removes syntactic concerns from the types of errors caused by unexpected interactions. Secondly, the problem of managing control relationships between the software components (a task that bedevils traditional objective-based decompositions) is significantly reduced. All agents are continuously active and any coordination or synchronisation that is required is handled bottom-up through inter-agent interaction.

From this discussion, it is apparent that the natural way to modularise a complex system is in terms of multiple autonomous components

that can act and interact in flexible ways in order to achieve their set objectives. Given this, the agent-oriented approach is simply the best fit to this ideal.

The Suitability of Agent-Oriented Abstractions.

A significant part of the design process is finding the right models for viewing the problem. In general, there will be multiple candidates and the difficult task is picking the most appropriate one. When designing software, the most powerful abstractions are those that minimise the semantic gap between the units of analysis that are intuitively used to conceptualise the problem and the constructs present in the solution paradigm. In the case of complex systems, the problem to be characterised consists of sub-systems, sub-system components, interactions and organisational relationships. Taking each in turn:

- Sub-systems naturally correspond to agent organisations. They involve a number of constituent components that act and interact according to their role within the larger enterprise.
- The case for viewing sub-system components as agents has been made above.
- The interplay between the sub-systems and between their constituent components is most naturally viewed in terms of high-level social interactions: “in a complex system... at any given level of abstraction, we find meaningful collections of objects that collaborate to achieve some higher level view” [Booch 1994] pg 34. This view accords precisely with the knowledge-level treatment of interaction afforded by the agent-oriented approach. Agent systems are invariably described in terms of “cooperating to achieve common objectives”, “coordinating their actions” or “negotiating to resolve conflicts”.
- Complex systems involve changing webs of relationships between their various components. They also require collections of components to be treated as a single conceptual unit when viewed from a different level of abstraction. Here again the agent-oriented mindset provides suitable abstractions. A rich set of structures are available for explicitly representing organisational relationships. Interaction protocols exist for forming new groupings and disbanding unwanted ones. Finally, structures are available for modelling collectives. The latter point is especially useful in relation to representing sub-systems since they are nothing more than a team of components working together to achieve a collective goal.

Flexible Organisational Structures.

Organisational constructs are first-class entities in agent systems – explicit representations are made of organisational relationships and structures. Moreover, agent-oriented systems have the concomitant computational mechanisms for flexibly forming, maintaining and disbanding organisations. This representational power enables agent systems to exploit two facets of the nature of complex systems. Firstly, the notion of a primitive component can be varied according to the needs of the observer. Thus at one level, entire sub-systems can be viewed as singletons, alternatively teams or collections of agents can be viewed as primitive components, and so on until the system eventually bottoms out. Secondly, such structures provide the stable intermediate forms that are essential for the rapid development of complex systems. Their availability means that individual agents or organisational groupings can be developed in relative isolation and then added into the system in an incremental manner. This, in turn, ensures there is a smooth growth in functionality.

2.2. Will Agent-Oriented Techniques Be Widely Adopted?

There are two key pragmatic issues that will determine whether agent-oriented approaches catch on as a software engineering paradigm: (i) the degree to which agents represent a radical departure from current software engineering thinking and (ii) the degree to which existing software can be integrated with agents. Each point will now be dealt with in turn.

A number of trends become evident when examining the evolution of programming models. Firstly, there has been an inexorable move from languages that have their conceptual basis determined by the underlying machine architecture, to languages that have their key abstractions rooted in the problem domain. Here the agent-oriented world view is perhaps the most natural way of characterising many types of problem. Just as the real-world is populated with objects that have operations performed on them, so it is equally full of active, purposeful agents that interact to achieve their objectives³. Indeed, many object-oriented analyses start from precisely this perspective: “we view the world as a set of autonomous agents that collaborate to perform some higher level function” [Booch 1994] pg. 17. Secondly, the basic building blocks of the programming models exhibit increasing degrees of localisation and encapsulation [Parunak 1999]. Agents follow this trend by localising purpose inside each agent, by giving each agent its own thread of con-

trol, and by encapsulating action selection. Thirdly, ever richer mechanisms for promoting re-use are being provided. Here, the agent view also reaches new heights. Rather than stopping at re-use of sub-system components (design patterns and component-ware) and rigidly pre-ordained interactions (application frameworks), agents enable whole sub-systems and flexible interactions to be re-used. In the former case, agent designs and implementations are re-used within and between applications. Consider, for example, the class of agent architectures that have beliefs (what the agent knows), desires (what the agent wants) and intentions (what the agent is doing) at its core. Such architectures have been used in a wide variety of applications including air traffic control, process control, fault diagnosis and transportation [Jennings and Wooldridge 1998, Parunak 1999]. In the latter case, flexible patterns of interaction such as the Contract Net Protocol (an agent with a task to complete advertises this fact to others who it believes are capable of performing it, these agents may submit a bid to perform the task if they are interested, and the originator then delegates the task to the agent that makes the best bid) and various forms of resource-allocation auction (e.g. English, Dutch, First-Price Sealed Bid, Vickrey) have been re-used in significant numbers of applications. In short, agent-oriented techniques represent a natural progression of current software engineering thinking and, for this reason, the main concepts and tenets of the approach should be readily acceptable to software engineering practitioners.

The second factor in favour of a wide-spread take up of agents is that their adoption does not require a revolution in terms of an organisation's software capabilities. Agent-oriented systems are evolutionary and incremental as legacy (non-agent) software can be incorporated in a relatively straightforward manner. The technique used is to place wrapping software around the legacy code. The wrapper presents an agent interface to the other software components. Thus from the outside it looks like any other agent. On the inside, the wrapper performs a two-way translation function: taking external requests from other agents and mapping them into calls in the legacy code, and taking the legacy code's external requests and mapping them into the appropriate set of agent communication commands. This ability to wrap legacy systems means agents may initially be used as an integration technology. However, as new requirements are placed upon the system, so bespoke agents may be developed and added. This feature enables a complex system to grow in an evolutionary fashion (based on stable intermediate forms), while adhering to the important principle that there should always be a working version of the system available.

3. CASE STUDY: PROVISIONING A VIRTUAL PRIVATE NETWORK

As an exemplar of a complex software system consider the task of dynamically provisioning a public communication network (such as the Internet) as a virtual private network for end users. To be concrete, let the task in question be setting up a video-conferencing meeting [Faratin et al. 2000]. This application involves a variety of different individuals and organisations (figure 3). There are the end users that are each represented by their personal communication agent (PCA). The providers of services on the network (such as setting up a video-conference, for example) are each represented by a service provider agent (SPA). Finally, there are the agents that represent the network provider on whose telecommunications infrastructure the services will actually be delivered (each represented by a network provider agent (NPA)). In setting up a video-conference call, the various PCAs negotiate (using one of the techniques described in [Jennings et al. 2001]), on behalf of their user, with one another in order to find a suitable time for the call. When they come to an agreement, one of the PCAs then contacts, and subsequently negotiates with, the various SPAs that offer the video-conference service (not all SPAs will do this). This negotiation revolves around the cost of the conference call and the quality of service that is desired. The SPA that wins the contract then negotiates with the various NPAs to determine which of them can deliver the desired quality and bandwidth at the best price.

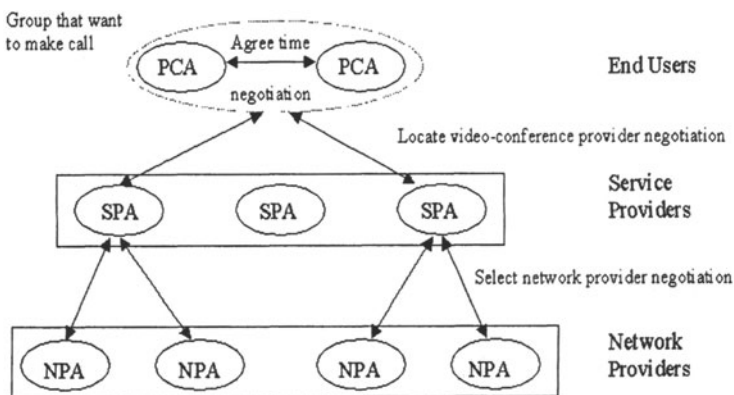


Figure 3. Dynamic Provisioning of Virtual Private Networks by End Users

This application highlights many of the benefits that are claimed above for an agent-oriented approach to software engineering. Firstly, autonomous agents are the most natural means of representing the distinct individuals and organisations that are present in the application. Each such entity is an active problem solver that has its own objectives to achieve and has control over the actions it chooses and the resources that it expends. The agents need to be responsive to changes in their environment (e.g. a NPA may need to buy in additional network capacity from another NPA in order to maintain its agreed upon quality of service if part of its network fails) and they need to be able to opportunistically adopt new goals as they present themselves (e.g. two SPAs may discover they have complementary service capabilities and may decide to act together in order to offer a new service).

Secondly, the agents need to engage in knowledge-level interactions in order to achieve their individual objectives. In this case, agents typically represent self-interested entities and so the main form of interaction is negotiation. Thus, to set the time of the video conference or to select a particular service/network provider the agents make proposals, trade offers, make concessions and, hopefully, come to agreements. This rich form of interaction is necessary because the agents represent autonomous (competitive) stake holders and also to ensure that agents can arrange their activities in a manner that is appropriate to their prevailing circumstances.

Finally, there is a very clear and explicit notion of organisational context. The application involves a number of different real-world organisations: individual end users, companies that provide the different types of services and network providers that control the underlying telecommunications infrastructure. These relationships directly affect the agents' behaviour. For example, if a SPA and a NPA are in fact part of the same organisation, then their negotiations are more cooperative in nature than if they represent two unrelated companies. Similarly, the PCAs that have agreed to hold a conference call, act as a team rather than a collection of individuals. Additionally, during the ongoing operation of the application new organisational groupings can appear and then disband. The PCAs of distinct end users form themselves into collectives when they require a particular service (e.g. all the participants of the video-conference). Individual SPAs combine their capabilities to offer new services that are beyond the scope of any individual provider. Competing NPAs form themselves into temporary coalition in order to respond to particularly large requests for network resources ⁴.

4. CONCLUSIONS

Agent-oriented techniques are increasingly being used in a range of telecommunication, commercial and industrial applications. However if they are to enter the mainstream of software engineering then it is vital that clear arguments are advanced as to their suitability for solving large classes of problems (as opposed to specific point solutions). To this end, this paper has sought to justify precisely why agent-oriented approaches are well suited to developing complex, distributed software systems. These general points are then made more concrete by showing how they apply in a specific telecommunications application. In making these arguments, it is possible for proponents of other software engineering paradigms to claim that the key concepts of agent-oriented computing can be reproduced using their technique. This is undoubtedly true. Agent-oriented systems are, after all, computer programs and all programs have the same set of computable functions. However, this misses the point. The value of a paradigm is the mindset and the techniques it provides to software engineers. In this respect, agent-oriented concepts and techniques are both well suited to developing complex, distributed systems and an extension of those currently available in other paradigms.

Notes

1. Booch actually uses the term "hierarchy" [Booch 1994]; however, this invariably gives the connotation of control. Hence the more neutral term "organisation" is used here.
2. The view that decompositions based upon functions/actions/processes are more intuitive and easier to produce than those based upon data/objects is even acknowledged within the object-oriented community (see [Meyer 1988] pg 44).
3. Although there are certain similarities between object- and agent- oriented approaches (e.g. both adhere to the principle of information hiding and recognise the importance of interactions), there are also a number of important differences. Firstly, objects are generally passive in nature: they need to be sent a message before they become active. Secondly, although objects encapsulate state and behaviour realisation, they do not encapsulate behaviour activation (action choice). Thus, any object can invoke any publicly accessible method on any other object. Once the method is invoked, the corresponding actions are performed. Thirdly, object-orientation fails to provide an adequate set of concepts and mechanisms for modelling complex systems: for such systems "we find that objects, classes and modules provide an essential yet insufficient means of abstraction" [Booch 1994] pg 34. Individual objects represent too fine a granularity of behaviour and method invocation is too primitive a mechanism for describing the types of interactions that take place. Recognition of these facts, led to the development of more powerful abstraction mechanisms such as design patterns, application frameworks, and component-ware. Whilst these are undoubtedly a step forward, they fall short of the desiderata for complex systems developments. By their very nature, they focus on generic system functions and the mandated patterns of interaction are rigid and pre-determined. Finally, object-oriented approaches provide only minimal support for specifying and managing organisational relationships (basically relationships are defined by static inheritance hierarchies).

4. In contrast, an object-oriented approach is less suitable for this problem because: it cannot naturally represent the autonomous problem solving behaviour of the constituent components (recall objects do not encapsulate action choice), it has nothing to say about the design of flexible problem solvers that balance reactive and proactive problem solving nor about inter-agent negotiation (other than the fact that it involves message exchanges) and it has no innate mechanism for representing and reasoning with the fact that the agents represent different stakeholder organisations (other than the fact that they are different classes).

References

- G. Booch (1994). *Object-oriented analysis and design with applications*. Addison Wesley.
- F.P. Brooks (1995). *The mythical man-month*. Addison Wesley.
- P. Faratin, N.R. Jennings, P. Buckle, and C. Sierra (2000). *Automated negotiation for provisioning virtual private networks using FIPA-compliant agents*. Proc. 5th Int Conf. on Practical Application of Intelligent Agents and Multi-Agent Systems, Manchester, UK, 185-202.
- N.R. Jennings (2000). *On Agent-Based Software Engineering* Artificial Intelligence 117(2): 277-296.
- N.R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra and M. Wooldridge (2001). *Automated negotiation: prospects, methods and challenges* Int. J. of Group Decision and Negotiation 10 (2) 199-215.
- N.R. Jennings and M. Wooldridge (eds.) (1998). *Agent technology: foundations, applications and markets*. Springer Verlag.
- A. Newell (1982). *The knowledge level*. Artificial Intelligence 18:87-127.
- J. Mayfield, Y. Labrou, and T. Finin (1995). *Evaluating KQML as an agent communication language*. In M. Wooldridge, J.P. Muller and M. Tambe, eds., Intelligent Agents II, Springer, 347-360.
- B. Meyer (1988). *Object-oriented software construction*. Prentice Hall.
- H.V. D. Parunak (1999). *Industrial and practical applications of distributed AI*. In G. Weiss, ed., Multi-agent systems, MIT Press, 377-421.
- H.A. Simon (1996). *The sciences of the artificial*. MIT Press.
- M. Wooldridge (1997). *Agent-based software engineering*. IEE Proc Software Engineering 144:26-37.
- M. Wooldridge and N.R. Jennings (1995). *Intelligent agents: theory and practice*. The Knowledge Engineering Review 10(2): 115-152.