

Agent-Based Control Systems

Nicholas R. Jennings¹ and Stefan Bussmann^{1,2}

¹ Dept. Electronics and Computer Science, ² DaimlerChrysler AG

University of Southampton,

Research and Technology 3, RIC/EN

Southampton SO17 1BJ, UK.

Alt-Moabit 96a, 10559 Berlin, Germany.

nrj@ecs.soton.ac.uk

stefan.bussmann@daimlerchrysler.com

INTRODUCTION

Modern control systems require sophisticated software in order to meet their ever more demanding requirements. These requirements stem from the need to deal with increased uncertainty, cope with more dynamic environments, and provide greater flexibility. This, in turn, means that control systems software is highly complex in that it invariably has a large number of parts that have many interactions [20]. This complexity means that state-of-the-art software engineering methods and techniques need to be employed. To this end, this paper will argue that analysing, designing and implementing such complex software systems as a collection of interacting, autonomous, flexible components (i.e., as agents) affords software engineers a number of significant advantages over contemporary methods. In seeking to demonstrate the efficacy of the *agent-oriented software engineering* approach [11], the most compelling argument would be to quantitatively show how its adoption improved the development process in a range of (control systems) projects. However, while there are now several deployed applications (see [15] [19] for overviews) 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 control systems. This general position is then augmented with two specific case studies in the domains of industrial process control (in particular, electricity transportation management) and manufacturing control (in particular, manufacturing line control) where the experiences of using an agent-based approach are assessed.

Before making the general case for agent-oriented software engineering, however, we first discuss

the characteristics of complex software systems (of which control systems are naturally an instance). We then go on to discuss the methods that software engineers have developed to help manage this complexity.

The role of any new software engineering paradigm is to provide structures and techniques that make it easier to handle complexity. Fortunately for designers, this complexity exhibits a number of important regularities [20]. Firstly, 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, etc.) can be identified. These relationships are not static: they often vary over time. Secondly, the choice of which components in the system are primitive is relatively arbitrary and is defined by the observer's aims and objectives. Thirdly, 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). Fourthly, 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.

Given these observations, software engineers have devised a number of fundamental tools of the trade for helping to manage this complexity [2]:

- *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*: 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). This covers phenomena such as inheritance (in object-oriented systems) and subroutines (in procedural languages). 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.

Having characterised complex software systems and identified the fundamental software engineering approaches that help manage this complexity, the case for agent-oriented software engineering can now be made.

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 [22]

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) [24].

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 [1]. 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 situated in a common environment. These interactions can vary from simple semantic inter-operation (information passing), through traditional client-server type interactions, to rich social interactions (the ability to cooperate [12], coordinate [10] and negotiate [13] 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 [17]). Consequently, interactions are conducted at the knowledge-level [16]: 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 ability 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 the agents' interactions. This context defines the nature of the relationship between the agents and/or the rules that must be adhered to during the course of the interaction. For example, the agents may be peers working together in a team, or one may be the manager of the others, or the agents must negotiate according to the rules of a particular auction house. To capture such links, agent systems have explicit constructs for modelling organisational relationships (e.g. manager, team member, auctioneer). In many cases, these rela-

tionships are subject to ongoing change: social interaction means existing relationships evolve (e.g. a team of peers may elect a leader) and new relations are created (e.g. a number of unrelated agents band together 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 [6], specified mechanisms to ensure groupings act together in a coherent fashion [10], [12] and developed structures to characterise the macro behaviour of collectives [15] [24].

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.

From a control perspective, this view of software systems has several similarities to work on heterarchical systems in distributed control [8]. The work of [9], for example, avoids the drawbacks of hierarchical control by distributing the decision making into intelligent parts. However, the work on heterarchical control tends to concentrate on the distributed systems nature of these control systems and, to a certain extent, on the autonomy of the individual components, rather than on the flexible, high-level nature of the interactions and the explicit representation of the organisational context.

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: (i) show that agent-oriented decompositions are an effective way of partitioning the problem space of a complex system; (ii) show that the key abstractions of the agent-oriented mindset are a natural means of modelling complex systems; and (iii) 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 (cross product) between the characteristics of a complex system and the key software engineering abstractions for handling complexity as they apply to agent-based systems. Each step is now dealt with in turn.

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 achieve. (Indeed 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 [18] pg 44).) 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” [18] 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 act and interact in flexible ways to achieve their 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” [2] 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 [6]. Interaction protocols exist for forming new groupings and disbanding unwanted ones. Finally, structures are available for modelling collectives [12]. 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.

The Need for Flexible Management of Changing Organisational Structures

Organisational constructs are first-class entities (in the programming language sense) 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.

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 objectives. 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” [2] pg. 17. Secondly, the basic building blocks of the programming models exhibit increasing degrees of localisation and encapsulation [19]. Agents follow this trend by localising purpose inside each agent, by giving each agent its own thread of control, 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 [23]. Such architectures have been used in a wide variety of applications including air traffic control, process control, fault diagnosis and transportation [15] [19]. In the latter case, flexible patterns of interaction such as the Contract Net Protocol [21] (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- and Second- Price Sealed Bid [25]) have been re-used in significant numbers of applications (see the manufacturing line control scenario). 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 existing software systems. Agent-oriented systems are evolutionary and incremental as legacy (non-agent) software can be incorporated in a relatively straightforward manner (see the electricity transportation case study). 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.

CASE STUDIES

Having discussed the potential benefits of agent-based systems for complex systems in general, we now turn to two specific agent-based control system applications. The aims in presenting these case studies are: (i) to ground the abstract concepts of agent-based computing in specific application contexts; and (ii) to highlight the practical advantages that can accrue from an agent-based solution. Moreover, the scope and applicability of agent-based solutions is emphasised by discussing two examples that are at fundamentally different levels of the control spectrum: one at the level of controlling an entire network and the other at the level of controlling an individual production line.

Electricity Transportation Management

This application was developed and deployed by the Spanish electricity utility Iberdrola (more details of the underpinning agent technology can be found in [14] and more details of its application in this domain can be found in [7]). Generally speaking, energy management is the process of monitoring and

controlling the cycle of generating, transporting and distributing electrical energy to industrial and domestic customers. Generation transforms raw energy (e.g. hydraulic, thermal, nuclear and solar) into a more accessible form that then needs to be *transported* from its generation site to the consumer. To minimise losses during transportation, the electrical voltage is made high (132 kV or above) before it is placed on a *transport network* and sent over many hundreds of kilometres. Finally, the voltage is lowered and electricity is delivered to the consumers using a *distribution network* which involves many kilometres of network (all below 132 kV) spread over a much smaller area.

To ensure the transportation network remains within the desired safety and economical constraints, it is equipped with a sophisticated data acquisition system (SCADA) and several conventional application programs that help the operator (a *control engineer*) to analyse it (these programs are primarily designed for normal operating conditions). The network's operation is monitored from a *Dispatching Control Room* (DCR) and whenever an unexpected event occurs hundreds of alarms are automatically sent to it by the SCADA system. Under these circumstances the operator has to rely on experiential knowledge to analyse the information, diagnose the situation, and take appropriate remedial actions to return the network to a safe state. To reduce the operators' cognitive load in such circumstances, and to help them make better decisions faster, Iberdrola first developed a number of (stand-alone) decision support systems (e.g., a real-time database that stores information about the state of the network and an alarms analysis expert system that diagnosed faults produced in the network based on the alarm messages received at the DCR). To improve this support, it was decided that these systems should be made to interoperate to produce a coherent view and that new functionality should be added (to enable the control engineer to actually perform and dynamically monitor the service restoration process and also to exploit the new data sources, such as chronological information and faster rate snapshots, which became available as the SCADA system was improved). What follows is a description of how this was achieved using agent technology.

Why use agent techniques for this application?

This application required the tried and tested decision support tools to be integrated and extended with new functionality. Two means of realising this system upgrading strategy were considered: extend the

existing systems to cover the new features or follow a distributed approach and allow the new functionality to be expressed as distinct computational entities that can interact with the existing systems through a common distribution platform. Here the second option was chosen because it was considered to be the most effective means of:

- permitting reasoning based on information of different granularity

Two types of alarm, non-chronological and chronological, need to be dealt with. In the former case, the time stamped is coincident with the time of acquisition by the control system (consequently it is conditioned by the control system's polling mechanism), whereas in the latter case the time stamped is coincident with the actual occurrence of the event. As chronological alarms represent a more accurate picture of events in the network they generally lead to a swifter diagnosis, however they have the disadvantage that chronological information has a low priority in Iberdrola's communication channels. Thus when the channels are saturated (as often happens during a disturbance) their time of arrival is unpredictable. For these reasons it was decided to build a new alarm analysis expert system that utilised chronological information and could subsequently integrate its results with those of the pre-existing system, rather than construct a monolithic system that received both types of data and had to embody both types of diagnostic knowledge. A similar situation occurs when considering service restoration. Two types of information are relevant to this activity: *snapshots* (which provide a comprehensive picture of the current state of all the network's components) and *alarm messages* (which show how the state of the components has changed over a period of time). The former can be produced relatively quickly and give a complete picture of the system's state, whereas the latter may take several minutes for a large disturbance but are needed to indicate the type of fault from which the system must be restored. Rather than trying to place both types of information and reasoning in a single system it seemed more natural to develop a service restoration subsystem that dealt mainly with snapshots and received the necessary high-level information about the equipment at fault from a diagnosis subsystem (rather than trying to deal with the raw alarm messages itself).

- allowing different network models to be included within the same system

Some of the problem solvers need to work on the SCADA model of the network, while others need the applications network model (a model that permits differential equations to be solved and takes the physical characteristics of all its components into account). Rather than trying to combine and harmonise these complex and disparate models at design time, it was decided that each subsystem should work on whichever model was most appropriate for its task. Then the various components can interact at runtime to resolve any inconsistencies that arise from their use of different network models.

- enabling a number of different problem solving paradigms to be utilised

The diverse range of activities that need to be performed mean there is no universally best problem solving paradigm: procedural techniques are required for algorithmic calculations like connectivity (to know which component is connected to which other) and load-flow analysis (solution of the differential equations), whereas symbolic reasoning based on heuristic search is best for diagnosis. A distributed approach enabled each component to be encoded in the most appropriate method.

- meeting the application's performance criteria

Transportation management is a time-critical application and as many different types of information can be processed in parallel, with only a small synchronisation overhead, the response time of the overall system can be improved through the use of a number of interconnected machines.

Having decided upon a distributed approach, a choice had to be made between using more conventional distributed processing techniques or agent-based techniques. Here the latter was adopted for the following reasons. (i) *Robustness*: As the subsystems have overlapping domains of expertise, the failure of one of them to produce an answer does not necessarily mean that no solution will be forthcoming (because one of the other systems may be able to produce at least a partial solution). However to achieve this back-up functionality in a flexible manner, the different problem solving components need to be intelligently coordinated in a context sensitive fashion (see the cooperative scenario for more details) - a task beyond present generation distributed processing systems. (ii) *Reliability*: The solutions

of the systems that overlap can be cross-referenced to enable the operator to be presented with more reliable information. Again, however, this cross-referencing functionality needs to be properly managed according to the prevailing circumstances and so requires dynamic and flexible reasoning to take place. (iii) *Natural representation of the domain*: An agent-based approach accurately represents the way the control engineers work when a large disturbance occurs. They specialise their roles - one looks after restoration, another tries to diagnose the problem based on different sources of information, and so on - and they then communicate relevant information to one another to ensure they are following a coherent course of action towards the overall objective of restoring the service.

Specification of the Agents

During normal working conditions, management of the network by the operator in the DCR consists mainly of routine and simple tasks. However, during emergency situations management becomes considerably more difficult because of the large number of constraints that have to be taken into consideration and the insufficient quality of the information that is available to make these decisions. Emergency situations typically originate from a short circuit in a line, bus-bar or transformer. They can be exacerbated by equipment malfunctioning (e.g. a breaker failing to open) or subsequent overloads (a domino effect can cause one line to fail because of an overload, this in turn increases the load on neighbouring lines so they become overloaded and subsequently fail, and so on). The situation can become even worse if power stations become disconnected as this will cause an imbalance in the network's power. Consequently, actions to restore service must be taken rapidly and accurately, so that what starts as a relatively minor problem does not escalate into a major disaster. In these circumstances, the actions that the operator can perform consist mainly of breaker operations, topology changes, and activation/deactivation of automatism and protective relays. For larger disturbances, however, actions on power plants may also be required. From this description of the control engineer's job, a top-down analysis identified that a comprehensive decision support system should cover the following activities:

- Detect the existence of disturbances; sometimes the operation of protective relays and breakers can be caused by routine maintenance and this should not be confused with genuine disturbances.

- Determine the cause, location and type of the disturbance; including identifying if any equipment is permanently damaged.
- Analyse the situation of the network once it arrives at a steady state.
- Prepare a restoration plan to return the network to its original operational state.

Allying this top-down analysis with the bottom-up perspective of examining the extant systems, it was decided to encapsulate the following pre-existing systems as agents - the alarms analysis expert system and the interface to the control system. As discussed above, the availability of chronological alarm messages necessitated a new diagnosis system which we decided to make available as an agent. Finally, it was always known that information about the initial area out of service (the *black out area*) could help constrain the search of the faulty equipment, however it was never deemed cost effective to develop a dedicated stand alone system for this purpose since the original alarm analysis expert system's performance was considered satisfactory (if somewhat slow). However through the use of agent technology much of the basic infrastructure to implement this functionality was now available from other agents and so it was considered economically viable to develop a system capable of producing this information.

In more detail, the operational system consists of seven agents running on five different machines (figure 3). This figure shows a small portion of the Iberdrola network which contains four substations (Sestao, Sodupe, Erandio and Achuri). Each of these substations has a corresponding remote transmission unit (RTU) which sends information to the DCR in Bilbao about the status of its bus-bars, breakers, and other electrical components. In the DCR, this information is collected by the front end computer and made available to the cooperating agents through the control system's interface functions.

BAI (Black-out Area Identifier) When a fault occurs, the network's protective relays and breakers automatically try to isolate the minimum amount of equipment possible; in an ideal case only the element at fault would be isolated. The BAI's objective is to identify which elements of the network are initially out of service as the actual element at fault must be within this region. It uses non-chronologi-

cal alarm messages as its information source and cooperates with the BRS and the AAA to increase the efficiency of the overall diagnosis process.

CSI (pre-existing Control System Interface) The CSI acts as the front end to the control system computers. Its objectives are to acquire and distribute network data to the other agents, to interface to the conventional management system application programs, and to monitor the restoration process to detect unexpected deviations. It is split into two physical agents: CSI-D which detects the occurrence of disturbances and preprocesses the chronological and non-chronological alarm messages that are used by the AAA, BAI and BRS agents; and CSI-R which detects and corrects inconsistencies in the snapshot data file of the network, calculates the power flowing through it and makes this information available to the SRA and the UIA.

BRS (Breakers and Relays Supervisor) The new alarms analysis expert system that detects the occurrence of a disturbance, determines the type of fault and its extent, generates an ordered list of fault hypotheses, validates hypotheses, and identifies malfunctioning equipment. In order to perform its analysis, it takes two types of inputs: chronological alarm messages and snapshots of the network that give the status of every breaker and switch.

AAA (pre-existing, non-chronological Alarms Analysis Agent expert system) This agent pursues similar goals to the BRS, however the quality of information it receives is inferior to that of the BRS. Although the alarm messages received by both systems relate to the same physical operations, those received by the AAA represent ± 5 seconds accuracy, while those received by the BRS are precise. This means that if the data is error free, then the BRS performs a better diagnosis than the AAA. However if some of the chronological information is lost (a distinct possibility when the SCADA system is busy) then the BRS may perform worse than the AAA. Therefore whenever incomplete or erroneous information exists, which is in most interesting cases, there is a need for cooperation between the two systems to make the overall system more robust and reliable (see below).

SRA (Service Restoration Agent) This agent devises a service restoration plan to return the network to a steady state after a blackout has occurred. To do this it takes into account the constraints imposed by the damaged equipment, as identified by the diagnosis agents.

UIA (User Interface Agent) This agent implements the interface between the users and the community of agents. It gives the user the facility to inspect the results produced by the diagnosis agents, display the alarms received, and browse through the log of analysed disturbances. From the point of view of restoration, the user can see the plan produced, modify it, run it in a simulated environment to see its predicted effect, and request the development of a new restoration plan that takes into account the actions that are deemed pertinent. Through the use of a distributed windowing system, the UIA presents the appropriate information on the consoles of the various control engineers who are working on the system (figure 3 shows two such control engineers - one working on restoration activities and one working on diagnosis activities).

This system design ensures all the tasks identified by the top-down analysis are performed by at least one agent. Robustness is achieved by having multiple agents that are able to provide the same (or at least some) overlapping results. Efficiency is obtained by the parallel activation of tasks. Reliability is increased because even if one of the agents breaks down the rest of the agents can often produce a result which, although not as good as the one provided by the complete system, is still of use to the operator.

Cooperative Diagnosis and Restoration

An important example of cooperation in this system involves the information interchange between the AAA, BRS and BAI agents. The AAA and the BRS produce the same result from different information sources, while the BAI applies different knowledge to produce a result that should be coherent with that of the AAA and the BRS.

Assume a block of non-chronological alarm messages has been provided by the SCADA system and these alarm messages have been identified as related to a disturbance by the CSI. Using its model of the other agents, the CSI will realise the alarms are relevant to the AAA and the BAI and so will voluntarily send them out as unsolicited data. Some time later, the same process will be repeated and the BRS will receive the corresponding chronological alarm messages. At this point, the AAA, BAI and BRS are all operating in parallel.

When the AAA receives the alarm messages, it starts its diagnosis process and a preliminary set of hypotheses are produced. During this time, the BAI would also have received the alarm messages and

would have started trying to identify the initial black out area. Again based on its models of the other agents, the BAI sends this information to the AAA. Simultaneously, although after a certain delay, the BRS agent starts working on the analysis of the chronological alarm messages. This will also result in a list of initial hypotheses being produced. The BRS checks whether any agents are interested in this information - again the AAA is noted and the hypotheses are sent to it. The BRS then continues with its diagnosis to try and validate the cause of the fault.

After producing its tentative list of hypotheses, the AAA proceeds with a detailed analysis to try and ascertain the precise cause of the fault. The following situations may then occur: (i) the initial black out area is available to the AAA, this triggers a refinement behaviour which may reduce the number of hypotheses to be validated because the BAI has given a focused view of the situation; (ii) the initial hypotheses provided by the BRS are available to the AAA, this triggers another refinement behaviour and obtains a better reordering of the hypotheses to be validated and a benefit in finding the element at fault; (iii) the validated hypotheses provided by the BRS are available to the AAA, this triggers yet another refinement behaviour, that has the same functionality as the previous one, but the reordering is based on validated hypotheses which are more accurate; (iv) if no information is available from the BAI or BRS, the AAA proceeds with its hypotheses validation as a standalone agent. Therefore, if the other agents are down or they are too slow to provide the information, the AAA will continue and find a faulty element although its diagnosis will be less reliable and will take longer.

The restoration process is activated whenever a disturbance is detected. Once the disturbance is identified, the disturbance identifier is sent to the CSI-R which acquires the snapshot of the network, corrects any inconsistencies that have arisen in its representation, and calculates the power flow solution of the current state. This information is then passed onto the SRA so that it can prepare for its restoration planning. The SRA waits until the diagnosis agents have informed it of the element suspected of being at fault and then proceeds to prepare a restoration plan. If, during this plan preparation, the SRA is informed that the equipment at fault is different from that originally indicated by either the AAA or the BRS, then it replans the restoration taking this information into account.

The UIA is the interface through which the user accesses the results produced by the agent commu-

nity. During the diagnosis phase, the user is presented with both the tentative (early) list of suspected hypotheses and the final (validated) list. During the restoration phase, the UIA supports a more participatory interaction between the user and the teams of agents. The user is presented with the restoration plan and can then decide to modify it, run a detailed simulation to see the effects of the plan or ask for a new plan to be devised taking into account new constraints which he specifies.

Observations and Reflections

This application afforded a number of benefits. Firstly, the agent system gives better results than its stand alone counterparts because it takes multiple types of knowledge and data into account and then integrates them in a consistent manner. Secondly, the agent system is more robust because there are overlapping functionalities; meaning partial results can be produced in the case of component (agent) failure. Thirdly, some results can be provided more quickly because cooperation provides a short cut. Fourthly, the functionalities of the different domain systems can be increased independently which makes them easier to maintain (see, for example, the argument for developing the BAI and the general point about stable intermediate forms). Fifthly, the control engineer is provided with an integrated view of the results of interest. Finally, the system has been designed to be open so that new agents can be added in an incremental manner.

One of the key features of this multi-agent system is the way it handles fault diagnosis by using two different types of data (the non-chronological alarms used by the AAA and the chronological alarms used by the BRS) and two different points of view (the typical diagnosis approach of hypothesis generation and validation used by the AAA and BRS, and the BAI's monitoring approach which provides a high level view of the status of the network). With this set-up, it is possible to dynamically select the solution method that is best suited to the current situation. For example, if the BRS is operational, but the AAA is not, the solution provided to the control engineer is the one created by the BRS; but if both the BRS and the AAA are running, the solution provided is the one that is mutually agreed between them. Also the fact that multiple agents are trying to generate the same results can be exploited to avoid repetition of certain tasks if it is deemed desirable in a particular context. For example both the AAA and the BRS can provide initial hypotheses, consequently if these hypotheses are provided by the BRS

and are made available to the AAA before it starts its own generation task, then this task need not be executed and the hypotheses provided by the BRS can be used instead. This ability to flexibly manage, at runtime, multiple sources of data and multiple problem solving perspectives provides enormous robustness to the overall system because if one of the agents crashes the others will still be able to provide some form of solution.

Manufacturing Line Control

This application was developed by a DaimlerChrysler led industrial consortium and was deployed on one of DaimlerChrysler's production lines in Stuttgart, Germany (more details can be found in [5]). The overall aim of the system is to provide a flexible and robust system for controlling a manufacturing line. In this process there are a number of basic parts (workpieces) that have various operations performed upon them by various machines.

The industry standard approach to manufacturing control is to devise a global schedule, typically covering one day, for the entire manufacturing process. This indicates when the various parts should be released from their stores, which machines they should be routed through, and what operations should be performed at the various machines. The problem with this centralised and pre-planned approach, however, is that plan formation is divorced from plan execution. Thus the schedule can rarely be adhered to in practice: machines and operations fail (sometimes in ways that are difficult to predict) and operations take longer than expected. When such disturbances occur, the plant controller either has to initiate a costly rescheduling exercise or use the out-of-date schedule as an approximate guide. Both of these options lead to inefficiencies in the manufacturing process and are, therefore, undesirable.

Why use agent techniques for this application?

To overcome the aforementioned problems, the control system needs to be made responsive to the prevailing situation of the manufacturing process. In this domain, a centralised controller is not a viable option—it would be too time-consuming to construct and maintain an up-to-date representation of what is going on in the whole system, it would be a severe bottleneck on the system's performance and it would represent a single point of failure (meaning the system is not robust).

Given the requirements for decentralisation, responsiveness and flexible contingency handling, it was

decided to adopt an agent-based approach. In this system, each manufactured part is represented by an autonomous agent that has the objective of getting itself to the end of the manufacturing line, having had a specified set of operations performed upon it. Each machine is also represented by an agent. Such agents have the objective of maximising their throughput and they do this by deciding what parts will be accepted in what order and what operations will be performed at what time. Thus, for a given part to have an operation performed upon it, its agent must negotiate with a machine agent capable of performing that operation. In short, resources are allocated dynamically on a just-in-time basis by a continuous coordination process among the relevant agents.

Specification of the agents and their interactions

To achieve robustness and flexibility the machines must have overlapping capacities. This means for every manufacturing step there must always be more than one machine that can perform it. Thus when machine breakdowns occur the redundancy provides the system with the flexibility of diverting the part to another machine. Diverting a part, however, is not possible without being able to bypass a machine (or machines). To this end, the concept of a modular manufacturing system was developed (figure 4). With this view, the entire manufacturing system is composed of standard modules. Each module consists of a machine, three one-way conveyors, and two transportation switches (figure 5). Every switch can move a part from any of its entry points to any of its exits. In figure 4, each of the intermediary switches has two entries and one exit on the left-hand side and two exits and only one entry on the right-hand side.

An arrangement of standard modules, as in figure 4, means a part can either enter a machine through the lowermost conveyor or else bypass the machine through the middle one. After having bypassed a machine, a part has two options: it can either proceed in a forward direction to a subsequent machine or move backwards using the topmost conveyor. If, for instance, the lowermost conveyor is already occupied, preventing a part from entering the target machine, then the part can move backwards and forwards in a circle until the lowermost conveyor is available again. In this way, the entire transportation system serves as a flexible buffer.

To control this flexible manufacturing system in a decentralised manner a specific agent is associated

with each workpiece, each machine, and each switch. A *workpiece agent* manages the state (which operations need to be performed and which operations have already been performed) of the workpiece attached. A *machine agent* controls the overall material flow through a machine, not just the work in progress. To this end, every machine agent manages what we call a *virtual buffer*. This buffer includes not only the machine's current work in progress, but also the outgoing flow of material; that is, all those workpieces that have already been processed by the machine without yet being able to find an appropriate new machine. The *switch agent* controls a particular switch; deciding which entry to serve first and where to move a part.

All these agents constitute parallel processes. These processes are, of course, not independent. They have to be coordinated. Here coordination is achieved by a negotiation procedure, which also takes place simultaneously. A single workpiece negotiates with the machines about which of them should process it next. In particular, the workpiece auctions off its current due operations by inviting machines to bid for these tasks. Every machine bid includes information about the current state of the machine's virtual buffer. If a workpiece awards one or more operations to a specific machine, then getting to this machine becomes the next goal of the workpiece. The routing of a workpiece is also organised through a sequence of bilateral communications, in each case between the workpiece and the next switch which the workpiece approaches. This continues until the workpiece eventually reaches its goal.

In more detail, the allocation of workpiece operations to machines is carried out by a first-price, sealed auction (the organisational structure). Each round involves three steps:

Step 1: The protocol is always initiated by a workpiece agent; in particular, whenever a workpiece first enters the manufacturing system and, thereafter, immediately after it leaves a machine. In any case, the workpiece determines its current task (next and subsequent operations to be performed on it) and all forward successors of the machine it has just left. The system is more efficient if parts move forward (left to right in figure 4). However in some cases this is not always possible (e.g., due to disturbances) and if workpieces do move backwards it is necessary to enforce an assignment in order to avoid deadlocks (see [5] for more details). The workpiece then sends an invitation to bid for its current task to all these machines.

Step 2: If a machine receives an invitation to bid for a current task, it checks whether it is able to perform it. If it can, then the machine issues a bid; otherwise, it simply ignores the call. Short-term disturbances of some of the machine's operations are ignored here. This is because the subject of the negotiation is a *future* allocation of a subtask and the current situation obviously does not tell us much about a machine's state when the workpiece enters the machine. The machine agent issues no bids without making sure that it is actually ready to accept a new workpiece; it therefore checks the capacity of its virtual buffer. If it does not have capacity it does not answer the call. If the agent does make a bid it includes: (a) the current size of its virtual buffer and (b) the maximum number of the desired operations it is able to perform.

Step 3: The workpiece agent collects all the bids for a specific call. If there are no bids, then it issues another invitation to bid, continuing with step 1; otherwise the workpiece selects the best bid. This selection is based on both components (a) and (b) of a bid, with (a) having a higher priority. In this case, the lower the current size of the virtual buffer, the better. The more operations the maximal subtask (b) contains, the better. The chosen machine agent is then informed and it then includes the relevant workpiece in its (virtual) input buffer.

Once a workpiece agent has selected its next target machine, the workpiece must be moved to its new goal. In a layout like the one depicted in figure 4, there are usually a vast number of different paths ultimately leading to the same goal. Of course, shorter paths are preferred. However, even more important than optimizing the routing, is the avoidance of any congestion (since this can have disastrous consequences on the performance of the overall system). In an unpredictable environment like a manufacturing system, jams can only be avoided by strictly separating the actual routing from the goal itself. In this system, such dynamic routing is ensured through a sequence of bilateral communications, each time between the workpiece and the next switch it approaches. A switch always tries to move a workpiece directly to its goal, thus trying to optimise the routing. If an exit is not available, then an alternative route is taken. In this case, however, the priority of the workpiece is incremented. These priorities are used to decide which workpiece to prefer if a switch has more than one possibility: the workpiece with the highest priority is always served first. This is to avoid indefinitely routing a workpiece along a

cycle rather than to its actual goal.

Observations and Reflections

To evaluate the system, DaimlerChrysler conducted a series of simulations, all of which are based on authentic product types and cycle times. The disturbance characteristics have been taken from existing machines. A typical configuration consists of 4 blocks of identical machines. The number of machines in a block ranges from 5 to 11, giving 36 machines in total. The simulations have shown that the agent-based mechanism is extremely robust against disturbances of machines as well as failures of control units. Moreover, its performance is nearly optimal; achieving about 99.7% of the theoretical optimum.

In addition to this simulation work, the control system has been installed as a bypass to an existing large-series manufacturing line for cylinder heads. The bypass, located in a plant in Stuttgart-Untertürkheim, Germany, is shown in figure 6. The bypass has undergone a series of performance tests which showed that the results of the simulations are valid under real manufacturing conditions. Moreover, the system has now been in routine operation for more than two years and has confirmed its robustness in day-to-day production situations.

The success of the system, both in terms of increased throughput and greater robustness to failure, can be attributed to a number of points. Firstly, representing the components and the machines as agents means the decision making is much more localised. It can, therefore, be more responsive to prevailing circumstances. If unexpected events occur, agents have the autonomy and proactiveness to try alternatives. Secondly, because the schedules are built up dynamically through flexible interactions, they can readily be altered in the event of delays or unexpected contingencies.

CONCLUSIONS

This paper has sought to justify precisely why agent-oriented approaches are well suited to developing complex software systems in general and control systems in particular. These general points are then made concrete by showing how they apply in two very different agent-based control systems—Iberdrola's electricity transportation management system and DaimlerChrysler's manufacturing line control system. 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.

Specifically in terms of the next generation of control systems, we believe that agent-based systems provide a number of advantages. They provide a decentralised solution based on local decision making that gives the system high degrees of flexibility and robustness. The downside of devolving the decision making to autonomous components, however, is that it is correspondingly more difficult to predict overall system behaviour. To this end, work is progressing on agent-oriented methodologies specifically for control applications [4]. In the systems described in this paper, the agents are under the control of a single organisation which makes it easier to tailor their behaviour so that desirable system properties emerge from their interplay. In the more general case of multiple organisations, however, producing predictable system-wide behaviour is still an area of active research.

References

- [1] A. H. Bond and L. Gasser (eds.) (1988) “*Readings in Distributed Artificial Intelligence*” Morgan Kaufmann.
- [2] G. Booch (1994) “*Object-oriented analysis and design with applications*” Addison Wesley.
- [3] F. P. Brooks (1995) “*The mythical man-month*” Addison Wesley.
- [4] S. Bussmann, N. R. Jennings and M. Wooldridge (2002) “Re-use of interaction protocols for decision-oriented applications” *Proc. 3rd Int Workshop on Agent-Oriented Software Engineering*, Bologna, Italy, 51-62.
- [5] S. Bussmann and K. Schild (2000) “Self-organising manufacturing control: An industrial application of agent technology” *Proc. 4th Int. Conf. on Multi-Agent Systems*, Boston, USA, 87-94.
- [6] K. M. Carley and L. Gasser (1999) “Computational organisation theory” in *Multiagent Systems* (ed. G. Weiss) MIT Press 229-330

- [7] J. M. Corera, I. Laresgoiti, and N. R. Jennings, (1996) "Using ARCHON, Part 2: Electricity transportation management" *IEEE Expert* **11** (6) 71-79.
- [8] D. M. Dilts, N. P. Boyd and H. H. Whorms (1991) "The evolution of control architectures for automated manufacturing systems" *Journal of Manufacturing Systems* **10** (1).
- [9] N. A. Duffie and R. S. Piper (1987) "Non-hierarchical control of a flexible manufacturing cell" *Robotics and Computer Integrated Manufacturing* **3** (2) 175-179.
- [10] E. H. Durfee (1999) "Practically coordinating" *AI Magazine* 20 (1) 99-116.
- [11] N. R. Jennings (2000) "On agent-based software engineering" *Artificial Intelligence* **117** (2) 277-296.
- [12] N. R. Jennings (1995) "Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems using Joint Intentions" *Artificial Intelligence* **75** (2) 195-240.
- [13] 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.
- [14] N. R. Jennings, E. Mamdani, J. Corera, I. Laresgoiti, F. Perriolat, P. Skarek and L. Z. Varga (1996) "Using ARCHON to develop real-world DAI applications Part 1" *IEEE Expert* **11** (6) 64-70.
- [15] N. R. Jennings and M. Wooldridge (eds.) (1998) "*Agent technology: foundations, applications and markets*" Springer Verlag.
- [16] A. Newell, (1982) "The knowledge level" *Artificial Intelligence* **18** 87-127.
- [17] J. Mayfield, Y. Labrou, and T. Finin (1995) "Evaluating KQML as an agent communication language" in M. Wooldridge, J. P. Müller and M. Tambe, eds., *Intelligent Agents II*, Springer, 347-360
- [18] B. Meyer (1988) "*Object-oriented software construction*" Prentice Hall.
- [19] H. V. D. Parunak (1999) "Industrial and practical applications of distributed AI" In G. Weiss, ed., *Multi-agent systems*, MIT Press, 377-421.
- [20] H. A. Simon (1996) "*The sciences of the artificial*" MIT Press.

- [21] R. G. Smith and R. Davis (1980) “Frameworks for cooperation in distributed problem solving” *IEEE Trans on Systems, Man and Cybernetics* 11 (1)
- [22] M. Wooldridge (1997) “Agent-based software engineering” *IEE Proc Software Engineering* **144** 26-37.
- [23] M. Wooldridge (2000) “*Reasoning about rational agents*” MIT Press
- [24] M. Wooldridge and N. R. Jennings (1995) “Intelligent agents: theory and practice” *The Knowledge Engineering Review* **10** (2) 115-152.
- [25] P. R. Wurman (2001) “Dynamic pricing in the virtual marketplace” *IEEE Internet Computing* **5** (2) 36-42.

Figure Captions

Figure 1: View of a complex system

Figure 2: Canonical view of an agent-based system

Figure 3: The Transport Network and Iberdrola's Agents

Figure 4: A Flexible Manufacturing System

Figure 5: A Standard Module

Figure 6: DaimlerChrysler's Prototype (Source: BLEICHERT Osterburken)

SIDEBAR: AGENTS VS. OBJECTS

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 [22]. 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” [2] 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).