Modelling Decentralised Decision Making

Love Ekenberg

The DECIDE Research Group
Department of Computer and Systems Sciences
Stockholm University and Royal Institute of Technology
Electrum 230, S-164 40 KISTA, SWEDEN

Abstract

A set of agents (human or computer process) are supposed to decide on an action given a finite set of possible strategies when the background information is vague or numerically imprecise. The agents evaluate the strategies and communicate their opinions about their utilities, their opinions about the credibilities of the reports of the other agents, as well as their opinions of the reliabilities of the other agents. A significant property of the theory is that it admits the representation of imprecise information at all stages. The strategies are evaluated relative to generalisations of the principle of maximising the expected utility. It is also demonstrated how decision trees can be integrated into the framework and we suggest an evaluation rule taking into account all strategies, credibilities, reliabilities, probabilities and utilities involved in the framework. This is done without presupposing the existence of a central co-ordinating agent that evaluates the different reports of the agents. In a distributed context there are no *a priori* reasons for such an agent and we show how it can be eliminated in our framework.

Background

In the area of Distributed Artificial Intelligence multi-agent co-ordination is approached in different ways depending on what relations sets of agents have to each other. One aspect of this is to what extent a group of agents is centrally co-ordinated (Ekenberg, Boman, & Danielson 1995), (Ekenberg, Danielson, & Boman 1995), (Katz & Rosenschein 1993), (Rosenschein 1982) and to what extent the control is distributed (Ephrati & Rosenschein 1991), (Zlotkin & Rosenschein 1991). Intermediate forms occur, where agents sub-contract their tasks to other agents (Kraus 1993); also in environments where the agents do not have a common goal and there is no globally consistent knowledge. Within the economic approaches, the Contract Net communication protocol (Smith 1980) has also been appreciated.

The approach in (Ekenberg, Danielson, & Boman 1996a), (Ekenberg, Danielson, & Boman 1996b) treats a centralised scenario where a decision making agent (DMA), which may be a human coordinator as well as another agent process, faces a situation involving a choice between a finite set of strategies $\{S_i\}$ having access to a finite set of autonomous agents $\{A_i\}$ reporting their opinions on the strategies to the DMA, see Figure 1. In such a scenario, the communication model required is restricted to the information flow between the DMA and the respective agents.

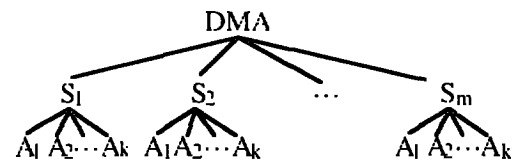

Figure 1 **A multi-agent decision model**

In a situation modelled as in Figure 1, some agents may be more reliable than others when evaluating the strategies involved, since different agents may have different capabilities to determine the utilities. The DMA also has access to assessments expressing the credibility of the different agents and is set on choosing the most preferred strategy given the agents' individual reports and the relative credibility of these. A main issue in such a context is how the local utilities of the agents are related to the global utility of the complete system (cf. (Rosenschein & Zlotkin 1994)). A common assumption in this area of the processes involved is that each of the participating agents has an individual preference relation over the set of strategies. Agents bypass negotiation by using a voting mechanism; each agent expresses its preferences and a group choice mechanism is used to select the result. For instance, (Ephrati & Rosenschein 1991), (Zlotkin & Rosenschein 1991) use a refinement of the Clarke tax (Clarke 1971). This is further developed in (Ephrati & Rosenschein 1993).

A feature in, e.g., (Ephrati & Rosenschein 1991) is that the utilities are communicated explicitly and are not dependent on particular local functions in defining a global (social welfare) function. Similar to this, the model in Figure 1 does not necessarily require particular utility functions of the agents, but rather that the utility estimates are expressed by sets of utility functions, determined by different constraints (Ekenberg, Danielson, & Boman 1996a). The decision rule used in agent deliberation has often been the principle of maximising the expected utility (PMEU) which has been equated with the concept of rationality (cf. (Gmytrasiewicz & Durfee

1993)). This equality is inspired by research that grew from efforts of Ramsey, von Neumann, and others (Fishburn 1981). It should be emphasized that even though this paper assumes the underlying use of the PMEU, this assumption can be relaxed (Boman & Ekenberg 1995), (Ekenberg, Danielson, & Boman 1996b).

Another important issue is that for a set of agents to perform adequately, it seems important that the agents are able to evaluate vague background information gathered from different sources. The dynamic adaptation taking place over time as the agents interact with their environment, and with other agents, is affected by the means available to assess and evaluate imprecise information. In this respect, our concern has been how problems modelled by numerically imprecise probabilities and utilities can be evaluated. Much work has been done in this area and the relation between our work and earlier approaches that attempt to incorporate imprecise estimates into probability theory is described in (Ekenberg & Danielson 1994). Our approach attempts to conform to traditional statistical reasoning by using the concept of admissibility (Lehmann 1959). The theory may still be used in vague domains and it avoids restricting assumptions such as common knowledge, precise utility measures, and a dependence on axiomatic characterisations (Rosenschein 1988).

An obvious disadvantage with the model in Figure 1 is that there are not always *a priori* reasons for choosing a specific DMA, and one purpose with this paper is to allow for the DMA to be removed. The next section presents an approach where the agents have access to a mechanism for communicating constraints on their individual utilities and a decision rule that evaluate a global utility from this. Since such a structure requires qualified reports from each of the agents, the agents should also have access to a general framework for evaluating the possible strategies from their own viewpoints. Such a framework is investigated in the second part of this paper, which describes how agents may use decision trees to evaluate different strategies with respect to imprecise background information.

Global Utility

The first step to eliminate the DMA is to distribute the possibility to estimate the credibility of the agents. As in the model in Figure 1, each individual agent may express its respective opinion about the utilities of the strategies under consideration. However, it may still be that different agents may have different capabilities to determine the utilities, and this section suggests how the possibility to make credibility assessments can be distributed.

Representation

The agents are asked to assess the competence of the other agents in two respects: (i) they may have opinions on the other agents' abilities to give adequate estimates on the utilities of different strategies, and (ii) they may have opinions on the reliabilities of the other agents as they in turn assess the abilities of other agents. Each agent has such capabilities of expression since there are no coordinating agents or global credibilities involved. This model is restricted in the sense that there are no good answers to why we should stop here. The agents could also have opinions on the reliability of the other agents when they in turn assess the reliability of the other agents. Furthermore the agents may have opinions on the abilities of the agents when stating the latter kinds of assessments, etc. Theoretically, this viewpoint implies an infinite regress and makes the problem intractable, but a reasonable, albeit unjustified, assumption is that the assessments converge after a few levels. The description in this section is restricted to a case with three kinds of assessments, but as will be seen later, the model can be extended to take account of an arbitrary number of assessments types. As an example how this model can be used, consider the following:

Example: A set of agents (A_1 and A_2) are to report on their respective assessments concerning the strategies for a risk policy of a company. The objective for the agents is to decide how to allocate resources for preventing potential losses of the company. The prevailing strategies are to prevent disruption of productions and services, to prevent obstruction of research and development, or to distribute the resources over both these objectives. These strategies are labeled S_1, S_2, and S_3 below. Assume that the agents A_1 and A_2 have reported the following utility assessments. The utilities involved could, for example, be monetary values. In that case, they are linearly transformed to real values in the interval [0,1].

For instance, the assessments according to agent A_1 could be the following:

- The utility of strategy S_1 is between 0.20 and 0.50
- The utility of strategy S_2 is between 0.20 and 0.60
- The utility of strategy S_3 is between 0.40 and 0.60
- The utility of strategy S_2 is at least 0.10 better than that of S_1

Agent A_2 can state similar assessments about the utilities of the strategies. Note that the utility estimates are treated only from a global perspective. Below, we will discuss how the agents may use decision trees in imprecise domains to individually perform evaluations of the strategies. (Risk evaluation methods that can be integrated into the framework are discussed in (Ekenberg & Danielson 1995), (Ekenberg, Oberoi, & Orci 1995)).

Moreover, the agents may estimate the credibility of A_1 and A_2 as numbers in the interval [0,1]. The number 0 denotes the lowest credibility and 1 the highest. Thus, the assessments according to agent A_1 could be:

- The credibility of agent A_2 is at least equal to that of A_1
- The credibility of agent A_2 is between 0.30 and 0.70

The assessments according to agent A_2 have a similar form. Furthermore, the agents have estimated the reliabilities of A_1 and A_2 as numbers in the interval [0,1]. The number 0 denotes the lowest reliability and 1 the highest. For instance, the assessments according to agent A_1 can be:

- The reliability of agent A_1 is between 0.50 and 0.70
- The reliability of agent A_2 is between 0.20 and 0.50 ■

One further reason for allowing interval as well as comparative assessments is that the agents' information may have different sources. For instance, intervals naturally occur from aggregated quantitative information while qualitative considerations often result in comparisons. Since the sources may be different, the assessments are not necessarily consistent with each other.

Example (cont'd): The reports provided by agent A_1 are translated to the following expressions.

$\mathbf{u_{11}} \in [0.20, 0.50]$
$\mathbf{u_{21}} \in [0.20, 0.60]$
$\mathbf{u_{31}} \in [0.40, 0.60]$
$\mathbf{u_{21}} \geq \mathbf{u_{11}} + 0.10$

The credibilities as well as the reliabilities of A_1 and A_2 are also represented as numbers in the interval [0,1], and the translation of the credibility and the reliability assessments above results in the following expressions.

$\mathbf{c_{12}} \in [0.30, 0.70]$
$\mathbf{c_{12}} \geq \mathbf{c_{11}}$
$\mathbf{r_{11}} \subset [0.50, 0.70]$
$\mathbf{r_{12}} \in [0.20, 0.50]$ ■

The sets of assessments are transformed into linear systems of equations that are checked for consistency. The credibility assessments, reliability assessments, and utility assessments constitute the *credibility base* $\mathcal{K}$, the *reliability base* $\mathcal{R}$, and the *strategy base* $\mathcal{S}$ respectively. A credibility base with k agents is expressed in the credibility variables $\{c_{11},\ldots,c_{1k},\ldots,c_{k1},\ldots,c_{kk}\}$ stating the relative credibility of the different agents when considering the strategies. The term c_{ij} denotes the opinion of agent A_i about the credibility of agent A_j's utility estimates. A reliability base with k agents is expressed in the reliability variables $\{r_{11},\ldots,r_{1k},\ldots,r_{k1},\ldots,r_{kk}\}$ stating the relative reliability of the different agents. The term r_{ij} denotes the reliability of agent A_j according to agent A_i. The credibility and reliability bases are also restricted by normalisation constraints

$$\sum_{i=1}^{k} c_{ij} = 1 \text{ and } \sum_{j=1}^{k} r_{ij} = 1,$$

for each agent A_i involved. A strategy base with k agents and m strategies is expressed in strategy variables $\{u_{11},\ldots,u_{1k},\ldots,u_{m1},\ldots,u_{mk}\}$ stating the utility of the strategies according to the different agents. The term u_{ij} denotes the utility of strategy S_i in the opinion of agent A_j. The three bases together constitute an *information frame*. It is assumed that the variables' respective ranges are real numbers in the interval [0,1]. Below, we will refer to an information frame as a structure $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$.

We will now provide a more detailed description of how problems, arising when eliminating the central position of a coordinating agent, are handled. We begin by discussing the representation and properties of assessments of agent credibility, agent reliability, and utilities of strategies according to the different agents.

Evaluation

In eliminating the DMA we obtain three different bases to consider in the evaluation process. The global utility of a strategy is a function of the reports of the individual agents, the agents' credibility when giving such reports according to the other agents, and all agents' opinions on each others reliability. The intuition for the formula expressing the global expected utility is that all assessments concerning the utilities of the strategy S_i are taken into account, each of which is balanced with the different credibility assessments. The credibility assessments are in turn balanced with the reliability assessments. For instance, $u_{i1}c_{11}$, i.e. the utility of strategy S_i in the opinion of agent A_1 weighted by the credibility of agent A_1 according to itself, is weighted by the reliability of agent A_1 according to all agents in the system. The resulting expression of this is

$u_{i1}c_{11}r_{11} + u_{i1}c_{11}r_{21} + \ldots + u_{i1}c_{11}r_{k1}$.

Similar expressions are received for all terms and the resulting expression when all assessments are taken into account is then (except for a normalisation factor)

$u_{i1}c_{11}r_{11} + u_{i1}c_{11}r_{21} + \ldots + u_{i2}c_{12}r_{11} + \ldots + u_{ik}c_{kk}r_{kk}$.

In the remaining parts of this section, the information frame $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$ is supposed to represent a system with k agents and m strategies.

Definition 1

Given an information frame $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$, the *global expected utility* $G(S_i)$ of a strategy S_i is:

$$G(S_i) = \frac{1}{k} \cdot \sum_{j=1}^{k} \sum_{m=1}^{k} \sum_{n=1}^{k} u_{ij} \cdot c_{mj} \cdot r_{nm}$$

where u_{ij}, c_{ij}, and r_{ij} are variables in $\mathcal{S}$, $\mathcal{K}$, and $\mathcal{R}$, respectively. ■

This expression should be evaluated with respect to the constraints in the information frame. However, since the agents are not necessarily consistent in their estimates, the first step of the evaluation procedure is to determine the solution sets to $\mathcal{K}$, $\mathcal{R}$, and $\mathcal{S}$.

Definition 2

A list of numbers $[n_1,\ldots,n_k]$ is a *solution vector* to a base X containing variables x_i, where $i = 1,\ldots,k$, if n_i can be consistently substituted for x_i in X. The set of solution vectors to a base constitutes the *solution set*. This set is a convex polytope, i.e. an intersection of a finite number of closed half spaces. ■

The solution sets to the bases can be determined by ordinary linear programming (LP) methods. In order to define an evaluation principle, a notation for instantiations of the global expected utility of an strategy is introduced below.

Definition 3

Given an information frame $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$. Let a, d, and e be three vectors of real numbers $(a_{i1},\ldots,a_{ik})$, $(d_{11},\ldots,d_{kk})$, and $(e_{11},\ldots,e_{kk})$.

$$^{ade}G(S_i) = \frac{1}{k} \cdot \sum_{j=1}^{k} \sum_{m=1}^{k} \sum_{n=1}^{k} a_{ij} \cdot d_{mj} \cdot e_{nm}$$

where a_{ij}, d_{ij}, and e_{ij} are numbers substituted for u_{ij}, c_{ij}, and r_{ij} in $G(S_i)$). ■

The evaluation an alternative can be based on the concept of admissibility. This intuitively means that a strategy is admissible iff no other strategy is better w.r.t. $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$. Thus, a strategy S_i is admissible if it has specific properties w.r.t. the information frame under consideration. The *t* in the expressions below is a real number in the interval [0,1] introduced for studying how large the differences between the expected utilities are.

Definition 4

Given an information frame $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$ and a real number t in the interval [0,1].

S_i is *at least as t-good* as S_j iff $^{a_i d_i e_i}G(S_i) - {}^{a_j d_j e_j}G(S_j) - t \geq 0$, for all a_i, d_i, e_i, a_j, d_j, e_j, where a_i and a_j are solution vectors to $\mathcal{S}$, d_i and d_j are solution vectors to $\mathcal{K}$, and e_i and e_j are solution vectors to $\mathcal{R}$.

S_i is *t-better* than S_j iff S_i is at least as t-good as S_j and $^{a_i d_i e_i}G(S_i) - {}^{a_j d_j e_j}G(S_j) - t > 0$, for some a_i, d_i, e_i, a_j, d_j, e_j, where a_i and a_j are solution vectors to $\mathcal{S}$, d_i and d_j are solution vectors to $\mathcal{K}$, and e_i and e_j are solution vectors to $\mathcal{R}$.

S_i is *t-admissible* iff no other S_j is t-better. ■

These definitions (for t = 0) conform to statistical decision theory (Lehmann 1959), and can be used below when comparing different alternatives. In the general case, determining admissibility is a fairly demanding task from a computational viewpoint. Consequently, general methods are not suitable for interactive use of the evaluation procedure. Instead, we use procedures for reducing problems of this particular kind (i.e. when the bases are separated) to linear systems, solvable with LP methods (Danielson & Ekenberg 1996).

Using the above definition, the strategies are evaluated with respect to $\langle \mathcal{S}, \mathcal{K}, \mathcal{R} \rangle$ and one of three situations may occur. For a real number t:

(i) No strategy is t-admissible

(ii) One strategy is t-admissible

(iii) More than one strategy is t-admissible

Case (i) has no solution. For case (ii) the task is done, since the only remaining strategy is superior. Case (iii) remains, which implies a set of seemingly reasonable strategies containing more than one element. When this case occurs, admissibility seems to be too weak to form a decision rule by itself. We have earlier suggested different principles for further discriminating rules for centralised agent systems (Ekenberg & Danielson 1994), (Ekenberg, Danielson, & Boman 1996a), but they can easily be adapted to the extended information frames. A class of techniques is also discussed in (Danielson & Ekenberg 1996). Using such procedures, incoming reports from agents with varying degrees of reliability can be evaluated in order to reach an informed decision based on possibly deviating report contents.

Decision Trees

The model used above requires much responsibility from the agents in the system since all agents are allowed to participate in the decision process. Consequently, for the agents to produce qualified deliberations, they should be able to evaluate the strategies under consideration in a qualified way. In the section above, the problem was modelled without taking into account how decision making agents arrived at their preferences and there were no requirements on the methods they used in this process. By extending the concept of strategy, and using techniques similar to those proposed above, more general decision models can be evaluated. Consider a traditional decision model in Figure 2.

	s_1	s_2	...	s_n
S_1	c_{11}	c_{12}	...	c_{1n}
S_2	c_{21}	c_{22}	...	c_{2n}
...	...	...	...	...
S_m	c_{m1}	c_{m2}	...	c_{mn}

Figure 2 **A state-consequence matrix**

The possible states $(s_1,\ldots,s_n)$ in the model describe a set of mutually exclusive and exhaustive descriptions of the world, not leaving any relevant state out. These states determine the consequences (such as c_{ij}) of the different strategies (alter-

natives) $S_1,\ldots,S_m$. The true state is the state that does eventually obtain the description of the actual world. Thus, if we adopt the strategy S_2 and if s_3 becomes the true state, consequence c_{23} will occur. The preferences among the consequences are supposed to be expressed by some kind of value function, for instance a utility function. If such a function exists, the value of consequence c_{ij} can be mapped onto a value v_{ij}, and the matrix can be evaluated with respect to different evaluation rules that have been proposed during the years (cf. (Ekenberg, Danielson, & Boman 1996b)). This model is formulated in terms of a single decision node, i.e., when a decision making agent have only one decision to make. This is unnecessarily restrictive, since an agent often confronts a decision situation involving paths with several decision nodes, (cf., e.g., (Raiffa 1968)). Such a situation can be modelled in a decision tree such as in Figure 3. The directed edges labelled by S in the figure denote alternatives, and the c's different consequences. The squares are decision nodes, i.e., where a decision has to be made by an agent. The circles denote chance nodes, from which edges lead to leaves or new decision nodes. Finally, the leaves correspond to ultimate consequences. A directed edge labelled by p denotes the probability of the node where the edge terminate, given that the alternative (leading to the chance node where p begins) occur.

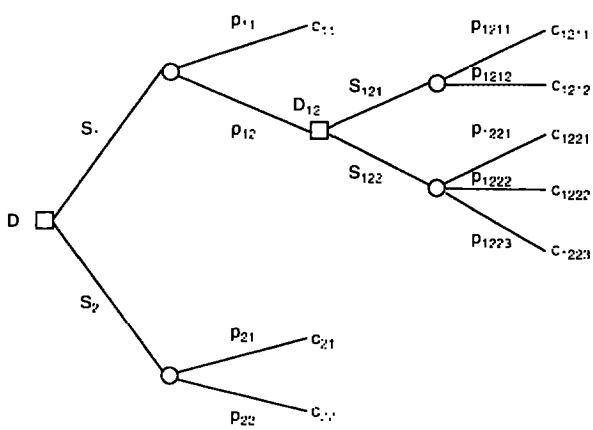

Figure 3 **A Decision Tree**

This model could be extended in a way similar to before by allowing for imprecise assessments such as "the probability of consequence c_{ij} obtaining is in the interval [40%,60%[" or comparative assessments such as "consequence c_{ij} is preferred to consequence c_{kl}". These assessments can be represented in systems of linear inequalities in probability variables p_{ij} and utility variables u_{ij}. Similar to before, the first assessment can be represented by the two inequalities $p_{ij} \geq 40\%$ and $p_{ij} \leq 60\%$, and the second one by $u_{ij} \geq u_{kl}$. To simplify the presentation in the sequel, it is assumed that, to each chance node, there is at most one directed edge leading to a decision node. The general case is very similar.

Definition 5

Given a decision tree. A set $\{c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}\}$ is an *alternative associated with a chance node* C_i, if the elements of the set are exhaustive and pairwise disjoint with respect to C_i. (This notation will be used even if an alternative does not contain an element $D_{i(s_i+1)}$). ■

Informally, this means that exactly one of $c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}$ will occur given that the alternative, represented by the directed edge to C_i is chosen.

Definition 6

Given a decision tree, a sequence of edges $[S_1,\ldots,S_r]$ is a *strategy*, if for all elements in the set, S_{i-1} is a directed edge from a decision node to a chance node C_{i-1}, and there is a directed edge from C_{i-1} to a decision node from which S_i is a directed edge. ■

Definition 7

Given an agent A, a decision tree associated with A, and a strategy $[S_1,\ldots,S_r]$, where each S_i is an alternative, i.e. $\{c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}\}$ associated with a chance node C_i. The *expected utility of* $[S_1,\ldots,S_r]$ *with respect to agent A*, $E^A(S_1,\ldots,S_r)$, is defined by the following:

$$E^A(S_i) = \sum_{j=1}^{s_i} p_{ij} \cdot u_{ij}$$

when S_i is an alternative $\{c_{i1},\ldots,c_{is_i}\}$, and

$$E^A(S_i,\ldots,S_r) = \sum_{j=1}^{s} p_{ij} \cdot u_{ij} + \left\{ p_{i(s_i+1)} \cdot E^A(S_{i+1},\ldots,S_r) \right\}$$

when S_i is an alternative $\{c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}\}$.

u_{ij} denotes the utility of the consequence c_{ij}, and p_{ij} denotes the probability of the consequence c_{ij} (or D_{ij}), given C_i. ■

Given a decision tree T, a decision node D in T can be considered a set $\{S_1,\ldots,S_q\}$ of alternatives, i.e. all directed edges from D. Two bases may be associated to D, one containing the probability variables of the edges from each S_i, and one containing the utility variables corresponding to possible leaves emanating from each S_i. Using such a structure, vague and numerically imprecise assessments can be represented and evaluated in a similar way as described in the previous section. The inequalities containing utility variables are included in the utility base *V(D)*, and inequalities containing probability variables are included in the probability base *P(D)*. These bases are composed to a *local decision frame corresponding to D and agent A* $\langle P^A(D), V^A(D)\rangle$. This framework for evaluating the expected utility of a strat-

egy can be combined with the framework described above and the total decision situation can be evaluated with respect to all credibilities, reliabilities, strategies, probabilities and utilities involved in the decision situation under consideration. Each agent in a multi-agent system may assert probability and utility assessments with respect to the tree. In this sense the probability and utility bases are local to each agent. What remains is to substitute the strategy utilities in section two with the expected utility of a strategy as defined in this section.

Definition 8

Given a set of agents $\{A_1,\ldots,A_k\}$, k decision trees – each associated with exactly one agent, and a strategy $[S_1,\ldots,S_r]$, the *global expected utility of* $[S_1,\ldots,S_r]$, $G(S_1,\ldots,S_r)$, is defined as:

$$G(S_1,\ldots,S_r) = \frac{1}{k}\cdot\sum_{j=1}^{k}\sum_{m=1}^{k}\sum_{n=1}^{k} E^{A_j}(S_1,\ldots,S_r)\cdot c_{mj}\cdot r_{nm}$$

where c_{ij} and r_{ij} are as in Definition 1. ■

Note that the definition does not presume that the decision trees for the different agents are identical. For some domains the tree could be the same for all agents and only the probability and utility assessments may differ. In other domains the agents may have constructed different decision trees involving the strategies under consideration. Similar to the previous section, the strategies are evaluated with respect to the information in the credibility and reliability bases. The difference here is that the strategy base is replaced by a set of probability and utility bases.

Consider the prerequisites in the definition above. Each S_i in the strategy $[S_1,\ldots,S_r]$ is an alternative on the form $\{c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}\}$, for each agent A. Each S_i is associated with a chance node C_i. Assume that the directed edge leading to C_i emanates from the decision node D_i, to which a local decision frame $\langle P^A(D_i),V^A(D_i)\rangle$ corresponds. Such a frame contains constraints representing the probability and utility assessments of agent A. Consequently, $E^A(S_1,\ldots,S_r)$ is associated with the set $\{\langle P^A(D_j),V^A(D_j)\rangle\}_j$, j = 1,...,r, in the same way as the strategy variables used in the previous section are associated with the strategy base.

Definition 9

Given an agent A, a decision tree T, and a strategy $[S_1,\ldots,S_r]$ in T, where each S_i is an alternative $\{c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}\}$ associated with a chance node C_i. Let $a_1,\ldots,a_r$ be a vector of real numbers $\{(a_{i1},\ldots,a_{i(s_i+1)})\}_{i=1,\ldots,r}$ and let $b_1,\ldots,b_r$ be a vector of real numbers $\{(b_{i1},\ldots,b_{i(s_i)})\}_{i=1,\ldots,r}$.

Now, ${}^{a_i b_i}E^A(S_i)$ is defined by the following:

$${}^{a_i b_i}E^A(S_i) = \sum_{j=1}^{s_i} a_{ij}\cdot b_{ij}$$

when S_i is an alternative $\{c_{i1},\ldots,c_{is_i}\}$, and

$${}^{a_1 b_1\ldots a_r b_r}E^A(S_i,\ldots,S_r) = \left(\sum_{j=1}^{s} a_{ij}\cdot b_{ij} + \left(a_{i(s_i+1)}\cdot {}^{a_{i+1}b_{i+1}\ldots a_r b_r}E^A(S_{i+1},\ldots,S_r)\right)\right)$$

when S_i is an alternative $\{c_{i1},\ldots,c_{is_i},D_{i(s_i+1)}\}$. ■

This may now be combined with the notation for instantiations of the global expected utility of an strategy in the previous section into the following:

Definition 10

Given a set of agents $\{A_1,\ldots,A_k\}$, k decision trees – each associated with exactly one agent, and a strategy $[S_1,\ldots,S_r]$.

Let $a_1,\ldots,a_k$, $b_1,\ldots,b_k$ be vectors of vectors $\{({}^ja_1,\ldots,{}^ja_r),({}^jb_1,\ldots,{}^jb_r)\}_{j=1,\ldots,k}$. The latter are vectors of real numbers $\{({}^ja_{i1},\ldots,{}^ja_{i(s_i+1)})\}_{i=1,\ldots,r}$, $\{({}^jb_{i1},\ldots,{}^jb_{i(s_i)})\}_{i=1,\ldots,r}$. Also let d and e be vectors of real numbers $(d_{11},d_{12},\ldots,d_{kk})$ and $(e_{11},e_{12},\ldots,e_{kk})$. Now,

$${}^{a_1 b_1\ldots a_k b_k d e}G(S_1,\ldots,S_r) = \left(\frac{1}{k}\cdot\sum_{j=1}^{k}\sum_{m=1}^{k}\sum_{n=1}^{k} {}^{{}^1a_1{}^1b_1\ldots{}^ra_r{}^rb_r}E^{A_j}(S_1,\ldots,S_r)\cdot d_{mj}\cdot e_{nm}\right)$$

where ${}^ia_{ij}$, ${}^ib_{ij}$, d_{ij}, and e_{ij} are numbers substituted for ${}^ip_{ij}$, ${}^iu_{ij}$, c_{ij}, and r_{ij} in $G(S_1,\ldots,S_r)$. ■

The different strategies can then be evaluated, for instance with respect to admissibility as before.

Definition 11

A *general decision frame* $\mathcal{F}$ is a structure $\langle\mathcal{T},\mathcal{S},\mathcal{L},\mathcal{K},\mathcal{R}\rangle$. $\mathcal{T}$ is a set of T_j's – decision trees associated with the agents A_j, j=1,...,k. $\mathcal{S}$ is the set of possible strategies modelled in the trees. $\mathcal{L}$ is a set of local decision frames $\langle P^A{}_j(D_i),V^A{}_j(D_i)\rangle$ corresponding to D_i and agent A_j, where D_i is a node in the tree T_j. $\mathcal{K}$ and $\mathcal{R}$ are credibility and reliability bases as in the previous section. ■

Definition 12

Given a general decision frame $\mathcal{F}$ and a real number t in the interval [0,1].

The strategy $[S_{i_1},\ldots,S_{i_r}]$ is at least as *t-good* as the strategy $[S_{j_1},\ldots,S_{j_q}]$ iff

$${}^{a_1 b_1\ldots a_k b_k d_1 e}G(S_{i_1},\ldots,S_{i_r}) - {}^{f_1 g_1\ldots f_k g_k d e}G(S_{j_1},\ldots,S_{j_q}) - t \geq 0$$

for all d_i, e_i, d_j, e_j, where d_i and d_j are solution vectors to $\mathcal{K}$, and e_i and e_j are solution vectors to $\mathcal{R}$. Furthermore, each ${}^{j}a_i$ in a_i and each ${}^{j}f_i$ in f_i are solution vectors to $P^{A}{}_{i}(D_i)$, and each ${}^{j}b_i$ in b_i, and each ${}^{j}g_i$ in g_i are solution vectors to $V^{A}{}_{i}(D_i)$.

The strategy $[S_{i_1},\ldots,S_{i_r}]$ is *t-better* than the strategy $[S_{j_1},\ldots,S_{j_q}]$ iff $[S_{i_1},\ldots,S_{i_r}]$ is at least as t-good as $[S_{j_1},\ldots,S_{j_q}]$ and

$$ {}^{a_1 b_1 \ldots a_k b_k d_i e_i}G(S_{i_1},\ldots,S_{i_r}) \quad {}^{f_1 g_1 \ldots f_k g_k d_j e_j}G(S_{j_1},\ldots,S_{j_q}) - t > 0 $$

for some d_i, e_i, d_j, and e_j, where d_i and d_j are solution vectors to $\mathcal{K}$, and e_i and e_j are solution vectors to $\mathcal{R}$, and for every a_i, b_i, f_i, g_i, $i=1,\ldots,k$, ${}^{j}a_i$ in a_i and ${}^{j}f_i$ in f_i are solution vectors to $P^{A}{}_{i}(D_i)$, and ${}^{j}b_i$ in b_i, and ${}^{j}g_i$ in g_i are solution vectors to $V^{A}{}_{i}(D_i)$.

The strategy $[S_{i_1},\ldots,S_{r_i}]$ is *t-admissible* iff no other strategy $[S_{i_j},\ldots,S_{r_j}]$ in $\mathcal{I}$ is t-better. ■

If the set of admissible strategies is too large, contraction methods similar to those suggested in the previous section can be used for investigating the stability of the result.

Concluding Remarks

We have shown how a set of agents can analyse and evaluate vague and numerically imprecise reports made by the individual agents when determining which strategies are reasonable to choose among. The approach considers a decision problem with respect to the contents and the credibilities of the reports as well as the reliability of the agents making the reports. These three aspects are modelled into an information frame consisting of three systems of linear expressions stating inequalities and interval assessments. The strategies are evaluated relative to generalisations of the principle of maximising the expected utility. We also demonstrated how decision trees can be integrated into the framework and suggested an evaluation rule taking into account all strategies, credibilities, reliabilities, probabilities and utilities involved in the framework. The rule is based on the classical concept of admissibility. The need for a central co-ordinating agent is eliminated and the framework is in this respect suitable for use in co-operating autonomous agent architectures. However, the approach has two shortcomings that should be emphasized.

(i) As was mentioned above, the model is restricted in the sense that there are no good justification why the agents should not have opinions on the reliability of the other agents in a more general sense. To model such a perspective would theoretically imply an infinite regress. Consequently, we have to assume that the assessments converge after a reasonable number of levels.

(ii) Another problem is that despite the development of fast algorithms for solving the kinds of problems that are described in this paper (Danielson 1996), solving very large problems may still be quite time-consuming. A natural line of research is to follow aspects such as discussed in (Good 1952), (Horwitz & Klein 1995), (Horwitz, Cooper, & Heckerman 1989), where the cost of the deliberation also is taken into account, and to investigate under what restrictions on the different bases the expected utility of a strategy converge. This will be treated in a forthcoming paper.

Acknowledgments

This project was supported by the projects *Risk Evaluations in Vague Domains* and *Coordination of Information in Cooperative Information Systems* (CISIS). Both of these are financed by NUTEK (Swedish National Board for Industrial and Technical Development).

References

Boman, M., and Ekenberg, L. 1995. Decision Making Agents with Relatively Unbounded Rationality. In Proceedings of DIMAS'95. I/28–I/35.

Clarke, E. H. 1971. Multipart Pricing of Public Goods. *Public Choice* 11:17–33.

Danielson, M. 1996. A Computational Framework for Decision Analysis, Royal Inst. of Technology, Ph.D. diss., DSV, Royal Inst. of Technology. Forthcoming.

Danielson, M., and Ekenberg, L. 1996. A Framework for Analysing Decisions under Risk. Submitted.

Ekenberg, L., Boman, M., and Danielson, M. 1995. A Tool for Coordinating Autonomous Agents with Conflicting Goals. In Proceedings of First International Conference on Multi-Agent Systems, 89–93.: AAAI/The MIT Press.

Ekenberg, L., and Danielson, M. 1994. A Support System for Real-Life Decisions in Numerically Imprecise Domains. In Proceedings of the International Conference on Operations Research '94, 500–505.: Springer-Verlag.

Ekenberg, L., and Danielson, M. 1995. Handling Imprecise Information in Risk Management. In *Information Security – the Next Decade*. Eloff, J. and Solms, S. von, eds.: Chapman & Hall.

Ekenberg, L., Danielson, M., and Boman, M. 1995. A Tool for Handling Uncertain Information in Multi-Agent Systems. In Proceedings of MAAMAW'94, 54–62.: Springer-Verlag.

Ekenberg, L., Danielson, M., and Boman, M. 1996a. From Local Assessments to Global Rationality. To appear in *International Journal of Intelligent and Cooperative Information Systems*.

Ekenberg, L., Danielson, M., and Boman, M. 1996b. Imposing Security Constraints on Agent-Based Decision Support. To appear in *Decision Support Systems International Journal*.

Ekenberg, L., Oberoi, S., and Orci, I. 1995. A Cost Model for Managing Information Security Hazards. *Computers & Security*, 14:707–717.

Ephrati, E., and Rosenschein, J. S. 1991. The Clarke Tax as a Consensus Mechanism Among Automated Agents. In Proceedings of Ninth National Conference on Artificial Intelligence, 173–178.: AAAI Press/The MIT Press.

Ephrati, E. and Rosenschein, J. S. 1993. Multi-Agent Planning as a Dynamic Search for Social Concensus. In Proceedings of 13th IJCAI, 423–429.: Morgan Kaufmann Publishers.

Fishburn, P. 1981. Subjective Expected Utility: A Review of Normative Theories. *Theory and Decision*, 13:139–199.

Gmytrasiewicz, P. J., and Durfee, E. H. 1993. Elements of a Utilitarian Theory of Knowledge and Action. In Proceedings of 13th IJCAI, 396–402.

Good, I. J, 1952. Rational Decisions, *J. R. Statist. Society B*, 14:107–114.

Horwitz, E., and Klein, A. 1995. Reasoning, Metareasoning, and Mathematical Truth: Studies of Theorem Proving under Limited Resources. In Proceedings of 11th Conference on Uncertainty in AI.

Horwitz, E. J., Cooper, G. F., and Heckerman, D. E. 1989. Reflection and Action Under Scarce Resources: Theoretical Principles and Empirical Study. In Proceedings of 11th IJCAI, 1121–1127.

Katz, M. J., and Rosenschein, J. S. 1993. Verifying Plans for Multiple Agents. *Journal of Experimental and Theoretical Artificial Intelligence*, 5:39–56.

Kraus, S. 1993. Agents Contracting Tasks in Non-Collaborative Environments. In Proceedings of 11th National Conference on Artificial Intelligent. 243–248.: AAAI/The MIT Press.

Lehmann, E. L. 1959. *Testing Statistical Hypothesis*.: John Wiley and Sons.

Raiffa, H. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*.: Random House.

Rosenschein, J. S. 1982. Synchronization of Multi-Agent Plans. In Proceedings of National Conference of Artificial Intelligence, 115–119.

Rosenschein, J. S. 1988. The Role of Knowledge in Logic-based Rational Interactions. In Proceedings of Seventh Phoenix Conference on Computers and Communications, 497–504.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*.: MIT Press.

Smith, R. G. 1980. The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver. *IEEE Transactions on Computers*, C-29(12):1104-1113.

Zlotkin, G., and Rosenschein, J. S. 1991. Incomplete Information and Deception in Multi-Agent Negotiation. In Proceedings of 12th IJCAI, 225–231.