

The Link between Autonomy and Organisation in Multiagent Systems

Michael Schillo, Klaus Fischer, Jörg Siekmann

German Research Center for Artificial Intelligence (DFKI),
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
{kuf, schillo, siekmann}@dfki.de

Abstract. Market-based approaches have a long tradition in supporting of task-assignment multiagent systems. Such systems consist of customer agents with jobs to assign, and provider agents that have the resources to perform these jobs. Jobs can be complex in the sense that they require the collaboration of several provider agents. We present a set of organisational forms of collaboration between firms that have the potential to increase performance through the structure they impose. This gain of structure, which comes with a loss of autonomy of the individual agents, is especially valuable in settings where communication has to be limited.

1 Introduction

The central setting for this paper is a market of two kinds of agents, *customers* and *providers*. The customer agents have jobs they want to be done by provider agents. These jobs may require more types of resources than a single agent can provide, hence providers need to collaborate. All agents are self-interested entities that do not necessarily have a common goal; we allow them to be designed and owned by different parties, which limits the possibilities for global control in this setting. In addition, we assume that provider agents do not know anything about the future orders of the customers.

If the system contains a large set of agents or the jobs to be assigned require the collaboration of many providers, assigning the jobs with standard auction systems is very time consuming and methods to improve the performance of the system become more important. If the same type of job needs to be assigned many times or parts of the jobs are the same over and over again, the system is be more efficient if this repeating structure on the demand side is reflected by an anticipating structure on the side of the provider's. That is, providers who are successful at completing a job should form relationships that facilitate long-term teamwork. This provider grouping can be formalized by the concept of *organisation*. Organisations are social structures that provide processes for conflict resolution, which results from previously resolved problems or conflicts [7]. They institutionalise anticipated coordination, which is especially useful for medium- and large-scale applications that require the limitation of the agents' communication behaviour.

Jennings writes that "the development of robust and scalable software systems requires autonomous agents that can complete their objectives while situated in a dynamic and uncertain environment, that can engage in rich, high-level social interactions, and

that can operate within flexible organisational structures” [11]. Agents acting in an organisational structure can encapsulate complexity of the subsystems (simplifying representation and design) and modularize functionality (providing the basis for rapid development and incremental deployment).

To model these aspects, we use the concept of a *holonic agent* or *holon* as introduced in [6, 8]. The concept is inspired by the idea of recursive or self-similar structures in biological systems [13]. A *superholon* consists of parts called *subholons*, which in turn may be superholons themselves, thus introducing recursion as a modelling technique. Any holon that is part of a whole is thought to contribute to achieving the goals of this superior whole. To the outside, each holon is represented by a distinguished *head* which co-ordinates the activities of the holon. Apart from the head, each holon consist of a (possibly empty) set of other agents, called *body agents*. The holonic agent may have capabilities that emerge from the composition of its agents and it may have actions at its disposal that none of its agents could perform alone. Body agents can give up part of their autonomy to the holon to enhance it’s overall performance. Committing to the participation in a holon, agents limit their possible future actions and therefore give up part of their autonomy. The degree to which they give up their autonomy is not fixed in advance but depends on the circumstances and is subject to negotiation between the agents participating in a holon. The least sacrifice of autonomy can be seen in holons forming a loose *federation* of agents. The long-term commitment in this form is actually so low that agents need to negotiate their coordination on a case-by-case basis. The federation is at one end of the autonomy spectrum and does not differ significantly from conventional multiagent systems (it merely institutionalises preference structures). At the other end of the spectrum, agents can give up all of their autonomy and merge into a single agent. Between the two extremes, there can be hybrid forms of different nuances.

However, up to this point little work has been done to elaborate these nuances. The contribution of this paper is to define a number of such hybrid forms motivated by inter-organisational networks found in human societies (Section 2) and investigate their properties with respect to the autonomy of involved agents (Section 3).

2 Organisational Networks

2.1 The Matrix of Delegation - A Grammar for MAS Organisation

Recent work on delegation (see e.g. [4] for an extensive treatment) shows that this is an interesting concept highly relevant for multiagent systems. The mechanism of delegation makes it possible to pass on tasks (e.g. creating a plan for a certain goal, extracting information) to other individuals and furthermore, it allows for the specialization of these individuals for certain tasks (functional differentiation and role performance, etc.). Representing groups or teams is also an essential mechanism for social processes of organisation, coordination and structuring. We distinguish two types of delegation: task delegation and social delegation. We call the procedure of appointing an agent as representative for a group of agents *social delegation*.

Social delegation is in several respects different from the well-known task delegation. For example it involves a possibly long-term dependency between delegate and represented agent, and the fact that another agent speaks for the represented agent may

incur commitments in the future, that are not under the control of the represented agent. Social delegation is more concerned with the delegate performing a certain role, than with producing a specified product. In holonic terms, representation is the job of the head, which can also be distributed according to a set of tasks to different agents. Just like fat trees (multiple bypasses to critical communication channels) in massive parallel computing, distributing the task of communicating to the outside is able to resolve bottlenecks. This makes social delegation a principle action in the context of flexible holons and provides the basic functionality for self-organisation and decentralized control.

Thus, we believe it is justified to differentiate two types of delegation: task delegation, which is the delegation of (autistic, non-social) goals to be achieved and social delegation, which does not create a solution or a product but represents a set of agents. Both types of delegation are essential for organisations, as they support independence from particular individuals through task and social delegation.

	Task Delegation	Social Delegation
Economic Exchange		
Gift Exchange		
Authority		
Voting		

Table 1. The delegation matrix showing two modes of delegation and four mechanisms for performing each mode. Theoretically, every combination of mode and mechanism is possible in multiagent organisation.

Given the two types of delegation, it remains to be shown how the action of delegation is performed. We observe four distinct mechanisms for delegation (see also Table 1):

- (i) Economic exchange is the standard mode in markets: the delegate is paid for the delegated task or representation. In economic exchange, a task is exchanged for money, where the parties involved assume that the value of both is of appropriate similarity.
- (ii) Gift exchange, as an important sociological mechanism [1], is a deliberate deviation from the economic exchange. The motivation for the gift exchange is the expectation of either reciprocation or the refusal of reciprocation. Both are indications about the state of the relationship between the involved parties. This kind of exchange entails risk, trust, and the possibility of conflicts (continually no reciprocation) and the need for an explicit management of relationships within an agent. The aim of this mechanism is to accumulate strength in a relationship that may pay off in the future.
- (iii) Authority is a well known mechanism, it represents the method of organisation used in distributed problem solving. It implies a non-cyclic set of power relationships between agents, along which delegation is performed. However, in our framework authority relationships are not determined at design time, but they are the result of an agent deciding at runtime to give up autonomy and to allow another agent to exert power.
- (iv) Another well-known mechanism is voting, whereby a number of equals determine that one of them is the delegate by some voting mechanism (majority, two thirds, etc.). Description of the mandate (permissions and obligations) and the particular circumstances of the voting mechanism (registering of candidates, quorum) are integral parts

of the operational description of this mechanism and must be accessible to all participants.

As suggested in Table 1, all four mechanisms work for both types of delegation: for example, economic exchange can be used for social delegation as well as for task delegation. This set of mechanisms is not necessarily complete, however, many mechanisms observed in human organisations that seem not to be covered here, are a combination of the above mechanisms.

2.2 The Spectrum of Organisational Forms

In general, we allow agents to be members of several organisations at the same time. In order to unambiguously determine which organisation is responsible for an incoming order, this general rule is restricted to all organisations an agent is engaged in being created for different types of orders. We will now describe (building on the matrix of delegation from the previous section) seven different forms of organisation and non-organisation for MAS in the order of increasing coupling between agents along a spectrum. Organisations may differ as agents interact either in a cooperative, in a competitive or in a authoritarian way (c.f. for example [14]) The names for the different forms are derived from the types of firms that are typically investigated in the field of organisational sociology.

Single, Autonomous Agents: This form of coordination is not particularly relevant but it provides the theoretical starting point, with fully uncoupled agents. All agents that provide services do not interact with each other to accomplish their tasks, the only interaction is between providers and customers.

Market: In the market-style interaction, agents exchange jobs directly and there is some kind of payoff (here represented as money). This does not necessarily imply that agents build a relationship or an organisation in the strict sense, as interaction is short term, case by case based. The provider agent that re-delegates parts of a job acts as the holon head for this specific job.

Virtual Enterprise: The virtual enterprise is a temporary network of legally independent companies to share skills, costs and access to each other's market. Virtual enterprises promise to offer the best of both worlds, flexibility and economy of scale. They are networks of legally and economically independent enterprises, each concentrating on its core competencies and out-sourcing the rest, modeled on the best-of-breed organisation. The virtual enterprise appears and acts like a single enterprise to the outside has to do world [2]. Moreover, there is no physical institutionalization of central management functions. The contract defining the relationship between the participating enterprises is deliberately left loose, in order to facilitate quick formation and greater flexibility in re-organisation. In our model, a virtual enterprise consists of provider agents with equal rights, there is no single designated head agent. A virtual enterprise is product-specific. Each member agent may accept jobs, but must start a new internal auction for each of its subtypes among its partners. This member agent becomes the head of the virtual enterprise for this specific job, other members may be heads of the holon for other jobs. There is no specific profit distribution other than the normal negotiation in the course of the internal auctions.

Alliance: An alliance as an organisational type is different than the virtual enterprise because of a long term contract between the participants that regulates a closer cooperation [9]. The relationship between the companies is formalized by a contract, which is the result of negotiation between the different companies. Alliances are not fully integrated economically and legally, therefore the profit distribution for all internal transactions is regulated in advance. Alliances are founded in order to create at least one new product. As the companies are only partially integrated they usually supply other products apart from the alliance as well. Thus, they are generally allowed to join other organisations apart from those which produce the same product as the alliance. As alliances are in some way legally integrated they need to appoint at least one CEO (representative), which is done by voting. The representation of the alliance incurs valuable reputation and contact to customer agents, hence it implies (economic) power. Quitting of one of the agents with many customer contacts may cause loss to the organisation, as customers may prefer to interact with the provider agent they already are acquainted with, no matter in which organisation it is in. To decrease the incentive to join the alliance solely for this purpose and for the stability of the organisation, and a focal participant, who is, due to his already powerful position, not reliant on this increase in reputation, is appointed by *social delegation* through *voting* to represent the alliance. The profit is distributed among the head (representative) and all body agents necessary for performing the task by using *economic exchange* and *gift exchange*. However on creation of the alliance agents agree on a ratio (which is in our case fixed by the designer) that describes how profit is split between the head agent and the body agents that are involved in performing the task.

Strategic Network: Strategic networks differ from virtual enterprises in that they use stronger legal contracts, and feature a *hub firm* that sets up the network, and takes proactively care of it [10]. The hub firm in a strategic network is usually significantly larger than the other members of the network. It coordinates activities in the strategic network, but the members retain their legal independence and autonomy. This network arrangement allows a participating firm to specialize in those activities of the value chain that are essential to its competitive advantage, reaping all the benefits of specialization, focus, and, possibly, size. The time frame and financial volume are usually larger than in the case of virtual enterprises, but firms have still the right to leave the network.

In our model, strategic networks consist of a head agent and body agents. If an incoming order matches the product of the strategic network, the rules that apply to the receiving agent discriminate whether it is the head or a body agent. Body agents may not directly accept a *call for proposals* (*cfp*) from outside, but must *bounce* it. Bouncing means that they refuse the order, but they send the name of their head inside the refusal message so the sender can resend the *cfp* to the head instead. The semantics of bouncing is that the head of the organisation is the one responsible for the organisation's interaction to the outside and all inter-organisational communication must be channeled through him. Heads can accept orders from the outside. They know about their body agents' schedules and resources, and can instruct them to do a job at any given time. Strategic networks are product-specific, so multiple memberships are allowed. The profit distribution is according to a fixed ratio.

Group: Groups are formed from enterprises that retain their legal independence, but are bound by contract to the authority of the central firm. Here, we mean group as the organisational structure of a firm as in "Bertelsmann group", not in the socio-psychological meaning of "team". In contrast to the strategic network, no multiple memberships are allowed, and usually there is no exit option for subordinate firms. All economic activities are focused on the group and subject to directions from the head enterprise. The interdependency between the firms is found in an authoritative hierarchy.

In our model, an agent who is a member of a group is not allowed to be a member of any other organisation. Body agents have to bounce incoming orders. Head agents may order body agents to do a specific job. This inclusion of all economic activity in the group results in the head agent always being up-to-date about its body agents' resource allocations. The head agent retains all the profit for orders completed by the group. Every round, it pays a fixed amount of money to each body agent.

Corporation: A corporation is the result of the complete inclusion of all legal and economic aspects of the original companies into a new entity. This organisational form marks the other extreme of the spectrum between market and hierarchy. Companies merging into a corporation give up all their autonomy. The process is usually not reversible; once inside a corporation, the former status cannot be regained. In the business world, the process of merging usually happens when a large company assimilates a much smaller one. We model corporations by letting the head assimilate the resources of its body agents. After the assimilation, the body agents are removed from the simulation. The head then acts like a normal single agent, except that it does not form new organisations.

We presented here multiagent organisations starting with the most autonomous form and proceeded to the one with least autonomy. The model provides a framework for the agents' decision at runtime. In theory, each agent can choose, depending on the situation in the MAS, whether it is in its interest to change its current status. As each organisational type has advantages and disadvantages, it may well be that a transition is not beneficial in the light of the current market situation.

2.3 Synopsis

Before we discuss the effect of the model on agent autonomy, we will explain the summarizing table of the critical features of the organisational forms. Table 2 gives a synopsis on the organisational types, it characterizes organisational forms in terms of their properties: the mechanism for task delegation (TD), social delegation (SD), membership limitations (M), profit distribution (PD), and the role of the holon head (HH).

Membership limitations can have the value "limitation on product", which means that the agent is free to choose other organisations to join, as long as they do not use the same set of resources. This parameter can also denote that there is no limitation (as in the market) or that an agent is only allowed to be a member of one single organisation (as with the group).

Profit distribution is on a per job basis using economic exchange, a fixed ratio between head and body agents (e.g. 20:80), or a fixed income which is paid by the head to the body agents regardless of the number of jobs performed (in this case, variable costs are paid by the head plus a fixed income chosen by the designer). The details of

	Market	Virtual Enterprise	Alliance	Strategic Network	Group	Corporation
M	No Limitation	Limitation on Product	Limitation on Product	Exclusive Membership	Exclusive Membership	N/A
PD	Case by case	Case by case	Regulation	Fixed Income	N/A	N/A
HH	One/All	All	One	One	One	One
TD	Economic Exchange	Economic/ Gift Exchange	Economic/ Gift Exchange	Authority	Authority	N/A
SD	Economic Exchange	Economic/ Gift Exchange	Voting	Authority	Authority	Authority
P	HCNCP	HCNCP	DCP	DCP	DP	None

Table 2. Overview of the five types of holonic organisation. Rows specify for each type the dependence introduced (Dep), membership limitations (M), the mode of profit distribution (PD), the role of the holon head (HH), and the protocol used (P).

the fixed ratio and fixed income must be agreed on at the time of creating the organisation. There is no profit distribution in the corporation, because the original agents have merged into a single entity.

The role or number of the holon heads (HH) for the market is to some extent up to interpretation: Although there is only one job per holon and only one agent communicating and coordinating for this holon, all agents in the system are allowed to accept jobs and then engage in coordination and communication. In the virtual enterprise all agents can receive incoming jobs and redistribute them. All other forms allow only a single point of access to the outside world. As shown in the table, depending on the organisational type we use three different protocols. The mechanisms for task delegation (TD) and for social delegation (SD) have been described in Section 2.1. We will come back later to the protocol used for task-assignment internal to the organisation (P).

3 Implications of organisational structure on agent autonomy

According to Castelfranchi [3], autonomy is not only a quantitative dimension (a question of more or less), but divided into several dimensions and it directly corresponds to (the lack of) dependence in some aspect. Some of these dimensions map precisely to the differences between organisations in our discussion. Along the spectrum we will now identify the dimensions of autonomy that are given up in the different organisational forms. Our theoretical starting point are the single, non-cooperating agents, which are fully autonomous.

Skill and resource autonomy: The agent that enters a pure market relationship divides a job and re-delegates those parts it cannot perform by itself. Therefore, it is dependent on the *skill* and *resources* of other agents, and thus loses skill and resource autonomy (emphasized words correspond to Castelfranchi’s terminology).

Conditional autonomy: When creating the virtual enterprise the agents specify which compound product this organisation is designed for. Entering the organisation implies that the agent commits to provide its resources for this compound product and that

any incoming orders will only be shared with other agents in the organisation. Hence, it may only collaborate with other agents under the condition that the result of this collaboration is not of the same type as the product of the organisation. Agents cannot arbitrarily choose collaborators, and therefore give up *conditional autonomy*.

Representational autonomy: The creation of an alliance requires the election of a single representative, which will fully take over communication to the outside of the organisation. No other agent is allowed to accept orders from other agents, they may not represent themselves, and hence lose *representational autonomy*.

Goal dynamics autonomy: Agents in a strategic network must announce the cost function and are subordinates to a single representative, which has the power to order them to perform tasks. This means that they cannot fully influence their own set of goals, and they lose *goal dynamics autonomy*.

Planning autonomy: The representative of a group has complete knowledge of the schedule of its subordinate agents. Therefore, it is not necessary to communicate for determining an agent that can perform a task. Planning can be done centrally by the representative, which then merely informs subordinates about when to perform which task. These agents have lost *planning autonomy*.

Processing autonomy: When agents have merged into a corporation agent, they give up individual computational resources and lose the last remaining autonomy, namely *processing autonomy*.

It is worth noting that each form of organisation builds on the previously described form, and introduces new dependencies corresponding to a loss of autonomy of the body agent. Therefore, we can speak of a total ordering of the organisational type and hence, a *spectrum* of organisational types.

This discussion is not just academic, but it has practical impact: As the basis for task-assignment we apply the *Holonic Contract Net with Confirmation Protocol (HCNCP)*. The HCNCP (see Figure 1) uses the standards of the *Foundation for Intelligent Physical Agents (FIPA)* as a reference (cf. [5]). It extends Smith's contract-net protocol (CNET) [15] and avoids the problem of committing too early, which often leads to suboptimal outcomes in the CNET, and can be applied in a cascading manner (for more detailed discussions see [12]). In the best case, the HCNCP requires six messages between a contractor and a contractee to assign a job. The HCNCP is used for inter-organisational communication, market interaction, and intra-organisational communication in the organisation form *virtual enterprise*. However, if agents are inside an *alliance* or a *strategic network* the initiator of an intra-organisational task-assignment protocol can save sending a cfp as a consequence of the loss of representational autonomy. It can directly proceed to the request for confirmation of being able to perform the task (Direction with confirmation protocol, DCP, cf. Figure 2). Although there is an authority relationship between representative and all other agents, they may still be a member of several organisations, and hence, the schedule of an agent is not fully known by the representative. In the group however, this is the case, and therefore a further part of the task assignment protocol can be saved (Direction protocol, DP, cf. Figure 3). Obviously, the different protocols have implications on the overall communication effort and are a direct result of the different levels of autonomy. In order to evaluate our organisational structures, we started an empirical study of the communication patterns

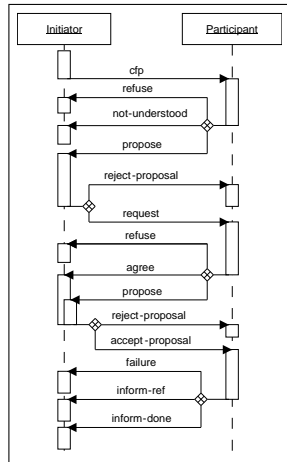


Fig. 1. The Holonic Contract-Net with Confirmation Protocol (HCNCP) which is used as the default protocol for efficient and cascading task assignment.

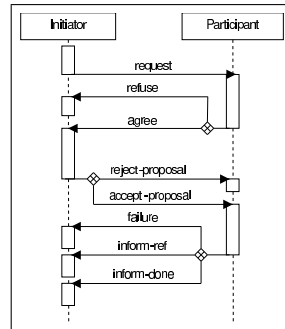


Fig. 2. The Direction with Confirmation Protocol (DCP) which is used in the strategic network, saves the announcement of cost proposals.




Fig. 3. The Direction Protocol (DP) which is used in the group, makes use of the knowledge of the initiator and the authority relationship.

inside of each of these organisations given a fixed task. In a round-based simulation, customers were created to announce a task that can consisted of three subtasks of types A, B, and C. For each organisation we tested one scenario, which contains 20 provider agents for each of the three types, each type occurring with production costs 4, 5, and 6 (a total of 180 provider agents). There are 60 customer agents emitting one order of type ABC each round. In each configuration, all occurring organisations are of the same form. Each organisation has three member agents, one for each of the three task types. One member agent has cost 4, one cost 5, and one cost 6. Each experiment ran for one hundred rounds. In this setting the aforementioned differences result in a diversification of the amount of communication required by each organisational form for solving the same task-assignment problem. For this setting, the number of messages for the virtual enterprise (ca. 14000) is higher than that for the strategic network (ca. 4200), which in turn is higher than that for the group (3.200). Corporation scenarios have the lowest number of messages (ca. 2200). The number of messages for agents in a pure market relationship is one order of magnitude higher (ca. 256000).

4 Conclusion

Multiagent systems have the power to model forms of collaboration inspired by real world organisations in a natural fashion and we contributed by laying out a framework of different organisational forms relative to the following parameters: mechanisms for task and social delegation, membership, profit distribution, number of representatives and protocol for task-assignment. This theory of how agents of different

capabilities can be tied together is inspired by wellknown sociological descriptions of inter-organisational networks and involves a distinct description of dimensions of autonomy and dependence agents have in these networks. While increasing coupling between agents increases performance in terms of numbers of messages, this coupling comes at a price. It requires the loss of skill/resource autonomy, goal dynamics autonomy, representational, planning, but be and finally, processing autonomy. So, the choice of an organisational form for a concrete application not only depends on the performance requirements but also on the necessity of privacy and the extent to which local decision-making is a requirement, a decision that involves considering a complex trade-off.

5 Acknowledgements

This work was funded by the Deutsche Forschungsgemeinschaft in the priority program *Sozionik* under contract Si 372/9-2. We are indebted to sociologists Dr. Michael Florian, Dr. Frank Hillebrandt, Bettina Fley and Daniela Spresny for many inspiring discussions. We also thank the anonymous reviewers of the conference for comments and advice.

References

1. P. Bourdieu. *Pascalian Meditations*. Polity Press, Cambridge, Eng., Oxford, Eng., 2000.
2. L.M. Camarinha-Matos, H. Afsarmanesh, C. Garita, and C. Lima. Towards an architecture for virtual enterprises. *Journal of Intelligent Manufacturing*, 9(2):189–199, 1998.
3. C. Castelfranchi. Founding agent's "autonomy" on dependence theory. In *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI 2000)*, pages 353–357, 2000.
4. C. Castelfranchi and R. Falcone. Towards a theory of delegation for agent-based systems. *Robotics and Autonomous Systems*, 24:141–157, 1998.
5. FIPA. Foundation for Intelligent Agents <http://www.fipa.org/repository/ips.html>, 2002.
6. K. Fischer. Holonic multiagent systems — theory and applications. In *Proceedings of the 9th Portuguese Conference on Progress in Artificial Intelligence (EPIA-99)*, LNAI Volume 1695, pages 34–48. Springer-Verlag, 1999.
7. L. Gasser. Social conceptions of knowledge and action: DAI foundations and open systems semantics. *Artificial Intelligence*, 47:107–138, 1991.
8. C. Gerber, J. Siekmann, and G. Vierke. Flexible autonomy in holonic multiagent systems. In *AAAI Spring Symposium on Agents with Adjustable Autonomy*, 1999.
9. R. Gulati and Martin Garguilo. Where do interorganizational networks come from? *American Journal of Sociology*, 104:1439–1493, 1999.
10. J. Carlos Jarillo. On strategic networks. *Strategic Management Journal*, 9:31–41, 1988.
11. N.R. Jennings. Agent-based computing: Promise and perils. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 1429–1436, 1999.
12. T. Knabe, M. Schillo, and K. Fischer. Improvements to the FIPA contract net protocol for performance increase and cascading applications. In *International Workshop for Multi-Agent Interoperability at the German Conference on AI (KI-2002)*, 2002.
13. A. Koestler. *The Ghost in the Machine*. Hutchinson & Co, London, 1967.
14. Walter W. Powell. Neither market nor hierarchy. network forms of organization. *Research in Organizational Behavior*, 12:295–336, 1990.
15. R. G. Smith. The contract net: A formalism for the control of distributed problem solving. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI-77)*, page 472, 1977.