Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

A Multi-Agent System Framework for Collaborative Decision Making Under Uncertainty

Ortiz-Arroyo, Daniel; Larsen, Henrik Legind

[Link to publication from Aalborg University](#)

A Multi-Agent System Framework for Collaborative Decision-Making under Uncertainty

Daniel Ortiz-Arroyo†, Henrik Legind Larsen‡
Department of Computer Science and Engineering
Aalborg University, Esbjerg
Niels Bohrs Vej 8, 6700 Esbjerg, Denmark
†*{do@cs.aue.auc.dk}*
‡*{legind@cs.aue.auc.dk}*

Abstract— In *Multi-Agent Systems (MAS)*, joint goals are achieved collectively through collaboration. In these systems no single agent owns all the knowledge required to solve a problem; the knowledge of a problem solution is distributed among all agents. Thus, collaborative decision-making must take into account the viewpoint of all group members. Moreover, solving complex problems in MAS usually involves making decisions with partial, contradictory, or uncertain information.

This paper presents a novel framework to guide and facilitate the decision-making process of multiple specialized agents working collaboratively to achieve common goals. Our proposed framework is able to estimate the consensus level existing among agents, jointly with the downside risk associated when performing decisions with partial or imprecise information. Additionally, mechanisms are included in the framework to suggest the best course of actions that will likely increase consensus and reduce risk.

Keywords: *multi-agent system, collaborative decision-making, risk, uncertainty.*

1. INTRODUCTION

An agent may be described as autonomous, social, reactive, and proactive software entity. Agents get information from the environment, make local decisions and perform some actions. In *Multi-Agent Systems (MAS)* no single agent owns all the knowledge required to perform complex tasks. Moreover, agents may use specialized knowledge, resources, and sources of information. Thus, in such environment agents have to coordinate their actions to meet global constraints and achieve common goals. Coordinating agents' activities requires some sort of a *collaboration* mechanism. *Organizational* structures enable agents to collaborate hierarchically. Other known mechanisms of coordination are *contracting*, *auctioning*, *negotiation*, and *argumentation* [5][6].

Hierarchical organizations give rise to the well-known *master-slave or client-server* coordination techniques. Using this type of organization, roles, authority and communication are known *apriori*. The master agent initially distributes fragments of a plan to its collaborators. Afterwards, agents accomplish their plans and report their results to the master agent.

Another well known organizational approach to coordination is *blackboard* negotiation. In this mechanism agents coordinate their actions through a *blackboard* in which they can post and read knowledge sources.

Organizational structures based on the application of economic concepts have been also proposed. Examples of this approach are the centralized and decentralized *markets* [5].

In the *contracting* protocol for MAS, a manager agent divides a problem into sub-problems and then searches for contractors. Contracts are awarded to the best *bid* sent by a contractor agent. In such contracting systems an agent may be simultaneously a manager and a contractor [5].

Auctioning mechanisms proposed for MAS are used to allocate goods to *bidder* agents.
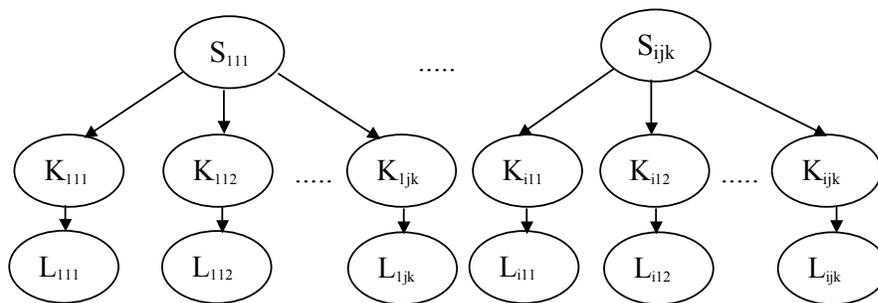
Figure 1 Bayesian Network used to handle uncertainty

Different models used in these auctioning mechanisms give rise to the *English, Dutch, Vickery* and *First-Price Sealed-Bid* type of auctions [8].

A common approach to coordinating MAS activities is *negotiation* [5]. Negotiation techniques have been either classified as based on *game-theory*, a *plan*, *AI*, or inspired by *human* behavior. Unfortunately, since some sort of negotiation is also included in most coordination protocols the distinction between negotiation and the other mechanisms is rather fuzzy. One significant characteristic of negotiation is that it allows conflict resolution. In order to negotiate effectively, agents may reason about the beliefs, desires, and intentions of other agents.

Negotiation via *argumentation* has been proposed as an effective coordination mechanism. In this protocol an agent attempts to convince other agents of performing some actions by exchanging proposals and counter-proposals. These proposals are backed by arguments that summarize the reasons why the proposals and counter-proposals should be accepted [7].

Collaborative decision-making (CDM) refers to a group of decentralized agents that collaborate to achieve common goals, which are beyond the capabilities of individual agents. In CDM the viewpoint of all group members should be taken into account since no single agent has all the knowledge required to make the best decisions. Moreover, solving complex problems that involve partial, contradictory or uncertain information creates new challenges in CDM.

Bayesian networks (BN) have been proposed to model uncertainty in intelligent systems [4]. Based on the probability theory, BN are directed acyclic graphs that represent conditional dependencies between random variables. BN provide an intuitive graphical representation of knowledge, allowing reasoning under uncertainty. Agents' design based on representing their knowledge with Bayesian Networks enables them to handle uncertainty [9].

In this paper we present a novel framework for multi-agent collaborative decision-making under uncertainty. The framework employs a hierarchical organization where agents participate using their knowledge to evaluate a group of alternatives. The uncertainty that occurs in the decision-making process is modeled using Bayesian Networks. Additionally, the level of consensus existing among a group of agents is calculated jointly with the risk involved in performing decisions under uncertainty. Mechanisms are put in place to provide the best course of action that will increase consensus and reduce risk.

In Section 2 it is explained how decisions are modeled in our framework. Afterwards, agent's decision-making architecture and the activities performed by those agents to collaboratively make decisions are explained. The process used to evaluate risk and reach consensus is explained in the following Section. Finally, in Section 6 we describe future work and provide some conclusions.

2. MODELING DECISIONS

Decisions are modeled (DM) in our collaborative decision-making framework with the following *tuple* of elements:

$$DM = <Ag,A,C,W, T,K,L> \text{ where}$$

$Ag = \{ ag_c, ap_1, ap_2,...ap_i\}$ is set of agents
$A = \{a_1, a_2,...,a_j\}$ is set of alternatives
$C = \{c_1, c_2,...,c_k\}$ is set of criteria
$W = \{w_1, w_2,...,w_k\}$ is set of preference models
$T = \{t_1, t_2,...,t_k\}$ are types associated to criteria

$K=\{k_{111},k_{112},...,k_{ijk}\}$ agent's   knowledge evaluations
$L=\{l_{111}, \quad l_{112},...,l_{ijk}\}$ agent's   confidence evaluations

Ag represents the group of agents involved in a collaborative decision. Among this group of agents, there is a single special *coordinator* ag_c that controls the whole decision-making process. *A* is the set of alternatives, which are evaluated individually by all the agents. *C* is the group of criteria that agents employ to evaluate each alternative. *W* represents the preference of each agent regarding the criteria. *T* is the criteria type, which may be *{qualitative, quantitative}*. Qualitative criteria are used to assess *{yes, no}* values whenever an alternative fulfills a criterion according to agent's beliefs.

A quantitative criterion may be one of two different types *{less is better, more is better}*. Quantitative criteria are used in situations where numerical data is available to evaluate an alternative. The data for quantitative criteria is represented in the model by the set of values *{min, best, max}*, containing the range of numerical values available for a criterion of this type. *Min* and *max* are the minimum and maximum values used to grade a quantitative criterion. The coordinator sends these values jointly with the criteria. However, *best* is a value that is assigned by a participant agent in the decision-making process.

K is the knowledge (in the range 0.5 to 1.0) an agent has on a particular criterion-alternative pair. A value of *K=0.5* means the agent does not have any knowledge. A value of *K=1.0* means the agent is an "*expert*" with perfect knowledge. *L* represents the confidence of an agent that an alternative will fulfill a criterion. *L* values are set in the range *0.0 to 1.0*; *L=0.0* meaning no confidence and *L=1.0* full confidence that an alternative will fulfill a particular criterion. *K* and *L* together represent the total *belief* of an agent concerning the alternatives.

The relationship between *K* and *L* variables is modeled using Bayesian networks. In the Bayesian Network shown in Figure 1 node S_{ijk} represents a satisfaction node, i.e. the certainty that alternative *j* will satisfy criteria *k* according to agent *i*. Node K_{ijk} has two degrees of freedom (*true, false*) and L_{ijk} has only one (*true*).

The probability distribution for the satisfaction is given by:

$$P(S_{ijk} = yes) = \alpha \prod_i (K_{ijk}L_{ijk} + (1 - K_{ijk})(1 - L_{ijk}))$$

being α the normalization factor:

$$\alpha = 1/\prod_i (K_{ijk}L_{ijk} + (1-K_{ijk})(1-L_{ijk})) + \prod_i (K_{ijk}(1-L_{ijk}) + (1-K_{ijk})L_{ijk}))$$

In previous equation K_{ijk}, L_{ijk} represent the knowledge and confidence respectively that alternative *j* fulfills criterion *k* according to agent *i*. The expected value of an alternative is given by:

$$E(A_j) = \sum_k W(C_k)P(S_{ijk} = yes)$$

Where $W(C_k)$ represents the criterion preference model of an agent. The preference model is a *weight* with range 0.0-1.0 that an agent employs to indicate how important a criterion is compared to all others. The coordinating agent will select the alternative with the highest expected value.

For a quantitative *less is better* criterion type, the value of *best* is preferred to be closer to *min*. Contrary, for a *more is better* criterion, the *best* value of a criterion is preferred to be closer to *max*. An agent uses a simple function to convert an estimated value of *best* in a quantitative criterion into its corresponding evaluation of knowledge. This function maps the range of values *{min, best, max}* into a single knowledge value given by:

$K_{lib}=1-(best-min)/2(max-min)$ for a *less is better* criterion

$K_{mib} = 1-(max-best)/2(max-min)$ for a *more is better* criterion

The confidence for quantitative criteria is also determined using *{min, best, max}* values according to the following equations:

$L_{lib} = (max-best)/(max-min)$ for a *less is better* criterion or

$L_{mib} = (best-min)/(max-min)$ for a *more is better* criterion

Being *max ≠ min* for both *K* and *L* equations describing quantitative criteria.

Figure 2 Participant agent decision-making architecture

Figure 3 Communication protocol between coordinator and agents

Notice that an agent is enabled to use new values for *max* and *min* instead of those sent originally by the coordinator. This may occur when an agent gets additional information from the environment or its knowledge base indicates that in some specific criterion-alternative pair more accurate values should be used for *min* and *max*. However, *best* should always be an intermediate value in that range and is always set by the evaluator agent. If an agent uses new values for *max* and *min,* these values will be sent to the coordinator agent jointly with its evaluations of *K* and *L*. As explained in Section 3, the coordinator will use these values to perform an evaluation of risk.

3. A FRAMEWORK FOR COLLABORATIVE DECISION-MAKING

Before exercising the collaborative decision-making mechanisms described in this paper, a series of steps need to be performed by the MAS. These steps are:
1. Agents start by identifying a potential to solve collaboratively a problem.
2. A group of agents is formed with decision-making purposes.
3. A coordinator agent generates alternatives and criteria, sending the data to participant agents.

Afterwards, during the iterative decision-making process the next steps are performed:
1. Agents independently perform their evaluations on the alternatives assessing the values of $<K,L,W>$ and send those evaluations to the coordinator.

2. The coordinator receives agent's evaluations and calculates the level of satisfaction for each alternative, the risk, and consensus level.
3. If a minimum consensus threshold has not been reached on a particular alternative and/or the risk obtained is too high, the coordinator agent determines which alternatives are more susceptible of changing a decision. Subsequently, it suggests participant agents the actions that will likely reduce risk and improve consensus level. Otherwise, if the risk is low and consensus is high on a particular alternative, the coordinator agent selects the best alternative and stops the decision making process.
4. Step 3 is repeated until the consensus level has reached a minimum threshold. However, when the consensus level does not reach the minimum level, the coordinator suggests agents to start a negotiation process to resolve conflicting viewpoints.

Figure 2 shows the decision-making model included in a participant agent's architecture.

As Figure 2 shows, agents perform an evaluation of their knowledge and confidence regarding each criterion-alternative pair using their preference model and knowledge base. Quantitative criteria is processed to obtain the values for $<K,L>$ in the way described in Section 2. These values are sent to the coordinator. The local planner uses the alternative sent by the coordinator to execute the appropriate set of actions.

An agent grades each criterion with respect to all other criteria, requiring $\frac{k(k-1)}{2}$ comparisons when k criteria are used. Moreover, a pair-wise evaluation of j alternatives with k criteria requires $\frac{k(j(j-1))}{2}$ evaluations for a total of $\frac{k((k-1)+j(j-1))}{2}$ pair-wise comparisons. To avoid the quadratic explosion in pair-wise comparisons when a relatively large number of criteria and alternatives are used, the coordinator generates a limited number of alternatives and criteria, being the number of alternatives generated larger than the number of criteria i.e. $k+1 < j$. As shown in [1], this assumption reduces considerably the number of pair-wise comparisons that each agent is required to perform. In general it is important to reduce the number of comparisons because the more comparisons are, the higher the communication bandwidth and the likelihood that erroneous data is introduced.

Agents perform evaluations to generate triplets $<K,L,W>$ for each criterion-alternative pair. Then, those evaluations are sent to the coordinator. Figure 3 shows the protocol used by the multiple agents when communicating with the coordinator.

With the information received from participating agents the coordinator builds dynamically a Bayesian Network that is used to calculate the level of satisfaction with which each alternative fulfills the criteria.

Figure 4 shows the decision-making architecture of the coordinator agent. As Figure 4 shows, the coordinator generates and sends the set of alternatives and criteria that participant agents will use to perform their evaluations. The alternatives and criteria generated depend on the current status of the task to be performed by the agents. A global planner keeps track of the current status of a task to make sure that agents reach the goals set by the global planner. Bayesian networks are used to model the uncertainty contained in the decision-making process and to calculate the satisfaction level with which every alternative fulfills the criteria. Sensitivity analysis is performed by adding nodes to the Bayesian network to determine which alternatives are more susceptible in changing a decision. The risk analysis module performs Monte Carlo simulations on agents' evaluation space to calculate how much the satisfaction level of an alternative changes with the current level of uncertainty in data.

4. REACHING CONSENSUS AND HANDLING RISK

The coordinating agent uses the data received from participanting agents to calculate the risk associated with each alternative. In a quantitative criterion the range $|max - min|$ is directly proportional to the risk associated in selecting an alternative. Risk is calculated by randomly sampling each criterion alternative pair using Monte-Carlo simulations. For qualitative criteria a uniform probability distribution is employed. For quantitative criteria a *Beta* probability distribution was used instead.

The *beta* distribution ($\beta(a,b)$) was implemented with parameters $\beta(3.0, f(min,max))$ and built from *Gamma* distributions using $\beta(a,b)=(\Gamma(a+b)/\Gamma(a)\Gamma(b))(1-x)^{a-1}x^{b-1}$. Figure 5 shows the graph of the *Beta* distribution with the quantitative *{min,best,max}* values used to define its shape.

The downside risk is found by combining the probability of being below the expected satisfaction, multiplied by the consequence of missing the ideal target value. This product is then used to compute how much lower the expected value might become if things do not go as planned. Downside risk calculations are performed for every alternative.

After assessing the risk, the framework suggests the best course of actions that will likely reduce risk using the rule-based expert system. To perform this task two new extra nodes are inserted in the Bayesian network shown in Figure 1. One node is set up with $K=1.0$, $L=1.0$ representing an *"expert"* indicating that the alternative will fulfill the criterion. Additionally, the second node is inserted with $K=1.0$, $L=0.0$ values. In this last case the new node simulates another *"expert"* indicating that the alternative will not fulfill the criteria. Using this technique the rule-based expert system located in the coordinator shown in Figure 4 is able to determine quantitatively, not only how much the satisfaction levels of an alternative will change when the information of an *"expert"* opinion is added, but also which alternatives have the potential (and are more susceptible) of changing a decision. Once this
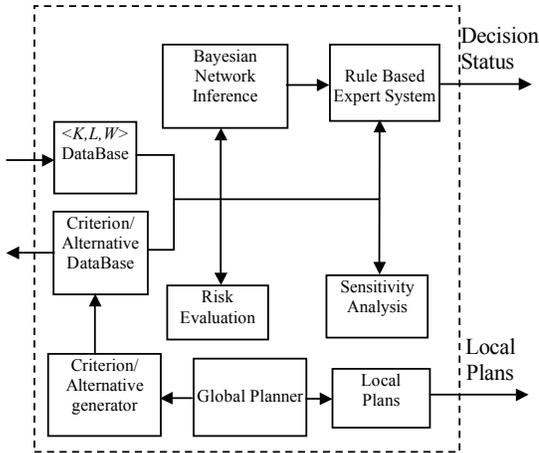
Figure 4 Coordinator's decision making
architecture

Figure 5 Beta distribution used for quantitative criteria

group of *sensitive* alternatives is detected, the expert system will indicates to all agents that further information on the sensitive alternatives is required to reach the minimum threshold of acceptable consensus and/or risk.

The level of consensus existing among the agents is calculated by considering a multidimensional space where all agents' evaluations consisting of multiple $<K,L>$ pairs (one per criterion/alternative pair) are placed. Then, the overall centroid (or mass center) of all evaluations performed by all agents is determined. The centroid of evaluations corresponding to each agent is also calculated separately. Finally, the average distance from the centroid of every agent's evaluations with respect to the overall centroid representing all agents' evaluations is calculated. This last value is the approximate measure used to represent the level of consensus existing among the group of agents.

The minimum threshold of *consensus* used in our framework is arbitrarily selected as 80%.

5. CONCLUSIONS

In this paper we proposed a new approach to reach consensus in collaborative decision-making problems for MAS. Our proposal is based on a hierarchical MAS organization in which a coordinator agent maintains control of the whole decision-making process. The coordinating agent tries to find iteratively the best alternative that will fulfill all criteria using all agents' evaluations.

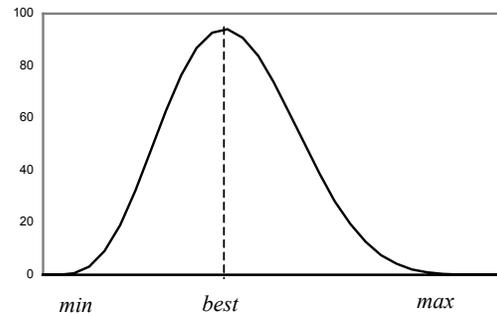The framework proposed in this paper is capable of processing quantitative and qualitative criteria, while handling uncertainty at same time. To assess the risk in performing a decision, Monte-Carlo simulations are employed. These simulations are executed to obtain the downside risks associated with selecting an alternative. The effect of achieving high risk on an alternative is to lower the satisfaction level of such an alternative making it less attractive for selection. When high risk and low consensus are found in the decision-making process, a rule-based expert system included in the coordinator indicates the steps that will likely reduce risk and increase consensus among the agents. However, when it is not possible to reach an adequate level of consensus the coordinator agent suggests agents to start a negotiation process to resolve their differences. Finally, once conflicts are resolved, the collaborative decision-making process is continued iteratively until a low risk, high consensus alternative is selected.

The proposed framework was tested in a Java-based simulation environment. Currently we are implementing the decision-making modules and protocols described in the paper on top of a FIPA-compliant multi-agent system.

Future work will address the problem of implementing heuristics to generate values for the *knowledge* and *confidence* variables, given the knowledge base of an agent, a process known as *aggregation*. In the current version of the framework these values are statically assigned for simulation purposes. Finally, we are evaluating other coordination mechanisms that could provide a more robust architecture for our collaborative decision-making framework.

6. ACKNOWLEDGEMENT

7. BIBLIOGRAPHY

[1] Triantaphyllou, E., "Reduction of Pairwise Comparisons in Decision Making via a Duality Approach," *Multi-Criteria Decision Analysis,* Vol. 8, No. 6, pp. 299-310, July 2000.

[2] Bruce D'Ambrosio, "Collaborative Decision Making", White Paper, Robust Decisions Inc., www.robustdecisions.com.

[3] Panzarasa, P., Jennings, N. R. and Norman, T. J. (2002). "Formalizing collaborative decision-making and practical reasoning in multi-agent systems" *Journal of Logic and Computation* 12(1):55-117.

[4] Jensen, F. V., *Bayesian Networks and Decision Graphs*. Ed. Springer-Verlag 2001.

[5] Nwana H.S., Lee L., Jennings, N. R., "Coordination in software agent systems". *BT Technology Journal*. Vol 14 No 4 October 1996.

[6] Sycara, K. P., "Multiagent Syatems"*, Journal of American Association Artificial Intelligence*, Summer 1998.

[7] Kraus, S., "Automated Negotiation and Decision Making in Multiagent Environments", *Lecture Notes in Computer Science*, Ed. Springer-Verlag June 2003.

[8] Anthony P, et al., "Autonomous Agents for Participating in Multiple Online Auctions", *Proceedings of IJCAI, Workshop on E-Bussiness and the Intelligent Web,* 2001.

[9] Sangüesa R., Cortés U., "The Bayesian Agent: an incremental step for learning agents working under uncertainty". *Vim Project Winter Workshop*, November 1