Working Paper

Reasonable Goals

Love Ekenberg

WP-96-136
November 1996

IIASA International Institute for Applied Systems Analysis • A-2361 Laxenburg • Austria
Telephone: +43 2236 807 • Telefax: +43 2236 71313 • E-Mail: info@iiasa.ac.at

Reasonable Goals

Love Ekenberg

WP-96-136
November 1996

IIASA International Institute for Applied Systems Analysis • A-2361 Laxenburg • Austria
Telephone: +43 2236 807 • Telefax: +43 2236 71313 • E-Mail: info@iiasa.ac.at

Contents

Reasonable Goals

Love Ekenberg

ekenberg@iiasa.ac.at
IIASA
International Institute for
Applied Systems Analysis
A-2361 Laxenburg, Austria
Telephone: +43-2236 807 209
Telefax: +43-2236 73147

Abstract. Assume that a number of autonomous agents are going to act in such a way that their respective goal states constitute a global plan. A main question that arises in this situation is whether there is such a plan at all, i.e. whether a solvable conflict prevails. In some sense, this means that the set of common goals is non-empty. Furthermore, if the agents are allowed to act in accordance with the result of some decision process, a situation may occur where subsets of their possible goal sets are consistent, but in actual fact the individual agents may nevertheless always terminate in states that are in conflict. We present a formal framework for the analysis of conflicts in sets of autonomous agents restricted in the sense that they can be described in a (first-order) language and by a transaction mechanism. This is also enriched by processes for evaluating decision situations given imprecise background information. The agent specifications are analysed with respect to a concept of consistency that requires the formulae of one specification together with a set of correspondence assertions to not restrict the models of an other specification, i.e. the agent system does not essentially restrict the individual agents. The main emphasis is on the specifications being compatible with respect to reasonable probable states, i.e. states for which it is reasonable to assume that they eventually will be reached.

1 Motivation

A conflict situation in a multi agent system is often considered to be one in which the negotiation set is empty, i.e. the ultimate goals of the individual agents are incompatible (cf. [23]). Even if we for a while disregard a precise definition of a conflict, it seems plausible that in a complex multi agent system it may be that it is practically impossible to determine whether a conflict actually prevails. This seems also to be the case even if there are no restrictions on the interactions or communication processes between them. Moreover, even if the agent system can be proven consistent, i.e. it has a common non-empty subset of goals, it may be difficult to determine whether this subset is meaningful in the sense that it contains goal states that the agents will reach with a reasonable degree of probability.

In [8] this problem was investigated in a restricted sense. The approach taken was that substantial parts of multi agent designs can often be formulated as processes (in the sense of algorithms). Consequently, aspects as beliefs, intentions, free will, wants, or consciousness was not taken into account. Thus, agent systems was considered as transition systems that could be specified in some kind of process based language. Parts of such a structure can be translated to a set of conceptual schemata, i.e. first-order formulae integrated with transition models. It was argued that this enables various methods for determining conflicts in specifications and allows analysis of the integrated model. The technical background of such a framework was introduced in [14], where an approach using integration assertions was chosen to demonstrate how different aspects of integration of such models formulated in first-order logic can be analysed. This approach makes the integration process simpler and easier to understand than the traditional restructuring approach [1]. This is further extended upon in [15], where it is demonstrated how general dynamic aspects of specifications are handled.

However, the work is limited in the sense that it does not separate between meaningful states and states that in actual fact could be reached when certain kinds of restrictions are imposed. In the context described herein, this means that the suggested consistency checks are a bit too weak to constitute a framework for investigating when specifications can be meaningfully integrated with respect to reasonable probable states. The current paper demonstrates how this deficiency can be removed. The next section briefly explains a formalism for extended finite state automata that can be used to model some aspects of agent behaviour. The particular representation is not essential for the main point of this paper and another formalism for representing algorithms could be used as well. This is also emphasised in Section 4 where the principle of reasonable freedom of conflicts is introduced. Section 3 describes how a framework for decision making agents can be introduced in the formalism for extended finite automata and explains how reasonable states are introduced in the framework. Section 5 concludes the paper.

2 Process Specifications

Fig. 1 from [8] illustrates schematically a procedure for conflict detection. Assume that there are initially only two specifications to be analysed. Aiming at facilitating logical analysis, they are first transformed into schemata which can be analysed separately. Thereafter they are merged, allowing them to be enriched with further invariants. At this stage the resulting schema can be analysed for static and dynamic consistency.

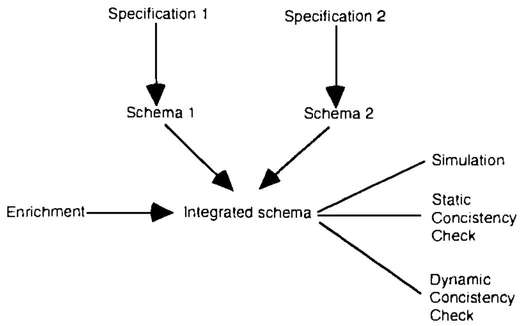

Fig. 1. Transformation, integration, enrichment, and analysis

A process can be described in a suitable formalism. Fig. 2 shows a specification of a process type expressed with SDL88's graphical notation.

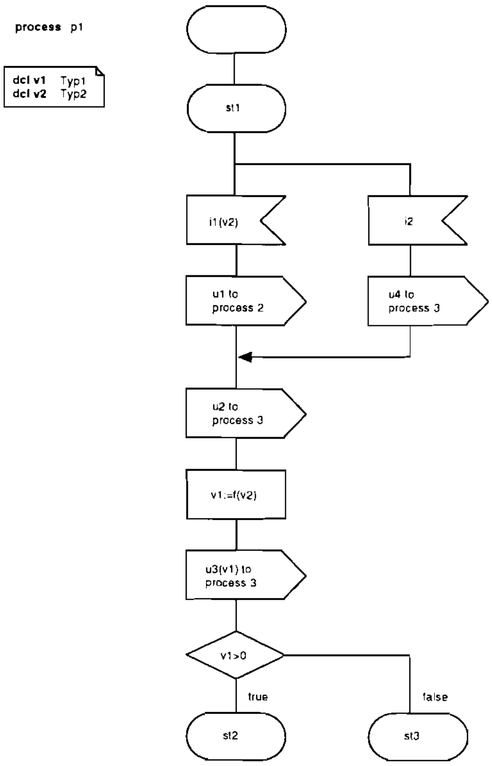

Fig. 2. A process specification in SDL

The basic symbols used in the graphical representation of SDL88 are as follows. An oval represents a state. A nicked box represents the consumption of a signal from the process' queue. A pointed box sends a signal. Plain boxes contain assignments and other operations on data contained in process variables. A diamond shape represents a decision node where the course of the transition depends on the outcome of evaluating some variable. The symbols are connected by a uni-directional arc.

A series of shapes and arcs builds a path that represents a procedural sequence. Paths follow a general downward direction. Every such path from one state symbol to the next represents a transition. The name of the process type *p1* is in the top left hand corner of the graph where the variables *v1* and *v2* are also declared and typed.

The first symbol in the figure is an empty oval representing the *start state* of the process. Every process has a start state. When an instance of a process is created, the instance automatically transits from the start state to the first state, carrying out whatever operations are specified along the way. In Fig. 2 there are no intermediary operations. Continuing with the example, the process instance waits in state *st1* until it receives an input signal of type *i1* or *i2*, upon which a transition is fired. If the input signal is of type *i1* then the variable *v2* is implicitly assigned the value of the parameter carried by the signal *i1*. The transition then proceeds to an output symbol which sends signal *u1* to process *2*. If the signal is of type *i2* then signal *u4* is sent to process *3*. The next four symbols are common to both paths. Signal *u2* is sent to process *3*, variable *v1* is assigned the value of *f(v2)* and a decision node is reached. If at this node *v1* has a value exceeding zero the transition terminates leaving the process instance in state *st2*, otherwise in *st3*. The process will remain in its new state until another transition is fired.

To not spend time on insignificant features of the particular formalism used to illustrate process specifications, we will not explain the features of SDL88 in detail. Instead we will rely on that the intuition for the semantics of a specification in SDL88 is sufficiently clear for the purpose of this paper. [4] provides the reader with the formal syntax and semantics of SDL88. In the next section we will focus on the decision nodes of the representation since they, in a certain sense, define the possible choice mechanisms in the processes.

3 Decision Making Agents

As described above, the actual paths that can be taken from a decision node are determined by a specific value that may have been acquired at an earlier stage of the process execution. At this node, the model can be extended by integrating more complicated structures into the communication model.

The framework used here is based on the idea that processes may receive information to update an information frame consisting of linear systems of inequalities. These systems can be treated in a variety of ways to determine reasonable courses of actions, but has not earlier been integrated in a model for agent processes. The basic ideas of decision making agents will not be recapitulated here. Instead, we only provide an example to illustrate the main ideas of possible treatments of information frames. The reader is referred to [9-13] for more careful treatments of a variety of aspects on this particular issue in centralised or decentralised architectures.

Assume that the agents may be 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. In the terminology suggested here, strategies correspond to uni-directional arcs from the decision node to the next symbol in the specification.

(ii) They may have opinions on the reliabilities of the other agents as they in turn assess the abilities of other agents. The arcs are also evaluated by the different agents with respect to their utilities according to the individual agents.

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

- The utility of arc 1 is between 0.10 and 0.30
- The utility of arc 2 is between 0.25 and 0.60
- The utility of arc 3 is between 0.45 and 0.60
- The utility of arc 2 is at least 0.15 better than that of arc 1

The other agents can state similar assessments about the utilities of the arcs. Moreover, the agents may estimate the credibility of the other agents as numbers in the interval [0,1], where the number 0 denotes the lowest credibility and 1 the highest. Thus, the assessments according to an agent could be:

- The credibility of agent 2 is at least equal to that of agent 1
- The credibility of agent 2 is between 0.35 and 0.75

Furthermore, the agents may estimate the reliabilities of the other agents as numbers in the interval [0,1]. One 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 analyses often result in comparisons. Since the sources may be different, the assessments are not necessarily consistent with each other.

The reports provided by the agents are translated to linear expressions. For instance, the utility assessments above are translated to the following:

$u_1 \in [0.10, 0.30]$

$u_2 \in [0.25, 0.60]$

$u_3 \in [0.45, 0.60]$

$u_2 \geq u_1 + 0.15$

Similarly, the credibilities as well as the reliabilities are also represented as numbers in the interval [0,1].

The sets of assessments are transformed into linear systems of equations. The credibility, reliability, and utility assessments, restricted by scaling and normalisation constraints, constitute an information frame. Consequently, the information frame consists of three different systems of linear inequalities, each expressed in a certain type of variables.

A decision problem may then be analysed with respect to various decision rules. For instance, the interesting choices may be the ones that are pareto optimal in certain respects, or the ones that do not violate specific thresholds. Often, there is also a need for further discriminating criteria. For instance, the decision may also be investigated with respect to stability constraints. That a decision is stable in this respect can mean that it is not affected when values close to the borders of the various intervals are ignored. Different kinds of algorithms for solving these kinds of problems have already been implemented in a tool for human decision makers [5].[1] Since they are independent of the particular agent architecture used, they can easily be adapted to automated evaluation in decision nodes in the sense described above.

In a specification allowing for the evaluation of information frames, a decision node may be of two kinds. It can be a regular node in the conventional meaning, but also a node involving a more intricate structure. An example of the latter is shown in Fig. 3, where the selection criteria is a bit more complex. If the expected value of arc S_1 ($E(S_1)$) is greater than the expected value of arc S_2 ($E(S_2)$), and this is the case when the information frame ($I(C,R,U)$) is contracted by 50%, the left arc is chosen. Otherwise, the right arc is chosen.

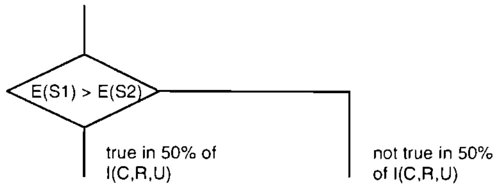

Fig. 3. Extending the decision abilities of an agent system

Thus, the agents are modelled with an underlying assumption of benevolence in the sense that they always obey the decisions taken at the decision nodes. An aspect which should be noted in the following is that the model is independent of whether the possible strategies given in a decision node have similarities with strategies in other nodes in other specifications. The only proviso in this respect is that in a multi agent system the agents have the *possibility* to express utility estimates concerning all possible strategies in a set of specification. The general framework is very liberal in this respect and the principles that should govern a particular implementation are naturally domain dependent. Needless to say, an information frame may also be changed, for instance, when the agent system is updated. However, we will assume that this is implicit in the information frames discussed below.

As we have seen, the paths worth consideration in a process description as described above are dependent of the particular information frames that are possible

[1] [6] also provides some comparisons with other models for representing and evaluating decision situations in imprecise domains.

in the different decision nodes. Thus, it is possible to characterise the reasonable decision space modulo a set of decision rules. The set of states that are accessible with respect to this space may then be considerably reduced. Only states that satisfy certain requirements defined by the information frames and the decision rules should be possible candidates for actual goals of the agents represented by the specifications.

For simplicity, in the definitions below we follow [12] by ignoring reliabilities and consequently suppose a centralised decision structure (cf. [17-21]). However, when introducing definitions similar to those in [9], a corresponding decentralised scenario is easily modelled.

Definition Let MA be a set of process specifications $\{P_1,\ldots,P_m\}$, and let N_{ij} be a decision node in the process P_j. An *information frame in node* N_{ij} is a structure $I_{ij}(C,U)$, where C and U are finite lists of linear constraints in the credibility and utility variables. These represent the credibility and utility estimates of the agents in MA.

As was mentioned above, the decision taken in a particular node depends also on the decision rule used at that node. One such candidate is a variant of the criterion of pareto optimality. First we define a weighted mean value of the values of the different arcs from a decision node. The weights in this sum are the various credibilities asserted by the agents (or in the simplified case a decision making agent) in the system. The intuition for the notions below is that the term c_i denotes the credibility of agent i, and that the term u_{ij} denotes the utility of arc i in the opinion of agent j.

Definition Given an information frame $I_{kl}(C,U)$, $E(p_i)$ denotes the expression $c_1u_{i1} + \ldots + c_{h_i}u_{ih_i}$, where c_i and u_{ij} are variables in C and U respectively.
Let a_i and b_i be two vectors of real numbers $(a_1,\ldots,a_{h_i})$ and $(b_{i1},\ldots,b_{ih_i})$.
${}^{a_i b_i}E(p_i)$ denotes the expression $a_1b_{i1} + \ldots + a_{h_i}b_{ih_i}$ (where a_i and b_{ij} are numbers substituted for c_i and u_{ij} in $E(p_i)$).

Next we define the concept of an admissible arc. An admissible arc is, in some sense, pareto optimal.

Definition Given an information frame $I_{kl}(C,U)$, an arc p_i is *at least as good* as an arc p_j iff ${}^{a_i b_i}E(p_i) - {}^{a_i b_j}E(p_j)$ _ 0, for all a_i, b_i, b_j, where $\{c_1 = a_1\}$ & ... & $\{c_{h_i} = a_{h_i}\}$ is consistent with the constraints in C and
$\{u_{i1} = b_{i1}\}$ &... & $\{u_{ih_i} = b_{ih_i}\}$ & $\{u_{j1} = b_{j1}\}$ &... & $\{u_{jh_i} = b_{jh_i}\}$ is consistent with the constraints in U.

p_i is *better* than p_j iff p_i is at least as good as p_j and ${}^{a_i b_i}E(p_i) - {}^{a_i b_j}E(p_j) > 0$, for some a_i, b_i, f_j, such that $\{c_1 = a_1\}$ &... & $\{c_{h_i} = a_{h_i}\}$ is consistent with C and
$\{u_{i1} = b_{i1}\}$ &... & $\{u_{ih_i} = b_{ih_i}\}$ & $\{u_{j1} = b_{j1}\}$ &... & $\{u_{jh_i} = b_{jh_i}\}$ is consistent with the constraints in V.

p_i is *admissible* iff no other p_j is better than p_i.

Observe that these concepts are defined locally, i.e. with respect to a given decision node in a given process specification. The next definition is a bit more informally stated. It defines the meaning of a path between the initial state of a specification and a particular state that may be reached from the initial state according to the specification. Two ordered symbols in a specification are consecutive if the second symbol immediately (along an arc) follows the first.

Definition Let P_i be a process specification and let s_j be a state in P_i. $S(s_j,P_i)$ is a set of ordered lists of consecutive symbols in P_i, from the initial state symbol to the state symbol denoting s_j.

The following definition characterises the set of reasonable states $CP(P_i)$ in a specification P_i. A state is a member in such a set if there is a path leading to the state, where all decision nodes in the path are followed by admissible arcs, i.e., it should be possible to choose them with respect to the information frames and the decision rules in the nodes.

Definition Let MA be a set of process specifications. Given a set of information frames $\{I_{ij}(C,U)\}$ in MA. A state s_k in the process P_j belongs to the set $CP(P_j)$ iff there is a list L in $S(s_k,P_j)$, such that for each decision node N_{ij} in L, the following symbol in L is connected to N_{ij} by an arc that is admissible with respect to $I_{ij}(C,U)$.

Again, it should be emphasised that, even if the presentation above presuppose a particular decision rule, there are no particular restrictions on the decision rules used by a decision node, and a variety of criteria and norms could be implemented. A framework for decision rules based on generalisations of the maximin and maximax criteria are discussed in [13] and could easily be adapted to the present framework. Different rules may also be grouped into hierarchies in the various nodes. The same applies to stability conditions as demonstrated in Fig. 3, where the decision is dependent on how stable the solutions are with respect to variations in the underlying information frame. Various kinds of stability conditions are discussed in [3, 6].

4 Conflict Detection

As we will see below, the conflict detection methods rely on theorem proving techniques. To simplify the use of these techniques the specification defined by state machines should be translated to a more suitable formalism. One way of representing significant features of a specification is by using conceptual schemata, and a treatment of the details of such a transformation is demonstrated in [7]. In this transformation, process specifications are transformed to formulae in first-order logic together with a transaction mechanism. Such a structure is referred to as a schema in the following sense.[2]

[2]The perceptive reader will note that the semantics of a schema is very loosely described and is referred to [15] for a more rigorous treatment of this.

Definition A *schema* S is a structure ⟨R,ER⟩ consisting of a *static part* R and a *dynamic part* ER. R is a finite set of closed formulae in a language L(S). ER is a set of *event rules*. Event rules describe the possible transitions between different states of a schema. The latter are called events for the schema. By a *diagram* for a set R of formulae in a first order language L, we mean a Herbrand model of S, extended by the negation of the ground atoms in L that are not in the Herbrand model. *A static integration assertion* expressing the schema S_2 in the schema S_1 is a closed first order formula: $\forall \mathbf{x}\ (p(\mathbf{x}) \leftrightarrow F(\mathbf{x}))$, where p is a predicate symbol in $L(S_2)$ and $F(\mathbf{x})$ is a formula in $L(S_1)$.

A description of a schema is a structure consisting of all diagrams for a schema, together with all possible transitions that are possible with respect to the basic events for the schema.

Definition The *description of a schema S* is a digraph ⟨R,E⟩, where D is the set of diagrams for S, and E is the set of basic events (i.e. arcs in the digraph) for S.

Fig. 4 illustrates a description of a schema. The arrows in the figure represent basic events and the dots represent the diagrams for the schema. If, for instance, the schema is in the state in the leftmost corner and an adequate event rule is initialised, the schema enters the leftmost upper state.

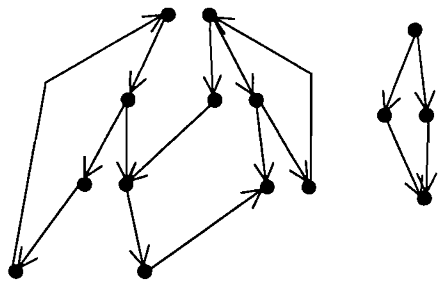

Fig. 4.

With respect to a set of conceptual schemata and integration assertions, the concept of conflictfreeness can be defined. Intuitively, two schemata are in conflict with respect to a set of static integration assertions if one of them together with the integration assertions restrict the set of diagrams for the other one. This means, for instance, that if the set of goals for an agent A_1 and the set of goals for an agent A_2 are (partly) incompatible, they are in conflict since the ultimate goals of both agents cannot simultaneously be fulfilled, i.e. they restrict each other in this sense. We will say that S_2 and S_1 are conflictfree with respect to IA iff for each diagram σ in D_1, there exists a diagram τ in D_2, such that $\sigma \cup \tau$ is a diagram for IA. However, to emphasise the independence of a particular underlying language, [8] formulates the concept of conflictfreeness somewhat more generally.

Definition Let IA be a set of static integration assertions expressing S_2 in S_1, and let the schemata be expressed as connected digraphs $\langle D_i,E_i\rangle$, where $E_i \subseteq D_i \times D_i$. Define a relation $R \subseteq D_1 \times D_2$. (There is no need to assume that the D_is are diagrams in the sense described above.) The definition of conflictfreeness between S_1 and S_2 then becomes an instance of R in the following sense: *S_2 and S_1 are conflictfree with respect to IA* iff $\forall\sigma\in D_1\ \exists\tau\in D_2\ \sigma R\tau$, where $\sigma R\tau$ iff $\sigma \cup \tau \models$ IA. Otherwise *S_2 and S_1 are in conflict with respect to IA.*

As [8] emphasises, the set of interesting goals could be decreased. In some cases only states including ultimate goals need consideration. If two schemata S_1 and S_2, representing the agents A_1 and A_2, are in conflict, and the conflicting states contain ultimate goals of the agents, then there may be a real conflict with respect to the goals. However, this set could now be further decreased by requiring that the set of goals should also be states that are on paths containing only admissible arcs. Therefore, we require that the goal set under consideration should be reasonable in the following sense.

Definition The *reasonable goal set* G_i in a schema S_i with respect to is the least set that fulfils the following clauses:
(i) if $\sigma \in G_i$, then σ is an ultimate goal in S_i, corresponding to an ultimate goal in the process modelled by S_i.
(ii) if $\sigma \in G_i$, then $\sigma \in CP(P_i)$, where P_i is the process modelled by S_i.

Thus, in case of a reasonable goal conflict detection, the relation R in the above definition could be modified in the following respect. (The definition does not take into account that several states may contain the same goal.)

Definition Let IA and $\langle D_i,E_i\rangle$ be as above. Further, let G_i be the reasonable goal set in a schema S_i. *S_2 and S_1 are conflictfree with respect to IA* iff $\forall\sigma\in G_1\ \exists\tau\in G_2\ \sigma R\tau$, where $\sigma R\tau$ iff $\sigma \cup \tau \models$ IA. Otherwise *S_2 and S_1 are in conflict with respect to IA.*

After each new update of the information frames, the schemata should be checked for conflictfreeness. It is possible that two specifications that were in conflict before such an update became conflictfree by the update, and vice versa. This is because a goal in a schema may cease to be a member in the reasonable goal set when new information is provided. It should also be observed that significant information concerning possible schema conflicts may be obtained even if the information frames are not entirely known. Note that an information frame containing only trivial constraints (i.e. scaling and normalisation constraints) is weaker than any possible information frame containing the trivial constraints. By generalising this observation, we can see that a characterisation of the possible information frames may in some cases be sufficient.

Definition Let MA be a set of process specifications. Given a set of possible information frames $\{I_{ij}(C,U)\}$ in decision node N_j in the process P_i in MA. The characterisation of $\{I_{ij}(C,U)\}$ is the union of the solution sets to the information frames in this set.

Consequently, the trivial information frame is a characterisation of all possible information frames. Needless to say, the more information that can be provided in terms of constraints in the characterisation, the more significant information can be derived about the reasonable goals in advance, but the main point here is that even when quite a weak characterisation prevails, information about reasonable goals can be acquired. Goals that are not reasonable with respect to a characterisation will never be reasonable with respect to an information frame in the characterisation (provided that the characterisation is correct, in the sense that it really corresponds to the possible information frames, of course).

A note on the complexity properties might also be in place here. Assume a first order language. With the (in most cases) reasonable assumption that there is a finite number of objects in the agent system as well as in its environment, the language restricted to include only a finite number of constants may be used. In that case we have a problem in second order propositional logic that is Π_2^P-complete. By restricting the formulae in IA as suggested in [15], we receive a problem that is NP-complete. Needless to say, even this is computationally demanding, but with an efficient theorem prover NP-complete problems can be solved (in most cases) within a reasonable time (cf. [22], [16]).

Note that the definition of conflictfreeness does only take into consideration the static aspects of the agent behaviour, i.e. how the static properties of one agent affect the static properties of another. The framework can be extended to the dynamic case in a similar way as described in [8]. The main proviso is that the new concept of reasonable goal set should replace the former meaning of a goal set, where no decision mechanism was introduced.

5 Concluding Remarks

The work described in this paper has two purposes. Firstly, it is motivated by the difficulties to determine whether complex agent systems in actual fact are in conflict. Secondly, the set of goals that are compatible when investigating the system from a meta-perspective may not be reasonable in the sense that the agents actually have a possibility to reach them. The latter may be the case when the agents are forced to abandon possible paths leading to goals because the utility is too low in a particular environment or when the probability to achieve them along a specific path is too low according to certain security constraints.

The framework described above can also be utilised in several ways in different phases in the execution of a multi agent system. First, assume that the possible sets of estimates can be characterised in a quite exact way at a certain point. Then the procedures suggested in the article can be utilised to exactly determine the reasonable goal sets at this point. If this cannot entirely be done, for instance when estimating future configurations in advance, the decision nodes may contain weaker constraints in terms of characterisations. Even these could be utilised for determining probable paths through the processes. If the frames in the latter situations turn out to be too tolerant, sensitivity analyses and contractions can be used to determine thresholds for possible estimates by evaluating the solution sets of the characterisations and determine, for instance, when an admissible arc ceases to be admissible. If these thresholds are violated, the goals along certain paths cease to be reasonable. This kind of constraints can also include stability properties of the results. For instance, if a decision is unstable in the sense that it is very sensitive already for small contractions of the intervals in the information frames, it may violate certain norms accepting only stable situations (cf. [2]), and for this reason is undesirable.

Acknowledgements

This work was supported by the Swedish Foundation for International Cooperation in Research and Higher Education (STINT). The author is also indebted to Magnus Boman and Mats Danielson, DECIDE Research Group, for discussions and remarks concerning the issues treated in this paper.

References

[1] S. Abiteboul and V. Vianu, "Equivalence and Optimization of Relational Transactions," *Journal of ACM*, vol. 35, pp. 130–145, 1988.

[2] M. Boman, "Implementing Norms through Normative Advice," Proceedings of ICMAS'96 Workshop on Norms, Obligations and Conventions, 1996.

[3] M. Boman and L. Ekenberg, "Eliminating Paraconsistencies in 4-valued Cooperative Deductive Multidatabase Systems with Classical Negation," Proceedings of Cooperating Knowledge Based Systems, pp. 161–176, Keele University Press, 1994.

[4] CCITT, *CCITT Recommendation Z.100: Specification and Description Language SDL (Blue Book, Volume X.1-X.5)*: ITU General Secretariat-Sales Section, Places des Nations, CH-1211 Geneva 20, 1988.

[5] M. Danielson, "A Computational Framework for Decision Analysis," PhD Thesis *Department of Computer and Systems Sciences*: Royal Inst. of Technology, under preparation.

[6] M. Danielson and L. Ekenberg, "A Framework for Analysing Decisions under Risk," to appear in *European Journal of Operational Research*, 1997.

[7] G. Davies and L. Ekenberg, "Using First-Order Theorem Proving for Analysing Transaction Systems," 1996, Submitted.

[8] L. Ekenberg, "Detecting Conflicts in Multi-Agent Systems," to appear in Proceedings of 2nd Australian DAI Workshop, LNAI, Springer Verlag, 1996.

[9] L. Ekenberg, "Modelling Decentralised Decision Making," Proceedings of Proceedings of 2nd International Conference on Multi-Agent Systems, ICMAS'97, 1996.

[10] L. Ekenberg, M. Boman, and M. Danielson, "A Tool for Coordinating Autonomous Agents with Conflicting Goals," Proceedings of 1st International Conference on Multi-Agent Systems, ICMAS'96, pp. 89–93, AAAI/The MIT Press, 1995.

[11] L. Ekenberg, M. Danielson, and M. Boman, "A Tool for Handling Uncertain Information in Multi-Agent Systems," in *Distributed Software Agents and Applications, Proceedings of MAAMAW'94*, J. W. Perram and J. P. Müller, Eds.: Lecture Notes of Computer Science, Springer-Verlag, 1996, pp. 54–62.

[12] L. Ekenberg, M. Danielson, and M. Boman, "From Local Assessments to Global Rationality," *International Journal of Intelligent and Cooperative Information Systems*,, vol. 5 Nos 2 & 3, pp.315–331, 1996.

[13] L. Ekenberg, M. Danielson, and M. Boman, "Imposing Security Constraints on Agent-Based Decision Support," to appear in *Decision Support Systems International Journal*.

[14] L. Ekenberg and P. Johannesson, "Conflictfreeness as a Basis for Schema Integration," Proceedings of CISMOD-95, pp. 1–13, Springer-Verlag, 1995.

[15] L. Ekenberg and P. Johannesson, "A Formal Basis for Dynamic Schema Integration," Proceedings of 15th International Conference on Conceptual Modelling ER'96, Lecture Notes in Computer Science, 1996.

[16] L. Ekenberg and F. Widebäck, "Formal Verification of the Supply of Power in a Nuclear Power-Plant using Propositional Logic," Proceedings of 2nd ENCRESS Conference, 1996.

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

[18] M. J. Katz and J. S. Rosenschein, "Verifying Plans for Multiple Agents," *Journal of Experimental and Theoretical Artificial Intelligence*, vol. 5, pp. 39–56, 1993.

[19] S. Kraus and J. Wilkenfeld, "The Function of Time in Cooperative Negotiations," Proceedings of Ninth National Conference of Artificial Intelligence, pp. 179–184, 1991.

[20] J. S. Rosenschein, "Synchronization of Multi-Agent Plans," Proceedings of National Conference of Artificial Intelligence, pp. 115–119, 1982.

[21] C. A. B. Smith, "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistic Society, Series B*, vol. xxiii, pp. 1–25, 1961.

[22] G. Stålmarck, "System for Determining Propositional Logic Theorems by Applying Values and Rules to Triplets that are Generated from a Formula," US. Patent No. 5 276 897 1995.

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