

A Knowledge Level Characterisation of Multi-Agent Systems

J. Stuart Aitken¹, Franz Schmalhofer¹ and Nigel Shadbolt² *

¹ German Research Center for Artificial Intelligence (DFKI),
P.O. Box 2080, D 67608 Kaiserslautern, Germany email:
aitken,schmalho@dfki.uni-kl.de Phone: (49)0631 205 3439

² Department of Psychology, University of Nottingham,
University Park, Nottingham NG7 2RD, England
email: nrs@psyc.nott.ac.uk

Abstract. This paper presents a knowledge-based analysis of a number of multi-agent systems. The analysis is based on the knowledge level hypothesis of Newell and Clancey. We claim that the models of problem solving which we obtain constitute descriptions of global coherence in multi-agent systems. We contrast the results of this analysis with the conventional views of coherence.

1 Introduction

One objective in the design of multi-agent systems (MAS) is for an agent to coordinate its actions with other agents in a coherent way. The concept of coherence has been studied both formally [24] [11] [26] and empirically [1]. In this paper we propose a knowledge-based analysis of coherence in distributed systems in which the resulting models can be formally specified.

The formal approaches typically combine two intentional attitudes: an information attitude and a pro-attitude (a definition of these terms and an introduction to the various formal approaches can be found in the review chapter of this volume). There are two objectives in the formal specification of multi-agent systems. Firstly, formalisation increases our understanding of multi-agent interactions by making intuitive concepts such as *common knowledge* and *group intention* more precise and revealing new relationships between them. Secondly, the ability to verify the behaviour of a real multi-agent system requires the existence of a formal model against which it can be compared. Achieving the second objective also entails being able to describe existing MAS in terms of the language and concepts of the formalism, this, however, can be problematic. One example is the concept of common knowledge [11] which can be shown to be prerequisite for agreement among agents, but which defines a situation which cannot be achieved in practice. ³

* This work was supported by DAAD Stipendium number 513 009 0293 at the DFKI, Kaiserslautern, Germany

³ A modified concept of common knowledge can be achieved by real systems, for a fuller discussion see [11]

An empirical analysis [1] of existing multi-agent systems concludes that coherence has a number of dimensions including solution quality and efficiency. The efficiency or minimality of the multi-agent solution is also emphasised in [26], and the *sharing of pro-attitudes* (goals, plans etc.) specified in some formal models can be compared with the mechanism of meta-level communication found in a number of implemented systems. Relating the formal models of coherence to the empirical analysis in a rigorous way is task which, in many cases, has yet to be undertaken. A number of criticism of the empirical approach are developed later in this paper, however, we share the objectives of Bond and Gasser [1] in that we aim to gain a better understanding of existing distributed AI systems.

This paper presents an analysis of multi-agent systems which attempts to clarify the concept of global coherence from the knowledge-based perspective. The analysis is based on the knowledge level hypothesis. The central feature of knowledge level (KL) characterisations is their focus on rational activity based on a functional notion of knowledge. This provides the theoretical background which we use to characterise a number of existing systems, and specifies an approach which abstracts away from implementational details to show the essential problem solving actions of the distributed system. While the resulting characterisations are generic, they can easily be related to the information states of the agent system and therefore this approach bridges the gap between theory and practice. We believe that our analysis increases our understanding of existing DAI systems and that it will aid the design of future, more capable, systems and encourage the reuse of models of multi-agent problem solving.

The knowledge level analysis is presented in Section 2. We first describe the KL hypothesis as proposed by Allen Newell [19], and discuss the modified KL theory proposed by William Clancey [3]. A logical notation for the knowledge level is introduced. The methodology is applied in Section 3 to a number of DAI systems whose architecture is currently influential. In Section 3.4 a conventional DAI description of the concept of global coherence is presented, analysed and contrasted with the knowledge level view. In Section 4 we discuss some related work and some conclusions are drawn in Section 5.

2 The Knowledge Level Analysis

This section restates the knowledge level hypothesis, highlighting the concept of agency in Newell's and Clancey's theories [19][3]. Subsequently, some notation is defined for the purpose of formalising knowledge level characterisations.

2.1 The knowledge level hypothesis

The knowledge level is defined by Newell [19] as a computer systems level lying immediately above the symbol level. This level is characterised by knowledge as the medium and the principle of rationality as the law of behaviour. Newell refers to the notion of agency in the definition of the principle of rationality: "If an agent has knowledge that one of its actions will lead to one of its goals

then the agent will select that action". The concept of an agent also appears in Newell's definition of knowledge: "Knowledge. Whatever can be ascribed to an agent, such that its behaviour can be computed according to the principle of rationality.". The concept of an agent is a simple one: "an agent is composed of a set of actions, a set of goals and a body". An agent is then a delimited computational process to which we can ascribe goals, knowledge and actions. For Newell, agency, rational behaviour and knowledge are bound together.

In contrast with previous analyses of computer systems, the knowledge level hypothesis focuses on the knowledge that can be attributed to a rational agent from the point of view of an observer [3]. A knowledge level characterisation is an abstraction made by an observer and hence is radically different from characterisations based on architectural considerations (e.g. those of DAI). It has been noted [3] that a knowledge level characterisation cannot be objective as we must (as the theoretician-observer) consider the intensions of the observer who makes the KL characterisation.

2.2 Clancey's redefinition of the knowledge level

In [3] Clancey modifies the definition of the knowledge level, he views it as a characterisation of a system of agents and not of an isolated agent.

"A KL description is about a situated system, not an agent in isolation. That is, the systems level being described is above that of individual agents. Therefore, a knowledge-level description cannot be identified with (isomorphically mapped to) something pre-existing inside an individual head, but rather concerns *patterns that emerge in interactions the agent has in some (social) world*.

... A KL description is always ascribed by some observer, and so is relative to the observer's frame of reference and is inherently subjective." [3]

This paper adopts Clancey's view of the knowledge level as an ascribed and idealised description of the rational behaviour of a system of (possibly) many agents. In contrast with Newell's definition, this alternative statement takes a more 'elevated' view of the world where agents are now symbol level entities.

The two versions of the knowledge level hypothesis are valuable for explaining different phenomena: the actions of an agent with respect to its knowledge and goals - in the case of Newell's definition, and the behaviour of a system of agents in the case of Clancey's definition.

As yet we have not defined exactly what is represented at the knowledge level. We propose that knowledge in the form of a specific method of problem solving, specified by its categorisations of knowledge, and the goal that is achieved define the knowledge level. This approach is based on the idea of an inference structure as proposed by Clancey [2]. However, we do not map this structure onto a single agent, rather we view the inference structure as a characterisation of the coherent problem solving activity of many agents.

The inference structure is an idealised description, but does not have the problem of omniscience that is associated with a purely logical characterisation of knowledge. It is interesting to note that this problem has arisen in knowledge acquisition research, where the straightforward logical view of inference has been

challenged, and also in agent modelling, where logics of restricted inferential capability have been investigated [8][14].

2.3 Notation for knowledge level characterisations

The notation of the inference structure was introduced to describe heuristic classification [2], a knowledge level characterisation of a problem solving method. The nodes of the inference structure diagram stand for propositions describing a class of objects e.g. Data Abstractions. The arcs are labelled with the name of the inferential process by which the nodes are related e.g. Heuristic Match. A similar notation is used in [25] where the nodes are called meta-classes and describe the role of domain objects in the problem solving process. Knowledge sources name the inferential step which relates the meta-classes. A knowledge source carries out a primitive inference step to produce a new piece of knowledge.

In this paper we formalise the inference structure by defining sets of domain terms, for example the set of symbols DA representing data abstractions, and defining a predicate symbol to denote instances of this class of symbols. We define the logical type a as instances of DA . We can now write that a_1 is a data abstraction by the atomic formula $Data-Abstraction(a_1)$ if a_1 is of type a ($a_1 \in DA$).

Relations between classes of objects are denoted by logical implications. These implications correspond to the arcs of the inference structure diagram. The intention of the logical formalisation of the KL is to specify conditions of consistency. That is, the sets of formulae presented in the following sections do not predict the symbol level behaviour, but specify the knowledge level such that we can determine whether or not the behaviour of the symbol level does indeed correspond to it. This KL formalisation retains the desired features of denoting the essential classes of domain knowledge, and their interrelation.

In this paper we adopt the terminology of Wielinga *et al.* [25] and use the term ‘knowledge source’ to denote inference processes and the term ‘meta-class’ to refer to classes of domain terms.

3 Characterising DAI Systems

A number of distributed AI systems are now characterised at the knowledge level. The systems include the contract net [23][17], a reactive planner [15] and DVMT, a multiple-blackboard system [6].

3.1 The contract net

The contract net [4][22][23] is typically described in terms of the bidding process, i.e. the issuing of a task announcement by a manager-agent, the response of contractor-agents with bids, and the awarding of contracts to the bidders judged most suitable. Agents can play the role of manager or contractor or both. Agents

have pre-defined capabilities, which with their level of activity, determine their response to task announcements.

The contract net is a framework in which decomposable tasks can be distributed among a group of agents in a flexible manner. Global coherence is to be achieved through negotiation as a mechanism for interaction, task decomposition and the common language shared by all agents [4]. However, it has been noted that if tasks cannot be decomposed into independent subtasks then the synthesis of results is problematic [1][17]. The synthesis of results is one dimension of global coherence, and hence we must doubt whether the mechanisms identified in [4] can guarantee global coherence.

The mechanisms of the contract net do not specify a particular solution method or the major categorisations of knowledge involved in solving the domain problem, these mechanisms therefore lie entirely at the symbol level. A concrete example of the use of the contract net for resource allocation is given in [17]. We now present a knowledge level analysis of the same example and contrast the results.

In the resource allocation problem, agents have resources which they require and resources which they are prepared to trade. Resources can be exchanged by a simple bidding process. This, however, does not lead to a globally satisfactory solution as longer sequences of resource exchanges are not explored.

An inference structure which describes this process is defined in Figure 1. The classes *ResourceSet* and *RequirementSet* contain the resources which are available and the resources which are required. The classes *Offer* and *Need* define a particular resource of the *ResourceSet* which is on offer, and an element of the *RequirementSet* which is a resource needed by some agent. The inferences which connect these meta-classes is *selection* in both cases. The final inference is called *assign*. An assignment is made if an *Offer* matches a *Need*.

Type	Description of set
<i>a</i>	Set of symbols representing agent names
<i>b</i>	Set of symbols representing resources
<i>c</i>	$a * b$
c^*	Powerset of <i>c</i>
<i>d</i>	Set of integers
e^*	Powerset of $b * d$
Selection of Need	
1	$(\forall x : c^*)(RequirementSet(x) \rightarrow (\exists y : a)(\exists z : b)Need(y, z))$
Selection of Offer	
2	$(\forall x : c^*)(ResourceSet(x) \rightarrow (\exists y : a)(\exists z : b)Offer(y, z))$
Assignment	
3	$(\forall x : a)(\forall y : b)(\forall z : a)(Need(x, y) \wedge Offer(z, y) \rightarrow Assignment(x, y))$

Fig. 1. A formalisation of the inference structure for resource allocation

The symbol level description is mapped onto the knowledge level as follows.

The decision of an agent to issue a request for a resource (as a task announcement) corresponds to the selection inference, i.e. the selection of a need. This is termed selection as if an agent has more than one resource which it requires it must, necessarily, make a choice between the alternatives. The decision of an agent to bid for the contract corresponds to the selection of an offer in the knowledge level description. The acceptance of a bid by the agent who issued the task announcement corresponds to the assignment inference.

The selection processes may simply be dependent on the order of elements in the sets, and under this assumption the formalisation produces the behaviour described in ([17] p301) for a simple contract net approach to the assignment problem.

It has been observed [17] that the quality of the global solution is improved if agents have information regarding the overall scarcity of resources. This requires that agents should not make decisions on purely local information. The knowledge of resource scarcity must be collected from all participating agents and a manager agent can be introduced to perform this task [17]. From the knowledge level point of view, the gathering of 'market statistics' represents a new type of knowledge which plays a role in the solution process. The method of problem solving is modified and we must redescribe the inference structure as a result. In specific, a new inference is added which calculates the *MarketStatistics*, and these statistics become a factor in the selection processes.

Calculate market statistics

$$4 \ (\forall x : c*)(\forall y : c*)(RequirementSet(x) \wedge ResourceSet(y) \\ \rightarrow (\exists z : e*)(MarketStatistics(z)))$$

Redefine selection of Need

$$1' \ (\forall w : c*)(\forall x : e*)(RequirementSet(w) \wedge MarketStatistics(x) \\ \rightarrow (\exists y : a)(\exists z : b)Need(y, z))$$

Redefine selection of Offer

$$2' \ (\forall w : c*)(\forall x : e*)(ResourceSet(w) \wedge MarketStatistics(x) \\ \rightarrow (\exists y : a)(\exists z : b)Offer(y, z))$$

Concretely, the selection process now takes place under an ordering relation defined by the market statistics. Scarce resources will now be traded for other scarce resources in preference to abundant resources (for the sake of simplicity we have not specified exactly how this is done in the formalisation).

In [17], knowledge about resources is distributed among many agents and agent activity occurs in four distinct phases; the calculation of the market statistics being the first, and the bidding process is the final phase. Precisely how the assignments are made is determined not only by the knowledge of the agents but also by the bidding protocol and hence is dependent on the bidding strategies of the agents (we have described only the simplest strategy). In [17] the improved method is described as delegated negotiation and is seen as an improvement in the organisation of the agents by the introduction of a manager agent. In contrast, the knowledge level characterisation focusses on the knowledge required to solve the problem and identifies the improved performance as being due to a new source of knowledge.

In this section we have described an application of the contract net at both the knowledge level and the symbol level. We have shown how the knowledge level specification can be used to explain why different organisations of agents produces different behaviours and differing qualities of solution.

3.2 A reactive planner

A reactive planner based on the RTA* algorithm [16] is described in [15]. The application is the 8-puzzle. The goal of the planner is to derive a sequence of moves which results in the goal state being reached, but the plan need not be optimal. The planner can be configured to be reactive, that is, at each state the best action is selected by calculating weights according to a simple distance function. Past moves are recorded and the weights associated with actions are modified if the state has been reached before. This prevents the planner from entering a loop. If two actions have equal weights then a random choice is made.

The RTA* reactive planner has properties which are highly valued in DAI, namely, that actions are executed based on local information only. This can be contrasted with the deliberative planning approach where the space of possible moves would be searched. There are a number of possible implementations of the RTA* planner:

- Option 1 As an off-line planner which returns a solution.
- Option 2 As a real time planner which executes each action when it is determined.
- Option 3 As a system on n off-line planners working in parallel
- Option 4 As a system of 8 agents each representing a square of the 8 puzzle.

Options 1-3 have been described in [15]. Option 3 is interesting as it makes use of the property of the RTA* algorithm that running the algorithm several times on the same problem results in a distribution of solution lengths. This is due to the random choice of equally weighted actions, some choices turn out to have been better than others. By running n planners in parallel, where n is greater than about 10, the probability is high that one planner will find a short solution, and this plan will be found first. There are great benefits in terms of computation time in this approach [15]. Option 4 is similar to the proposal of [5].

The distinctions of knowledge made in all of the above implementations are characterised in Figure 2. We can distinguish the classes of the current state, *State*, possible future states, *FutureStates*, the association of weights with future states, *HeuristicEstimate*, and the immediately following state, i.e. the one with the greatest weight, *NextState*. Formulae 1-3 in Figure 2 characterise the selection of one action in a specific state, in order to describe a sequence of actions we could add a temporal argument to the predicates or use a temporal logic.

There are clearly many differences in where the knowledge of states, weights etc. is represented at the symbol level. For example, in a single agent planner all knowledge resides in that program. In a multi-agent system such as Option 4, knowledge of future states need only reside in agents which are capable of

Type	Description of set
a	Set of symbols representing states
a^*	Powerset of a
b	Set of integers representing heuristic estimates
c	The set $a * b$, ordered by the value of b
c^*	Powerset of c
	Generate possible future states
1	$(\forall x : a)(State(x) \rightarrow (\exists y : a^*)FutureStates(y))$ Estimate values for future states
2	$(\forall x : a^*)(FutureStates(x) \rightarrow (\exists y : c^*)HeuristicEstimate(y))$ Select the state with the maximum estimate
3	$(\forall x : c^*)(\forall y : c^*)(HeuristicEstimate(x) \wedge Top(x) = y \leftrightarrow NextState(y))$

Fig. 2. A formalisation of the inference structure for a reactive planner

moving, knowledge of actions executed in the past can be stored in the agent which executed that action, but all agents need to know the current positions of all squares (agents) as this is necessary for the weights to be correctly modified when that global state has been visited before.

In order for the RTA* algorithm to operate correctly the estimation function must not overestimate the value of an action [20]. When considering the multi-agent implementation of Option 4 it is clear that each agent must have the same weighting function and must always respect the outcome of any negotiation over which agent can move. If these conditions are not met then the system will not behave according to RTA*. In terms of the knowledge level characterisation, rule 3 will show a contradiction if the next state is not the top element of the set of heuristic estimates.

In this section we have presented a knowledge level characterisation of a reactive planner and discussed four symbol level implementations of it.

3.3 Multiple-blackboard systems

The basic idea of the blackboard model is that there are a number of knowledge sources, viewed as experts, which contribute to the solution of a problem by reading and writing data to a central data structure, the blackboard. Implementational details can be found in [7], of relevance here is the hierarchical organisation of the blackboard and the partitioning of domain knowledge according to content into distinct knowledge sources. These features result from a model-based view of problem solving [7] which includes an explicit categorisation of domain knowledge. Therefore we can, in certain instances, associate the levels of a blackboard with meta-classes and the knowledge sources of the blackboard with knowledge sources in the KL sense.

In the distributed vehicle monitoring testbed (DVMT) of Durfee *et al.* [6] each agent is a blackboard system. The DVMT consists of between 1 and 13 agents whose task is to identify the track of a vehicle from acoustic data sensed

by a number of agents from adjacent or overlapping regions. Durfee *et al.* investigated the effect of communication policies and organisational structure on the efficiency of resource utilisation. The use of planning and the exchange of meta-information were also investigated with regards to improving the local control of agent activity e.g. preventing agents from duplicating work and increasing the priority of potentially rewarding tasks.

Each agent in the DVMT has the same architecture, a blackboard system. More significantly from the knowledge level point of view, each agent has the same global model of the problem solving task, namely the meta-classes and knowledge sources (levels and knowledge sources) of the blackboard. This means that the inferences performed by one agent, for example, one which adds information to the meta-class (blackboard level) 'vehicle location', can be broadcast to other agents who are able to interpret it correctly. This is possible because all agents have the same blackboard levels. All agents share a global, knowledge level view of the problem solving process and their activity is therefore globally coherent by design. There is a marked difference between the contract net and DVMT in this regard, as agents in the contract net have no global view. The activity of DVMT agents may not be optimal as regards the use of computational resources (this is termed 'coherence' in [6]).

The DVMT experiments can be characterised as assuming a global problem solving model, shared by all agents, and investigating a number of symbol level techniques for the assignment of functions to agents - lateral or hierarchical agent organisation, and for the optimal guidance of search within agents, given the global state of problem solving.

3.4 An alternative characterisation in terms of DAI concepts

Bond and Gasser ([1] pages 19-25) define the coherence of a multi-agent system as having four dimensions: solution quality, efficiency, conceptual clarity of system behaviour and graceful degradation⁴. Coherence can be achieved by a number of means, including the following:

- ☐ the assignment of roles to agents
- ☐ planning in order to align agent activity
- ☐ management of communication
- ☐ management of resources
- ☐ data abstraction
- ☐ the use of meta level information

These mechanisms include agent roles - an attributed or conceptual property, planning - a problem solving paradigm, and data abstraction - a problem solving step common to a range of domains. This view of coherence is not a generalised, analytical view, it is a documentation of what occurs in existing DAI systems. This is unsatisfactory if we wish to gain a deeper understanding of DAI systems. Bond and Gasser do not consistently identify these methods with the specific

⁴ It should be noted that the analysis of DAI systems in [1] is of much greater scope than that presented here.

dimensions of coherence that they define, and this detracts further from clarity. They do refer to the specific DAI systems which utilise planning etc. but do not consistently refer to the domain problem, hence the problem solving context is lost. The theoretical generalisations based on this method are weak: coherence is related to resource management and communication, resource management is related to communication, communications can be planned etc. The problem is that the mechanisms of DAI are analysed outwith their context in the solution of problems. The division of the knowledge level from the symbol level proposed in this paper is an attempt to clarify the purpose of the multi-agent system as a whole and in doing so, to distinguish the various roles of the symbol level functions. Such distinctions can play a useful role in the design of distributed systems, as they have done in the model-based design of knowledge-based systems.

4 Related Work

It is noted by Gasser [9] that DAI has taken the agent as the focus of analysis. Two attempts to break from the agent-centered approach are the Open Information Systems (OIS) Semantics of Hewitt [12] and the six principles for social-DAI outlined in [9]. Gasser proposes the existence of multiple actors as a fundamental concept in a 'more social' DAI. The social perspective entails grounding DAI principles in the group rather than the individual agent. The social perspective does not view the autonomy of agents as a concept which is completely divorced from group concepts. Presumably, it should not focus on the mechanisms within agents, but rather on the joint activity of a system of agents.

When we consider the proposals for a 'more social DAI' then it appears that the KL analysis satisfies some of the goals of this enterprise by proposing a global model of behaviour. This model is above the agent level (which we consider to be the symbol level) and this type of characterisation is radically different from the conventional DAI viewpoint(s).

In contrast with [13] we have not sought to add a layer above Newell's knowledge level in order to account for cooperative problem solving. Instead, we have used Clancey's redefinition of the knowledge level to provide the basis of our analysis. This view appears to be consistent with the situated view of knowledge as advocated in [21] as the knowledge level description need not be reflected in specific structures at the symbol level. The analysis of this paper does not lead us to conclude that reasoning about knowledge [11] or planning are necessary features of a knowledge level description. However, in agreement with [13] we do hypothesise a common agent goal which persists through an episode of problem solving. Finally, we acknowledge the importance of planning and reasoning about knowledge as mechanisms for guiding agent activity i.e. as sophisticated control mechanisms. Constructing a plan may also be the purpose of agent activity, in which case we would expect to find knowledge level models of planning. These models may be reactive or deliberative.

The problem of designing the interaction between a knowledge-based system and the user, i.e. a two agent system, from an inference structure is described in

[10]. The method we outline generalises this approach, making use of Clancey's recent review of the knowledge level hypothesis. The problem of implementing a KL description in a blackboard architecture has been addressed in [18] in practical terms. Thus, there are examples which suggest that the analytical approach we advocate can also be used as a basis for the design of multi-agent systems.

5 Conclusions

This paper has presented characterisations of a number of distributed systems at the knowledge level. The systems we have considered have been of a restricted range, distributed knowledge-based systems, and we have emphasised the problem solving domains and solution methods of these systems.

The separation of the knowledge level from the symbol level entails a clear delineation of the roles of symbol level processes and of the concept of agency. Our analysis shows that the global coherence of a number of existing systems can be understood by this analysis. Therefore, we conclude that the knowledge level analysis is a valuable addition to our understanding of distributed information processing systems.

Our conclusions are that the concept of global coherence is best understood as the idealised, rational behaviour of a system of agents. This can be described in terms of the deduction of new knowledge units within a model of problem solving. The model need not specify an optimal solution, nor the intentional attitudes of the agents. It is required that by following the model some solution is identified, and for global coherence, all participating agents must act consistently with the model. It is not necessary that the model be represented within the agents themselves, however, this could produce a flexible exchange of roles among agents by enabling agents to reason about their own role and that of others.

References

1. Bond, A.H. and Gasser, L. (1988) An analysis of problems and research in DAI. in *Readings in distributed artificial intelligence*. (eds) Bond, A.H. and Gasser, L., Morgan Kaufmann, 1988 :333-356
2. Clancey, W.J. (1985) Heuristic Classification. *Artificial Intelligence* 27 (1985) :289-350
3. Clancey, W.J. (1991) The frame of reference problem in the design of intelligent machines. in (ed) Vanlehn, K. *Architectures for intelligence*. Lawrence Erlbaum Associates, 1991
4. Davis, R. and Smith, R.G. (1983) Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence* 20, 1983 :63-109
5. Drogoul, A. and Dubreuil, C. (1991) Eco-problem-solving model: Results of the N-puzzle. *Proc. 3rd European workshop on modelling autonomous agents and multi-agent worlds (MAAMAW 1991)* (eds) Steiner, D.D. and Müller, J., D.F.K.I. report No. D-91-10
6. Durfee, E.H. Lesser, V.R. and Corkhill, D.D. (1987) Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers* C-36 :1275-1291

7. Englemore, R.S. (1988) *Backboard Systems Addison Wesley 1988*
8. Fagin, R. and Halpern, J.Y. (1985) Belief, awareness and limited reasoning. *Proc. IJCAI 1985 :491-501*
9. Gasser, L. (1991) Social conceptions of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence 47 (1991) :107-138*
10. de Greef, P. and Breuker, J.A. (1992) Analysing system-user cooperation in KADS. *Knowledge Acquisition (1992) 4 :89-108*
11. Halpern, J.Y. and Moses, Y. (1990) Knowledge and common knowledge in a distributed environment. *Journal of the ACM, Vol 37. No. 3 :549-587*
12. Hewitt, C. (1991) Open information systems semantics for distributed artificial intelligence. *Artificial Intelligence 47 (1991) :79-106*
13. Jennings N.R. (1992) Towards a cooperation knowledge level for collaborative problem solving. *Proc. ECAI 1992 :224-228*
14. Konolige, K. (1986) A deduction model of belief. *Pitman, London, 1986*
15. Knight, K. (1993) Are many reactive agents better than a few deliberative ones ? *Proc. IJCAI 1993 :432-437*
16. Korf, R. (1990) Real-time heuristic search. *Artificial Intelligence 42(2-3) 1990*
17. Lenting, J.H.J. and Braspenning, P.J. (1993) Delegated negotiation for resource re-allocation. in *Lecture Notes in A.I. No. 671, (ed) Ohlbach, H.J., Springer Verlag, 1993 :299-311*
18. Major, N. Cupit, J. and Shadbolt, N. (1994) Applying the REKAP methodology to situation assessment. *Proc. 4th KADS Meeting, GMD, Sankt Augustin, Germany 24-25 March 1994, (eds) Voss, H. and Studer, R. Arbeitspapiere der GMD 832*
19. Newell, A. (1982) The knowledge level. *Artificial Intelligence 18 (1982):87-127*
20. Reinfeld, A. (1993) Complete solution of the eight-puzzle and the benefit of node ordering in IDA*. *Proc. IJCAI 1993 :248-253*
21. Rosenschein, S.J. (1985) Formal theories of knowledge in AI and robotics. *New Generation Computing 3(1985) :345-357*
22. Smith, R.G. (1977) The contract net: a formalism for the control of distributed problem solving. *Proc. IJCAI 1977 :472*
23. Smith, R.G. (1980) The contract net protocol: high level communication and control in a distributed problem solver. *IEEE Trans. on computers Vol. C-29 No. 12 December 1980 :1104-1113*
24. Werner, E. (1990) What can agents do together ? A semantics for reasoning about cooperative ability. *Proc. ECAI 1990 :694-701*
25. Wielinga, B.J. Schreiber, A.T. and Breuker, J.A. (1992) KADS: a modelling approach to knowledge engineering. *Knowledge Acquisition (1992) 4, :5-53*
26. Wooldridge, M. (1994) Coherent social action. *Proc. ECAI 1994 :279-283*