

An Agent-Based, Multilevel Welfare Assessment Model for Encompassing Assignment and Matching Problems

Antoine Nongillard, Sébastien Picault

► **To cite this version:**

Antoine Nongillard, Sébastien Picault. An Agent-Based, Multilevel Welfare Assessment Model for Encompassing Assignment and Matching Problems. PAAMS'2017 - 15th International Conference on Practical Applications of Agents and Multi-Agent Systems, Jun 2017, Porto, Portugal. pp.196-208, 10.1007/978-3-319-59930-4_16 . hal-01536632

HAL Id: hal-01536632

<https://hal.inria.fr/hal-01536632>

Submitted on 12 Jul 2017

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.

An Agent-Based, Multilevel Welfare Assessment Model for Encompassing Assignment and Matching Problems

Antoine Nongaillard¹ and Sébastien Picault^{1,2}
name.surname@univ-lille.fr

¹ Univ. Lille, CNRS, Centrale Lille, UMR 9189 – CRISTAL (SMAC), F-59000 Lille, France

² Bioepar, INRA, Oniris, La Chantrerie, 44307 Nantes, France

Abstract. Multi-Agent Systems (MAS) have been applied in recent years to assignment or matching problems in order to enhance privacy in preferences and constraints for individuals, and to facilitate the distribution of solving. A further step in this direction consists in using the organisational structures provided by MAS. Thus, in this paper, we rely upon the capability of Multilevel MAS to reify intermediate viewpoints between the individual and the collective levels, in order to encompass matching or assignment problems. Therefore we define a meta-model for assessing the welfare of agent groups with respect to relevant metrics, so that these groups are able to elaborate solutions that improve the collective well-being without forcing them to disclose all their private information. Finally, we outline the general principles for distributed solvers designed for this type of modelling.

Keywords: Multi-Level Agent-Based Modelling, Social Choice Theory, Assignment and Matching Problems

1 Introduction

Resource allocation and matching problems are two large and generic families of problems. They have been deeply studied by various methods, both for their modelling and their resolution. They also have been extremely specialised to meet many contexts and heterogeneous application, resulting in an abundance of algorithms, that rival in front of increasingly complex instances.

In this paper, we propose a multilevel agent-based modelling of these problems. We assume that each component of the system modelled is represented by an agent. Moreover, each component is able to calculate its individual satisfaction with respect to its environment, its own constraints and subordinate agents. This approach is intended to elicit objectives and constraints for each system element. We also propose the principles of a distributed solving method, based upon the satisfaction of each agent towards its own constraints, which allows some privacy.

The paper is organised as follows. In the next section, we review the definition of allocation problems and matching problems. We also discuss the conventional modelling and solving methods. We show the limits of their expressiveness and we present our contribution. We explain the meta-model we propose for measuring multilevel welfare

considering all these aspects. Then, we sketch the principles of resolution algorithms. Before concluding, we discuss advantages and limitations of our approach. Eventually, we present the pending questions that requires further works.

2 Literature and contribution

2.1 Centralised approaches

A problem of resource allocation aims at distributing a resource set among a population of individuals, optimising a goal defined as the aggregation of individual measures. A matching problem aims at grouping individuals, by optimising a goal defined as the aggregation of assessments made by each individual on other members of their group. In both cases, objectives are often expressed as an aggregation of individual evaluations using the social choice theory [2, 3] for instance. While these two families have similarities, they are always addressed as problems of a different nature, in particular because of the possibility or not for “resources” to express preferences towards other members of their group. They are thus resolved by algorithms dedicated to specific families, or even specific sub-problems.

Various paradigms can model these problem families including the constraint satisfaction problem (CSP) multi-objective optimisation (which aims at finding compromise between various potentially conflicting objectives), or multi-criteria optimisation (which seeks to optimise a composite metric).

In terms of resolution, the most common methods are centralised and based on complete information. Indeed, all preferences or constraints are public and manipulated by an overall solver. Among the most known algorithms in this field, we can mention of course the Hungarian algorithm for assigning resources [9] and the Gale-Shapley algorithm for matching problems [8].

These methods usually focus in identifying optimal solutions in absolute terms, by making the strong assumption of a full and public information. Mechanisms for achieving such a solution from a given starting point is not a concern of these methods.

2.2 Distributed approaches

In recent years, these problems have also been modelled using the multi-agent paradigm. On the one hand, one can consider the problem of resources/tasks allocation within a population of agents: MAS are then just an application area [5, 1, 18], for which it seems natural to seek distributed solving methods [10]. On the other hand, MAS can be used as a tool to solve allocation or matching problems in a distributed way.

Amongst the latter approaches, some are a mere distribution of computations [14, 4] and do not consider the *behaviours* required for agents to achieve a solution, but rather a protocol to set up (e.g., on the assumption of individual rationality) and the solution characteristics (Pareto-optimality, envy-freeness...). Other methods, on the contrary, seek to strengthen the privacy of preferences and constraints: for instance, distributed algorithms guaranteeing privacy of such information were proposed by [15] for allocation problems and [7] for matching problems. In addition, the MAS approach is especially suited to clarify the point of view of each actor involved in the resolution. This

“more natural modelling” facilitates the acquisition of expertise on the preferences and constraints, and provides opening towards multi-criteria or multi-objective optimisation. However, the measurement of welfare can only be either individual (agent level) or collective (MAS level).

However, these approaches are primarily a *distribution* under agents of individual preferences and a *distribution* of the allocation/matching solving. In fact, MAS allow to go beyond this mere distribution, taking into account relational aspects for instance (such as networks of acquaintances to specify which agents can negotiate), behavioural aspects (by introducing a variety of strategies in partners or resources selection) and organisational aspects (by the explicit introduction of agent groups sharing some interests). These factors are rarely considered in classic allocation problems or matching problems. Indeed, it presupposes a form of distribution that classic methods exclude, but it is equally rare in MAS methods.

2.3 Contribution

The proposal described here explores a complementary approach. Instead of addressing the very issue of distributing either the preferences or the resolution among a population of agents for a particular problem, we rather try to build a generic multilevel structure which allows the modelling of assignment or matching problems, and the expression of multiple concerns at the same time.

The multilevel agent-based meta-models [13, 16, 6, 17] took a growing importance within MAS in recent years. Indeed, from their origin in the field of simulation, they have spread to other areas, such as distributed constraint solving [12]. Thus, we think that they can also benefit to the types of problems we have discussed above, especially by introducing a capability to represent intermediary stages between “atomic” individuals and the overall system.

In this paper, we propose a multilevel modelling of these problems and lay the foundation of a solving method that clarify goals and constraints of each system element. This distributed solving method is based on the satisfaction by each agent of its own constraints, and consider also some privacy. Our modelling increases the variety of constraints and preferences that can be expressed at every organisation level, using the ability to build composite metrics. Our longer-term objective is to design mechanisms for agents to negotiate a path between an initial situation and a solution displaying desirable properties.

3 Proposed model

3.1 Multilevel Formalism

In this section, we present the multilevel modelling setting which underlies our approach. Especially, we assume that all relevant entities of the model (e.g. “individuals”, “resources”, “groups”...) are represented by agents and linked by a belonging relationship. Those relations are not necessarily hierarchical, since for example some problems may imply resource sharing (one resource belonging to several individuals), or allow individuals to be members of several groups at the same time.

In order to implement these features, the formalism we chose is a multilevel agent-based simulation meta-model called “PADAWAN” [16]. It has already been successfully adapted for handling distributed constraint solving in the context of cartographic generalisation [12]. Besides, in the context of allocation or matching problems, some simplifying assumptions can be made. For instance, environments, which are considered a first-order abstraction in simulation, can here be reduced to a mere set of agents, according to the “AgentSet” Design Pattern proposed in [11]. Thus, the key feature in this meta-model is that any agent can contain other agents, and conversely an agent can belong to several agents.

In the rest of the paper, we use the following definitions. The set of all agents is denoted by \mathbb{A} ; $a_1 \sqsubset a_2$ means that agent a_1 is **contained** by agent a_2 (or a_2 is **host** to a_1). Since the agents can represent individuals as well as resources or groups, the \sqsubset relationship expresses very generic belonging links (in general, a_1 and a_2 have different types, but this is not mandatory). An agent can be contained in several other agents at the same time (in order to represent multiple memberships). The \sqsubset relationship induces a *hosting graph* between the agents, which is an oriented graph. According to the original meta-model hypotheses [16], this graph must be cycle-free. Figure 1 provides an example of such a graph.

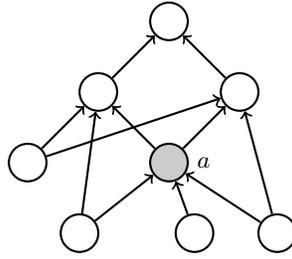


Fig. 1: Example of a hosting graph between agents. The arrow link represents the \sqsubset relation. Agent a belongs to two hosts and contains 3 other agents.

For any agent a we also define its *hosts* and conversely the set of all agents *contained* in a respectively as:

$$hosts(a) = \{a_i \in \mathbb{A} \mid a \sqsubset a_i\} \quad (1)$$

$$content(a) = \{a_j \in \mathbb{A} \mid a_j \sqsubset a\} \quad (2)$$

The *level* of an agent in the hosting graph is defined as:

$$level(a) = \begin{cases} 0 & \text{if } hosts(a) = \emptyset \\ 1 + \min_{h \in hosts(a)} level(h) & \text{otherwise} \end{cases} \quad (3)$$

This structuring of the MAS through the belonging relationship leads to an original modelling of the welfare of the agents.

3.2 Multilevel Welfare Model

Our approach consists in grounding the computation of welfare values in the very belonging relations in the MAS. To do so, we decompose the individual welfare of any agent into three contributions which can be aggregated in various ways (depending on the application domain). They represent respectively the agent as *an individual*, as *the neighbour of other agents*, and as *the host to other agents*. Thus we have:

$$w(a) = f_a(\sigma(a), \mu(a), \gamma(a)) \quad (4)$$

The f_a function can be chosen arbitrarily, depending on the situation to be modelled. Since the hosting graph is cycle-free, a consistent computation of the welfare values can be done level by level.

As an individual The $\sigma(a)$ contribution (“situation”) represents *the satisfaction of agent a as an individual, which is also situated in a given structure*. Thus, this value can be computed from the state of agent a , but also according to the *perceived* properties of its hosts and of their own “ancestors” (as shown on Figure 2-a), and defined as follows:

$$\mathcal{H}(a) = \{h \in \mathbb{A} \mid a \sqsubset h \vee \exists h' \in \mathcal{H}(a), h' \sqsubset h\} \quad (5)$$

The corresponding value can then be computed by using an operator \bigoplus^a , *specific to agent a* (and to be defined for each concrete situation), in order to aggregate the perceptions by agent a of its *situation*, i.e. its own properties and those of all agents that contain a (either directly or transitively):

$$\sigma(a) = \bigoplus_{h \in \{a\} \cup \mathcal{H}(a)}^a \chi_a(h) \quad (6)$$

χ_a , here and in what follows, is a way to express the “subjectivity” of agent a , in the sense that, depending on the target application, it can be either properties that can be accurately *measured* by agent a observing other agents, or *estimations* made by agent a .

As a neighbour The $\mu(a)$ contribution (“membership”) represents *the satisfaction of agent a as member of a group*, in other words the contribution of externalities due to the presence of other agents within the same hosts. Thus, this value must take into account the perception by a of the properties of its *neighbours*, in every host (as shown on figure 2-b), i.e. the set of agents defined by:

$$\nu(a) = \bigcup_{h \in \text{hosts}(a)} \text{content}(h) \setminus \{a\} \quad (7)$$

The perceived properties of those neighbours are then aggregated using another operator which we denote by \bigodot :

$$\mu(a) = \bigodot_{n \in \nu(a)} \chi_a(n) \quad (8)$$

As a host The $\gamma(a)$ contribution (“group representation”) represents the satisfaction of agent a as the representative of a group of agents, namely in relationship with the agents which are contained in a (i.e., $\text{content}(a)$). This satisfaction as a group is intended to measure the *collective welfare* of the agents contained in a , which can be done by aggregating the perceived properties of the agents contained in a , using a third domain-dependent operator, denoted by \bigoplus (as shown on figure 2-c):

$$\gamma(a) = \bigoplus_{m \in \text{content}(a)} \chi_a(m) \quad (9)$$

In most problems, the welfare is computed according to an “ascending” order: the collective welfare, as a measure of the quality of the assignment or matching, is calculated from the individual welfare values. In that case, the “perceived properties” that are relevant to calculate γ are simply the individual welfare of the agents contained in a :

$$\forall m \in \text{content}(a), \chi_a(m) = w(m). \quad (10)$$

To ensure the consistency of this choice, welfares must indeed be calculated by sorting agents in descending order of levels, as defined in equation (3).

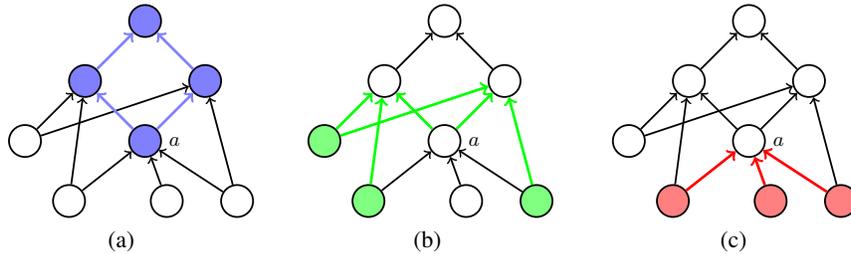


Fig. 2: Contributions to the welfare of agent a : (a) as an individual $\sigma(a)$: a and its direct or transitive hosts; (b) as a neighbour $\mu(a)$: neighbours of a in each of its hosts; (c) as a host $\gamma(a)$: agents contained in a .

4 Applications

This example is an overview of the capabilities of our meta-model to represent various points of view (hence, various objectives), specific to each agents family or to each level. It also provides a concrete case of multiple membership.

Here we consider individuals (\mathcal{I}) allowed to enrol in several associations (\mathcal{A}). These associations can gather into federations (\mathcal{F}) and can be funded either by municipalities (\mathcal{M}) or by regions (\mathcal{R}).

The objectives of such agents are clearly quite different. The individuals seek to maximise their participation in the associations offering their preferred activities, according to private time or budget constraints. The associations and federations intend

to assert their size (enrolment) so as to defend their grant requests. The municipalities aim at allocating their available budget as fairly as possible between their local associations. Similarly, the regions try to do the same between municipalities and federations. A typical situation is described by the hosting graph on Figure 3.

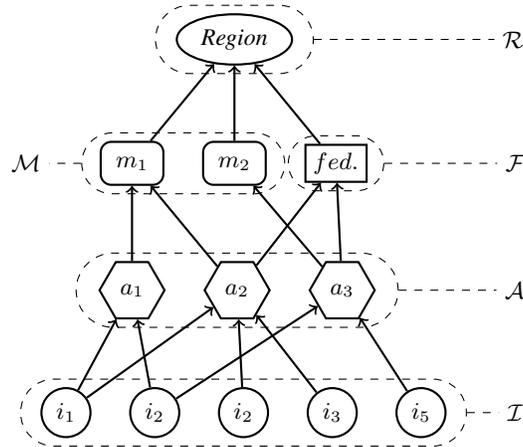


Fig. 3: Hosting graph representing a common associative network, with individuals (\mathcal{I}) enrolled in several associations (\mathcal{A}), possibly grouped into federations (\mathcal{F}) and funded by municipalities (\mathcal{M}) or regions (\mathcal{R}).

We can assume in this example that the grant allocation is based on a Nash welfare, in order to prevent the largest structures from monopolising the budget, and yet taking their enrolment into account, hence: $\oplus_{\mathcal{R}} = \oplus_{\mathcal{M}} = \prod$.

An association a which only intends to assert the number of registered members (without consideration for their satisfaction), can simply use $\chi_a(i) = 1$ for every member i and $\oplus_{\mathcal{A}} = \sum$. The weight of an association obviously depends on its belonging to a federation, thus for each of its hosts h , we assume $\chi_a(h) = 1$ if $h \in \mathcal{F}$, and $\chi_a(h) = 0$ otherwise, with $\oplus_{\mathcal{A}} = \sum$.

The individuals are rather in search of a trade-off between the participation to their preferred activities (e.g. using an affinity model) and the cumulative cost of these activities (which requires a cost matrix $(c_{i,a})$ which reflects the cost for person i to enrol in association a). Again, $\oplus_{\mathcal{I}} = \sum$ can be used. But, in addition, people are highly sensitive to their neighbours, i.e. the other individuals enrolled in the same associations. This can again be handled for instance through an affinity-based approach, and an aggregation operator $\odot_{\mathcal{I}}$ either optimistic (max), or pessimistic (min).

To summarise, the diversity of objectives within the MAS is made explicit by using a large diversity of metrics, within a structure where all agents are otherwise homogeneous. We show in the next section that this systematisation is the basis for defining generic resolution principles, which can deal with quite different situations, in opposition of classical approaches (either centralised or distributed) where each specific

problem is to be solved by its specific method. Indeed, we assume that our meta-model allows the construction of intrinsically multi-agent resolution methods, i.e. relying on local perceptions and interactions between agents, and on generic behaviours subject to context-dependent settings.

5 Solving method

At this stage, we mainly focused on the *modelling* of assignment or matching problems using a multilevel agent-based formalism. This allows an accurate expression of the viewpoints of every entity involved in the problem. Nevertheless, the advantages of such a modelling have to be also supported by concrete resolution methods.

Most of the time, the specificities of the problem to address, or of the application domain, have a deep impact on the resolution algorithms to use. The approach we propose cannot bypass this issue, and there certainly are many open questions, some of them being discussed in the next section.

Yet, we believe that our meta-model, due to its generic structure, can guide the construction of solving algorithms through the use of generic principles. The guidelines we present below are based on our preliminary work in the matter, and we are currently carrying out a thorough study on this subject.

Besides, a resolution algorithm can be assessed according to a wide range of criteria, from e.g. its computational cost and the size of manipulable instances, to the volume of exchanged messages, or the capability to optimise several metrics at the same time. Thus, searching for a convenient all-purpose method would be quite illusory.

5.1 Principles

Instead, we describe below a protocol that allows two agents at the same level to exchange a subset of their content (i.e. other agents). This protocol relies upon some simple principles drawn from [15]:

1. agent s (the *seller*, which starts the negotiation) should be ready to initiate a dialogue by offering “resources” (other agents contained in it) to agent b (the *buyer*, interlocutor of s) – and to do so, *even if this gift is potentially unfavourable*
2. agent b can propose a counterpart, but it does only if it finds this useful
3. to assess the exchange, some principle of subsidiarity applies: if both agents agree upon the favourable or unfavourable nature of the exchange, s and b can take a joint decision; otherwise, they request a mediation from their host.

5.2 A bilateral protocol

In what follows, we denote by $\wp(\mathcal{S})$ the power set of set \mathcal{S} (sometimes denoted by $2^{\mathcal{S}}$). $\Delta_a w(+X)$ (resp. $\Delta_a w(-Y)$) denotes the the variation in the welfare of agent a when it adds to (resp. removes from) its *content* the set of agents X (resp. Y). $\Delta_h w(s \xleftrightarrow{X} b)$ denotes the variation in the welfare of an agent h , host to s and b , when s and b exchange

the sets of agents X and Y . We discuss in the next section several methods to assess those values.

On this basis, the protocol for two agents s and b sharing the same host h runs as follows:

1. agent s calculates the “best” non-empty subset X^* of the agents it contains, which it can part with:

$$X^* = \arg \max_{X \in \wp(s.content) \setminus \{\emptyset\}} \{\Delta_s w(-X)\} \quad (11)$$

2. if no such subset exists, agent s can propose nothing, the negotiation is over; otherwise, s sends b the message `propose(+X)`.
3. agent b calculates a possible counterpart to the proposal of s :

$$Y^* = \arg \max_{Y \in \wp(b.content)} \{\Delta_b w(+X - Y)\} \quad (12)$$

4. if $\Delta_b w(+X - Y) > 0$, agent b sends s the message `offer(-X+Y)`: it is a firm offer, since the exchange of X against Y is favourable to b ; otherwise, b sends `propose(-X+Y)`.
5. two “simple” cases can occur:
 - $\Delta_s w(-X + Y) > 0$ and b sent `offer`: in that case, the exchange is favourable to both agents s and b , thus it can be immediately accepted
 - $\Delta_s w(-X + Y) \leq 0$ and b sent `propose`: in that case, the exchange does not benefit anyone, thus it can be immediately rejected
6. in other cases, only one agent benefits from the exchange, thus s asks its host h to pronounce on the collective interest of the exchange. The host calculates $\Delta_h w(s \xleftrightarrow{X} b)$: the exchange is accepted iff this value is strictly positive.

5.3 Primitives and experiments

This protocol relies upon primitives which appear to be fully independent from any application domain and from any kind of problem. Basically, in the multilevel meta-model, all agents can be endowed with a capability of adding other agents to their content, of removing other agents from it. From these two building blocks, several primitives can be built, such as a unilateral gift, an exchange, the transfer of agents from one level to another, etc. All these primitives compose a generic basis for the design of agents behaviours aimed at solving particular problems.

For now, this bilateral negotiation protocol, intended for agents sharing the same host, has been successfully tested in the context of a “table assignment problem”, on “small” instances (containing 9 to 15 agents to be able to compare with the optimal assignments). Tables are allowed to interact for exchanging guests, exactly as if they were mere resources: they can be handled like that because the calculation of the welfare of the tables includes the preferences of the guests through the γ contribution. Besides, in order to cope with realistic solutions, we introduced a limitation in the size of the X/Y subsets which are proposed and/or accepted during the negotiations, so as to take into account the number of seats available at each table. This reflects a situation that can occur in many concrete problems, and show that the generic protocol can be easily extended to cope with such constraints, without requiring dramatic changes.

6 Discussion and future work

The modelling and resolution principles proposed above are a first step towards a systematic study of distributed solving algorithms for assignment and matching problems modelled through multilevel agent-based representations. We have shown that such a modelling provides a substantial enhancement in the capability of expressing subtle nuances in welfare assessment, and takes explicitly into account the point of view of intermediate organisation levels. This approach to the modelling of assignment or matching problems is generic and requires new resolution methods.

As we built a meta-model which benefits from the intrinsic features of a multilevel MAS, we intend to keep those interesting properties in the solving phase, especially by endowing agents with generic behaviours, based on their local perceptions and belonging or neighbourhood relationships with other agents. The negotiation protocols which can be built under these assumptions are yet unexplored. Then, an extensive work is to be conducted so as to evaluate them in the light of classical criteria, such as the number of operations (moves) to reach a good solution, the nature of the solution itself (optimal or approximate), the volume of exchanged messages, or the quantity of individual information made public.

We have laid out the first blocks in this process by proposing the core principles of any negotiation, presenting a bilateral protocol for agents sharing a same host, and carrying out the corresponding experiments. Our ongoing work on the matter suggests several open questions. We discuss the three main of them below.

The first key issue consists in specifying which agents are allowed to interact with each other so as to exchange part of their content. In the simplest case, which we addressed above, the interacting agents share the same host, which ensures that this host is able to provide an direct arbitration when the negotiating agents disagree about the assessment of a transaction.

But, in some situations, it may be more relevant to allow more agents to interact, using for instance an acquaintance network to specify their potential negotiation partners. In that case, the arbitration mechanism has to be “recursive”, i.e. may involve upper levels up to a common host which is able to correctly assess the exchange.

The nature or level of the interacting agents is another aspect of this same question. For now, we assumed negotiation partners at the same level. But, negotiations between agents from different levels are conceivable as well. For instance, in the “table assignment” problem, the guests could be initially contained in the restaurant, which then is in charge of proposing guests to tables; or, the guests contained in the restaurant may be allowed to interact directly with tables. We are currently investigating such inter-level negotiations protocols.

The second key question concerns the capability of agents to assess the variation of their welfare value, induced by a potential exchange (Δw). It is rather difficult to propose an all-purpose solution that would allow a prior evaluation, since it is tightly dependent on the application domain, on the nature of aggregation operators, and on the way the perceived properties are calculated. If prior knowledge is available and if most functions are linear, the variations of welfare may be computed before any effective move.

When prior assessment is not realisable, a general solution consists in a *simulation* of the moves, carried out by the negotiation partners themselves. This process allows each agent to “test” its interest in accepting the exchange, without requiring further communication on its preferences and evaluation functions. The process has of course to be reversible through a “backtrack” mechanisms.

A last point we would like to emphasise, is that a resolution method based on agents behaviours should be able to deal with both static and dynamic problems. Indeed, each agent makes a “local” decision, with regards to its own evaluation of the situation. Thus, from the point of view of the agent, there is no noticeable distinction between a close, static problem where all agents (individuals, resources, groups, etc.) are given once and for all, and an open, dynamic problem where the objectives or properties of the agents, or even the agents themselves, can change over time.

Agent-based resolution methods, in general, are able to operate gradual modifications in response to the reconfiguration of agents, instead of starting the resolution again from scratch. It may prove less efficient of course than a centralised method, but allows to address open systems such as those found in the real world. Hence, in the long term such solving methods can often be transposed to build control mechanisms in distributed systems.

7 Conclusion and perspectives

In this paper, we proposed a uniform approach to model assignment or matching problems through a multilevel multi-agent system. We have shown on examples the capability of this formalism to express a broad diversity of objectives, representing the viewpoints of the actors of the system. Besides, we started an investigation of the resolution principles and algorithms which can be built for this agent-based structure.

Ongoing work now focuses on the elaboration and evaluation of several protocols based on the nature of the hosting graph between agents, so as to avoid being specific to a problem or to a particular family of problems. This requires to identify couplings between the nature and structure of the addressed problems, and the behaviours to give the agents to enable them build a solution in an incremental way, and should also lead to the development of methodological guidelines.

Finally, we would also mention a long-term, but promising, perspective of our work. As we explained, most agent-based multilevel meta-models were initially developed in the field of simulation. One of the key issues in multilevel agent-based simulation consists in techniques enabling the automatic aggregation or disaggregation of agents, either for reducing the computational cost of the behaviours, or to focus the observation of the simulation on a relevant sub-system, or to enhance the realism of the simulation outcome.

Based on the multilevel decomposition of the satisfaction of the agents proposed in our meta-model, metrics can be built at each level in order to reflect the balance between the state of the simulated entities and external criteria: CPU load, interaction with an observer, compliance of the simulation outcome with reference data, etc. In response to low welfare levels, the agents would decide to aggregate or to disaggregate, in order

to enhance their individual or collective satisfaction. We thus think that our multilevel welfare assessment model offers new perspectives to contribute to this issue.

References

1. Airiau, S., Endriss, U.: Multiagent resource allocation with sharable items. *JAAMAS* 28(6), 956–985 (2013)
2. Arrow, K.: *Social Choice and Individual Values*. J Wiley (1963)
3. Arrow, K., Sen, A., Suzumura, K.: *Handbook of Social Choice and Welfare*, vol. 1. Elsevier (2002)
4. Brito, I., Meseguer, P.: Distributed stable matching problems. *LNCS – Principles and Practice of Constraint Programming* 3709, 152–166 (2005)
5. Chevaleyre, Y., Dunne, P., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-Aguilar, J., Sousa, P.: Issues in multiagent resource allocation. *Informatica* 30, 3–31 (2006)
6. Drogoul, A., Amouroux, E., Caillou, P., Gaudou, B., Grignard, A., Marilleau, N., Taillandier, P., Vavasasseur, M., Vo, D.A., Zucker, J.D.: GAMA: multi-level and complex environment for agent-based models and simulations. In: Gini, M., et al. (eds.) *Proc. Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS)*. pp. 1361–1362 (2013)
7. Everaere, P., Morge, M., Picard, G.: Casanova : un comportement d’agent respectant la privacité pour des mariages stables et équitables. *Revue d’Intelligence Artificielle* 26(5), 471–494 (2012)
8. Gale, D., Shapley, L.: College admissions and the stability of marriage. *American Mathematical Monthly* 69, 9–14 (1962)
9. Kuhn, H.: The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 83–97 (1955)
10. Macarthur, K., Stranders, R., Ramchurn, S., Jennings, N.: A distributed anytime algorithm for dynamic task allocation in multi-agent systems. In: *AAAI Conf. on Artificial Intelligence* (2011)
11. Mathieu, P., Picault, S., Secq, Y.: Design patterns for environments in multi-agent simulations. In: Chen, Q., et al. (eds.) *18th Conf. on Principles and Practice of Multi-Agent Systems (PRIMA)*. LNCS, vol. 9387, pp. 678–686. Springer (2015)
12. Maudet, A., Touya, G., Duchêne, C., Picault, S.: Representation of interactions in a multi-level multi-agent model for cartography constraint solving. In: Demazeau, Y., et al. (eds.) *12th Int. Conf. on Practical Applications of Agents and Multi-Agent Systems (PAAMS)*. pp. 183–194. No. 8473 in LNCS, Springer (2014)
13. Morvan, G., Veremme, A., Dupont, D.: IRM4MLS: the influence reaction model for multi-level simulation. In: *Multi-Agent-Based Simulation XI*, LNCS, vol. 6532, pp. 16–27. Springer (2011)
14. Netzer, A., Meisels, A., Zivan, R.: Distributed envy minimization for resource allocation. *JAAMAS* 30(2), 364–402 (2015)
15. Nongaillard, A., Mathieu, P.: Reallocation problems in agent societies: A local mechanism to maximize social welfare. *J. Artificial Societies and Social Simulation* 14(3) (2011)
16. Picault, S., Mathieu, P.: An interaction-oriented model for multi-scale simulation. In: Walsh, T. (ed.) *22nd Int. Joint Conf. on Artificial Intelligence (IJCAI)*. pp. 332–337 (2011)
17. Siebert, J., Ciarletta, L., Chevrier, V.: Agents & artefacts for multiple models coordination: Objective and decentralized coordination of simulators. In: *Proc. ACM Symposium on Applied Computing*. pp. 2024–2028 (2010)
18. Weerdt, M., Zhang, Y., Klos, T.: Multiagent task allocation in social networks. *JAAMAS* 25(1), 46–86 (2011)