



Coalition-based Multi-agent Approach for Implementing Ethics: An assistive application case-study

Nadia Abchiche-Mimouni, Etienne Colle

► To cite this version:

Nadia Abchiche-Mimouni, Etienne Colle. Coalition-based Multi-agent Approach for Implementing Ethics: An assistive application case-study. The Eleventh International Conference on Adaptive and Self-Adaptive Systems and Applications (ADAPTIVE 2019), May 2019, Venice, Italy. hal-03226425

HAL Id: hal-03226425

<https://hal.science/hal-03226425>

Submitted on 15 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Coalition-based Multi-agent Approach for Implementing Ethics

An assistive application case-study

Nadia Abchiche-Mimouni

IBISC, Univ Evry, Université Paris-Saclay,
91025, Evry, France
email: nadia.abchichemimouni@univ-evry.fr

Etienne Colle

IBISC, Univ Evry, Université Paris-Saclay,
91025, Evry, France
email: etienne.colle@univ-evry.fr

Abstract— This paper presents an adaptive multi-agent approach based on coalitions for ambient assisted living applications. Adaptation is crucial because the challenge is to deal with a dynamic environment in order to provide adequate services to an elderly or a sick person at home. Moreover, it is necessary to take into account constraints such as degree of urgency of the service, intrusion level of the system and person's privacy. Ethical dimension is then important for the acceptability of such applications. The evolution of the degree of intrusion based on the degree of urgency and the availability of the communication devices of the ambient environment are particularly targeted by considering ethical dimension. The results show that not only ethics consideration allows better acceptability of the system, but also the performances are improved.

Keywords-adaptation; agents coalitions; ethics.

I. INTRODUCTION

Adaptivity is widely studied as a capability that makes a system able to exhibit intelligent behavior. Moreover, software increasingly has to deal with ubiquity, so that it can apply a certain degree of intelligence. Our specific context is to assist an elderly or a sick person in loss of autonomy at home by providing assistive applications based on cooperation among a robot and Communication Objects (CO). Maintaining such people at home is not only beneficial to their psychological conditions but helps to reduce the costs of hospitalizations. Ambient assistive robotics can be defined as an extension of ambient intelligence, which integrates a mobile and autonomous robot and its embedded sensors and the CO present in the house. The interaction among the components in such systems is fundamental. Arnand & al. [2] the authors presented a coalition-based multi-agent system (MAS) for implementing an ambient assistive living framework that takes advantage of an Ambient Environment (AE): a robot and its embedded sensors, cooperating with a network of COs. The aim is to provide a service to the person in an adaptive way. A coalition of agents proposes a set of data and the way of combining these data in order to offer the desired service. Adaptation is needed because the context is dynamic and difficult to predict. Depending on the context, the same service can be achieved by different combination of the data. A MAS reifies the sensors, the CO and the robot, allowing the cooperation by means of coalitions formation. The agents

combine the data according to their availability and the relevance. Moreover, the system has to deal with privacy and intrusion level so that one minimizes causing inconvenience. This work is based on our previous system COALAA (Coalitions for Ambient Assisted living applications), which is a coalition-based approach for implementing ambient assisted applications [1]. An improvement is proposed by: (1) embedding a Rule-Based Reasoning (RBR) module in the agents in order to reason about the coalition formation criteria, and (2) extending the scope of the adaptiveness to ethical, functional... The new approach can be considered as a general approach for implementing adaptation in ambient assisted applications. New CO can be added in a dynamic way and the way of forming the coalitions can be tuned by the user by introducing new rules in the system.

The rest of the paper is organized as follows: Section II presents an ambient assisted living approach based on multi-agent coalitions. Section III presents a generalized approach which deals with ethical dimension. Section IV highlights its benefits and shows the results validation. Section V concludes with some improvements and perspectives.

II. COALITION-BASED APPROACH AMBIENT

The principle of coalitions aims at temporarily putting together agents for reaching a common goal. The works [6][8][9] illustrate the relevance of coalition-based approaches for adaptiveness. The methods are various: either incremental, random or centralized. But, all of them proceed in two stages: (1) the formation of agent coalitions according to their ability to be involved in achieving a goal and (2) the negotiation stage among the coalitions in order to choose the one that provides the closest solution to the goal. The interests of the coalition-based formation protocols are the flexibility with which coalitions are formed and straightforwardness of the coalition formation process itself. The coalitions can get rid of dynamically reorganize with local and simple rules defined in the agents.

A. COALAA

COALAA is MAS-based on a coalition-based approach for ambient agents. Each agent in COALAA encapsulates a CO and decides in a local and proactive way when and how to contribute to the required service to the person. A more general notion than a service, called an *effect* has been introduced. An effect can be a particular lighting at a precise

place of the residence or the localization of a robot. The MAS configures itself for providing a solution according to the availability of the CO and the respect of criteria. Note that the goal is not to find the optimal solution but a solution close to the required effect. In the coalition formation protocol, the obligation to obtain the required effect and an intrusion level depending on the urgency of the situation, are the most important considered criteria. They are also used during the agent reorganization while trying to achieve a desired effect. The effect obligation criteria is used in priority while the level of intrusion is modified only if needed, i.e., to acquire new data and thus to activate the sensors likely to cause discomfort to the person.




Figure 1. Architecture of COALAA

As shown in Figure 1, several kinds of components are necessary to deal with the complexity of COALAA. An effect is modelled in the form of a triple $\langle t; c; f \rangle$ where:

- $t \in T$, $c \in C$
- T a set of task labels: localization, enlightening...
- C a set of criteria: accuracy, efficiency, neighbourhood
- F a list of factors: intrusion level, urgency degree...

The designer of the system statically assigns the criteria, while the influencing factors are assigned by the end-user.

The information handled by the system is classified into two types. This so-called persistent information, related to the application domain, puts together data about the structure of the residence and the features of the CO. The second type concerns volatile data mainly the measures

provided by the sensors and the orders sent to actuators. The volatile data are distributed in each agent, while persistent data are stored in an ontology named AA (Ambient Assistance) [5][4]. The AA ontology contains four categories of information related to the application domain: The Home category for defining the structure of the environment, the CO category for knowing their characteristics and their operating mode, the User category for defining the user profile and the Task category that puts together the tasks and Services achieved by the system.

The Gateway is a module for the standardization of information exchanged between the ambient environment and the MAS. Its role is to make the agents manipulating the common information format. This standardization is necessary because of the heterogeneity of protocols from different manufacturers.

B. Agent internal architecture

The agents of the MAS are created according to the ontologies concepts. Each agent is assigned an internal architecture able to take in charge the agent adaption and reactivity by using three main parameters that are: neighborhood, history, and ability. The neighborhood sets the list of agents that are close to this agent at a given time, according to the topological distance. The history stores previous perceived information that comes from the sensors. This is a simple succession of perceived data, which helps to consider the timescale during the process of coalition formation. The ability identifies the skills of the agent, which are directly related to the encapsulated CO.

C. Agent behaviors

In the process of the coalition formation, an agent may be either initiator or candidate. Any agent whose ability can partially meet the desired effect can be a coalition initiator. The initiator exchanges messages with other agents, potential members of the coalition, called candidate agents. The communication is based on exchanges of messages between the initiator agent and candidate agents. As soon as the overall ability of the coalition is close to the desired effect, the initiator agent is pending the negotiation phase. At the end of the coalition formation, each initiator agent that is the referent of a coalition is negotiating with other initiators agents to select the winning coalition. The coalition whose ability is the closest to the desired effect is the winning one. The concept of ability is generic. In the localization application example, it is instantiated by the measure precision. The principle is simple. Each initiator agent sends a message that contains the ability obtained by its coalition. On receipt of this message, each initiator agent compares the ability of the coalition it received to its own one. If its ability is lower than that received, the coalition will be no more considered, otherwise, it is a winning coalition up to receiving a new message. Apart from the desired effect, the formation of coalitions uses other criteria such as the topological neighborhood to reduce the response time or the obsolescence of a measure when the desired effect depends on sensor data. Thus, the first step is the identification of candidate neighbors according to its own

location in the environment (defined by the topological distance) and the desired effect. The aim of this strategy is to ensure that a result will be provided. For that purpose, the first selection criteria considered is the topological distance. Once all candidate agents are known, each initiating agent continues the selection of candidates based on the recent measures criteria. When no coalition is able to meet the desired effect, a new search for a successful coalition is restarted after having relaxed the constraints on certain criteria. Indeed, it is possible to increase the level of intrusion of the system despite of the tranquility of the person at home. This authorization to increase the level of intrusion allows, for example, operating a pan-tilt camera of the robot in order to acquire new measures and restart the process of searching for a winning coalition. This point is sensitive because there is a risk of violating the person's privacy. The protocol of coalition formation is composed of two distinct steps. The first step consists in forming coalitions of agents according to their ability. The second step is a negotiation and refining phase so that the best one, in satisfying the desired effect criteria, is chosen. Figure 2 summarizes the agents' behaviors. The baseline algorithm proceeds in three steps. After initialization, the exchanges among agents follow three main actions: formation of all possible coalitions for each referent, selection of the best coalition according to the coalition precision, deployment of the winning coalition.

The agents' interaction semantic is based on speech act theory [11], allowing the agents to assign a semantic to each message by defining a message type. The most important types are: Request, Response, Initiate, Acknowledge, Accept and Negotiate.




Figure 2: Agent behaviors

D. Discussion

COALAA shows the feasibility and the relevance of coalition-based MAS for ambient assisted scenario. It also shows that it is possible to deal with privacy criteria while building the coalitions. This is due to the high degree of adaptability of the coalition formation algorithms. To fit the obligation for the system to give a result, COALAA required the user for manually assign a priority to the criteria and the bounds for the values of the criteria. The next section illustrates this weakness and shows a generalized way of solving this problem.

III. A GENERALIZED CRITERIA MANAGEMENT COALITION FORMATION (COALAA-GEN)

Figure 3 shows an example scenario. A robot in the person's home; the patient has fallen. To move towards her/him and to guide its camera to the remote caregiver, the robot has to be located first. A visual contact will then help the remote caregiver to perform a correct diagnosis of the situation. Depending if the robot is the room P1 or the room P2, the CO required are different.




Figure 3: Fall detection scenario

Figure 4 illustrates how the MAS solving this problem. More details can be found in [3]. Three kinds of CO are involved: a robot pan-tilt camera, a fixed camera and a presence detection sensor. Three respective ambient agents encapsulate these three CO: a Presence Detector Agent (APD), a Fixed Camera Agent (AFC) and a Pan-Tilt Camera Agent (APTC). Visual markers like Data matrix are associated with each camera. Following the fall of the patient, a request for a localization effect is generated in the form of a triple $\langle t; c; f \rangle$ where t is a localization task which matches with the desired effect, c matches with a singleton containing the precision criterion needed for the localization task and f matches with a set containing two influencing factors that are the intrusion level and level of urgency. In the considered scenario, we have considered a precision equal to 0.1, a level of urgency equals to 3 (three levels of urgency are considered: low=1, medium=2, high=3) and an intrusion level initialized to 0 (the less intrusion level). So, the triple becomes: $\langle \text{Locate}; f_0:1g; f_3:0g \rangle$. The Interface agent (AI) has received the desired effect and then broadcasts the request $\text{InitCoal} (\langle \text{Locate}; 0:1g; f_3:0g \rangle)$ to all the agents of the MAS. As soon as each agent receives the desired effect, it checks its ability. As all sensors in the environment have a precision that is not better than the desired effect, each agent initiates a coalition with immediate neighborhood. In this figure, only interactions with APD agent are shown. Assuming that all agents are topologically close, APD broadcast a coalition formation request by sending an

InitCoal message. Each agent receiving the initialization message checks if its ability is adequate with the request of coalition formation. If yes, it sends an acceptance message labelled AcceptCoal to be a candidate. Such a message contains the precision of the agent. APD adds progressively answer acceptance, and accumulates the abilities, which are the precision in the considered localization task. By this way, it calculates the overall ability of the coalition until it reaches that of the desired effect. Then, it sends ACKCoal acceptance to confirm the membership of the candidate to the formed coalition. The next step is to activate the coalition. The robot moves to the place designated by the coalition and guides its pan-tilt-camera to the remote caregiver. First of all, the distant user has to verify that the person is in his field of vision, so it can perform a correct diagnosis of the situation and adopt an adequate action. Conversely, if the person is not well located the system restarts searching for a new result, after having increased the intrusion level. This allows the cameras to be moved randomly so that the chances of getting a visual marker are increased. The consequence will be improving the precision of the result.




Figure 4: Interaction diagram

A. Agent rule-based reasoning module

The previous scenario shows that criteria management is critical. Indeed, obtaining a successful coalition depends on the order in which the criteria have been considered. In the above scenario, if the first considered criterion was the level of intrusion (instead of the precision), then the first result would have been the correct one. Then, the question could be the following: why can one not have a management criteria step integrated in the coalition formation process? This is the main contribution of this work. We have introduced into each agent of COALAA a Rule-Based Reasoning (RBR) module responsible of determining a priority of the criteria to consider according to the context. The RBR is also responsible of assigning and adjusting the criteria values. The RBR is used for interleaving the execution of the behaviors in a dynamic way.

The RBR is composed of a Knowledge Base (KB) and an inference engine. The KB contains a set of rules and a set of facts. The rules are given in the form of implications. The facts describe the state of the world. The inference engine is a special interpreter that controls the triggering of the rules according to the KB. The form of a rule is: *IF* <antecedent> *THEN* <consequent>, <antecedent> is the condition that must be satisfied to trigger the rule, <consequent> is the performed action when the rule is triggered. Antecedent is satisfied if the condition matches the facts in the KB.

Instead of having a procedural control, each behavior is modelled by a production rule whose activation condition is precisely the context of its execution. The behaviors of the agents are associated with trigger conditions. These conditions represent the context that makes behaviors possible to be executed. Explicit chaining between the behaviors is no more needed since the inference engine triggers the rules. For example, the AcceptCoal behavior is chained with the InitCoal behavior. So, the InitCoal behavior is executed once the AcceptCoal behavior is terminated. The rules below express in Jess syntax [12] that if an agent has in its working memory an InitCoal message and if the agent has an ability ?x, so the rule can be triggered. In this case, the core of the behavior associated with the rule is executed.

```
(defrule check-ability "accepts to join coalition if required ability"
(Message InitCoal ?x)(Ability ?x)=>(assert (Behavior AcceptCoal ?x)))
(defrule perform-ability "Create an accept message" (Behavior
AcceptCoal ?y) => (bind ?m (createMessage (AcceptCoal ?y)))
```

As said earlier, agent architecture is endowed with a RBR module responsible of a declarative reasoning process. It consists in an inference engine that implements a decision module. The facts represent the knowledge that have been extracted from the ontology, the perceived data and the exchanged messages among the agents. For that purpose, a set of rules is defined to determine, depending on the context, the most relevant criteria to consider first at each step of the coalition formation process. When the coalition proposed by the system is not a correct one, the RBR is in charge of determining the most relevant criteria to relax or to modify. The involved rules in this case are some kinds of heuristics that guide the coalition process in managing the criterion. For example, if a coalition does not include a CO whose precision is sufficient, it is advisable to relax the intrusion level. This increases the degree of freedom of the system regarding to its actions allowing the cameras to be activated or lights to be switched on. Another use of the RBR for the management criteria concerns the addition of new criterion such as data freshness. It is sometimes more relevant to consider not sufficiently precise data if they are very recent. For example, a presence detector can only inform that the person is situated in a particular room. Suppose that a particular presence detector "informs" that the person is in the room R1 and a camera "shows" that the person is in the right corner of the room R2. Obviously, the information given by the camera is more accurate, but if it is too old it should be obsolete and will not help correctly locating the person. It is suggested here to consider the date

of perceived information for determining the priority of the criteria.

Note that exception handling is not provided in the current version of the system.

IV. ADAPTIVENESS MULTI-DIMENSIONALITY

The results are obtained in a real environment composed of heterogeneous sensors and markers. The platform includes several sensors of the market and dedicated sensors developed in our laboratory. The environment is composed of three rooms equipped with a set of sensors and the robot with its own sensors. The localization is based on goniometric measurements provided by robot on-board sensors and environment sensors. These can provide localization information allowing the localization of the robot in its environment using real-time data either from the robot on-board sensors or from the sensors in the environment. COALAA-GEN has been implemented using the Jade multiagent platform [5], where each agent embeds an instance of Jess. The production rules are given as a text file input parameter to the agents.

A. Computational adaptiveness

COALAA-GEN and COALAA have been compared to the well-known CNP protocol [7]. The Figures below shows the obtained results. The tests have been performed with a dozen scenarios. Each scenario has been executed with CNP, COALAA and COALAA-GEN. For COALAA and CNP, different values for the criteria have been experimented. COALAA-GEN has been tested with the same collected data, without any user intervention for criteria management. The figures below summarize the results.

Figure 5 shows the number of formed coalitions depending on the number of agents present in the MAS. The preferred strategy in our approach is to obtain a maximum number of coalitions that meet the selection criterion. The goal is to maximize the number of solutions to meet the request to increase the chances of securing a result. The number of coalitions is less than or equal to the number of initiators. In terms of the number of formed coalitions, the Contract Net protocol is less efficient than COALAA. COALAA-GEN gives the result with fewer numbers of agents. This can be interpreted by the fact that "intelligent" criteria management helps the agents to be more relevant for coalition formation. The response times are compared (see Figure 6). This time corresponds to the time spent in calculating the coalitions, including the message exchanges. The fact that the number of coalitions that the CNP can form is lower than the number of initiators has a direct effect on the response time. It also impacts the number of exchanged messages (Figure 7). The curve representing the number of exchanged messages follows the same rate for CNP, COALAA and COALAA-GEN. However, COALAA-GEN shows a higher number of exchanged messages. Unlike the CNP, COALAA and COALAA-GEN avoid system crashes, by a progressive coalition formation, which in contrast increases the number of exchanged messages. In terms of time response COALAA, COALAA-GEN and CNP are

almost similar; CNP is slightly better in terms of response time. But in terms of obtained COALAA-GEN is the best. Indeed, a failure can be catastrophic and thus the few milliseconds delay in the response time may be insignificant, if success to complete the task is assured. This is explained by the fact that COALAA-GEN continues reorganizing until a solution is found (even with deteriorated criteria).




Figure 5. Formed Coalitions




Figure 6. Response time




Figure 7. Exchanged messages

B. Methodological and functional adaptiveness

The genesis of the MAS is done automatically in COALAA-GEN. This is a very important feature of the system. In fact, modifying the AE, by adding or suppressing CO, automatically updates the ontology and triggers automatic MAS reconfiguration. In case of such modifications, the user does not need to do any specification to make the system adapting its architecture to AE dynamic updating. This ability is qualified by methodological adaptiveness. We refer to functional adaptiveness while dealing with services that the system can offer to the user. The description of the ability of the CO used by the agents

to construct services according to the "effect description" is included in the "task" ontology part. This allows the agents to perform an automatic detection of their ability to perform an effect.

C. Ethical adaptiveness

An original specificity of our system is that it deals with ethical dimension in an adaptive way. Adding ethical values as criteria for forming the coalitions ensures this specificity. The level of intrusion of the system is modelled in such a way that it is upgraded only in case of emergency and if the user wishes to. Moreover, the personal data are stored in the equipment of the house and are uploaded only if needed by the distant caregiver and if the user has agreed. The degree of intrusion of each CO is modelled in the ontology as an attribute associated to the CO concept. The personal data are kept locally in the agent and are not stored in the distant ontology. But if the distant caregiver needs it (in case of emergency), the private data are uploaded, with a special status that is, volatile. This means that they are deleted from the distant storage as soon as they have been used. In the presented scenario, only two ethical criteria have been considered: the level of intrusion and the data privacy. They have been modelled as criteria for coalition formation. Adding new criteria is performed by adding new rules:

```
(defrule crit-manag-001 "add new criteria" (Crit ?type ?name)
=>(assert (Coal Crit ?name)))
(defrule crit-manag-002 "assign new criteria for coalition formation"
=>(modify (Coal Crit ?name)))
```

D. Control adaptiveness

The fact that an inference engine has been employed instead of a procedural algorithm has a direct effect on the intelligence of the system. The behaviors are involved only when their associated rules are triggered, which are themselves triggered when some declarative conditions are met. Since the conditions of the rules can be modified without any procedural modification, the control of the execution of the behaviors is completely adaptive. The user can control and modify the execution of the behaviors even at run time. Furthermore, the system is also able to detect missing information that is able to lead to the execution of a particular behavior. This is ensured by backward chaining rules. The engine seeks steps to activate rules (when necessary) whose preconditions are not met. This is illustrated by the given below:

```
(defrule ctrl-001 "alarm occurred, but no behavior to trigger"
(Alarm ?x ?y)(not (Behavior ?z ?t))
=>(assert (Backward ?z ?t)))
```

More generic the rules are, more the system intelligence can be improved.

V. CONCLUSION AND PERSPECTIVES

We have introduced a new general approach for improving adaptiveness in ambient assistive applications by adding ethical dimension. A RBR module has been embedded in the agent architecture to dynamically assign

the criterion to consider during the coalition formation process and we proposed to deal with the adaptation at different levels. The adaptiveness has been considered according to four dimensions: (1) computational dimension: during the coalition formation process, (2) functional and methodological dimension: while service modelling, (3) ethical dimension: associating the intrusion level to the degree of emergency, (4) control dimension: for behaviors triggering and criteria management. We have compared the obtained results with those previously obtained without the RBR, and we have observed that the adaptiveness has been improved without any performance degradation. The feasibility of this general approach has been showed on a usage scenario to remove the doubt of a false alarm in fall detection. The first results are promising. Current and future work concerns modelling of ethical criteria in the ontology so that one can deal with various situations and contexts. Indeed, recent works [5][10] links ethics and automated reasoning in autonomous systems and artificial intelligence.

REFERENCES

- [1] A. Andriatrimoson., N. Abchiche-Mimouni, E. Colle, and S. Galerne, "An adaptive multi-agent system for ambient assisted living," in ADAPTIVE 2012, Copyright (c) IARIA, 2012. ISBN: 978-1-61208-219-6 ThinkMind, July, pp. 85–92, <http://www.thinkmind.org/>
- [2] R. Anand, E. M. Robert, H. C. Roy, M. and Dennis, "Use of ontologies in a pervasive computing environment," in Knowledge Engineering Review, vol. 18, 2003, pp. 209–220.
- [3] F.L. Bellifemine, G. Caire, D. Greenwood, Developing Multi-Agent Systems with JADE. Wiley Eds. 2007.
- [4] V. Bonnemains, C. Saurel, C. Tessier, "Embedded ethics: some technical and ethical challenges," Ethics and Information Technology. SSN: 1388-1957 (Print) 1572-8439 (Online) 10.1007/s10676-018-9444-x, 2018, pp. 41–58.
- [5] A. Kivela and E. Hyvonen, "Ontological theories for the semantic web,". in Semantic Web. Finland, May, 2002, pp. 111–136.
- [6] M. Sims., C. Goldman, and V. Lesser, "Self organization through bottom up coalition formation," Proc. of the 2nd international joint conference on Autonomous agents and multiagent systems, 2003, pp. 867-874.
- [7] R. Smith, "The contract net protocol: High-Level Communication and Control in a Distributed Problem Solver," Journal IEEE Transactions on computers, Volume 29 Issue 12, December 1980, pp. 1104–1113.
- [8] L.-K. Soh and C. Tsatsoulis, "Reflective negotiating agents for real-time multisensor target tracking," IJCAI'01, 2001, pp. 1121-1127.
- [9] L. Soh and C. Tsatsoulis, "Allocation algorithms in dynamic negotiation-based coalition formation," AAMAS02 Workshop 7 "Teamwork and coalition formation", pp. 16–23.
- [10] Y. Han, S. Zhiqi, M. Chunyan, C. Leung, V. R. Lesser and Q. Yang, "Building Ethics into Artificial Intelligence," in Proc. of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18), 2018, pp..
- [11] J. Searle, "Speech acts. an essay in the philosophy of language," Cambridge University Press, 1969. 5527-5533
- [12] E. Friedman-Hill, Jess in Action: Java Rule-Based Systems (In Action series) Paperback – July, 2003 by Ernest Friedman-Hill.