A Multi-Agent Resource Negotiation For Social Welfare

Antoine Nongaillard*†
*Dept. Computer Science and Software Engineering
Concordia University, Montréal, Canada
antoine.nongaillard@lifl.fr

Philippe Mathieu†
†Laboratoire d'Informatique Fondamentale de Lille
Université Lille 1, Villeneuve d'Ascq, FRANCE
philippe.mathieu@lifl.fr

Abstract

This study seeks to provide scalable and distributed algorithms to solve the resource allocation problem within an agent community. We propose an approach that can be applied to any kind of contact network, any range for the utility values, for the most important social welfare notions, avoiding the centralized approach drawbacks. In that purpose, we study various agent behaviors. We show that there exists in each case a simple behavior leading the negotiation process to a socially optimal resource allocation as an emergent phenomenon, or to a socially close solution if the need arises. We give, for each social welfare notion, the agent behavior to implement in order to solve the problem.

1. Introduction

The allocation problem has been initially studied as an optimization problem. The purpose of such centralized approaches is first to determine the outcome maximizing an objective function, and then to allocates the resources consequently. These approaches suit well with the solution of many auction problems. Indeed, entities report their preferences to an auctioneer who then notifies the final resource allocation. Mathematical analyses, models and algorithms are provided to solve various kind of auctions such as in [1]. However, strong assumptions are made leading to several important drawbacks. First, such solving processes require the preferences of all entities. No privacy notion is then possible while nowadays people accept less and less to reveal private information, especially with Internet applications. Centralized approaches also assume that any entity can communicate with all others. Such an assumption is not plausible for many concrete applications. For instance, in a social network, someone only knows a small subset of the whole population. Finally, centralized approaches lack of adaptability. A small variation of initial data leads to a new solving process. It may be critical in dynamic applications like peer-to-peer networks, where peers enter and leave continuously.

Alternative approaches have then been offered, focusing on agents [2]. The resource allocation problem is solved using local negotiations among agents where the initial allocation evolves little by little thanks to deals. [3] describes the classes of deals and establishes theorems on the existence or not of deal sequences leading to an optimal resource allocation. Restricted communication possibilities are not considered. Classes of utility and payment function have been studied in order to design convergent negotiation processes [4]. Some authors also study the impact of negotiation processes on the resources, using different notions of social welfare [5]. They establish convergence results depending on the classes of transactions that are allowed. Negotiation protocols have also been designed [6], where no common knowledge is available for the agents. However, the communication possibilities of the agents are always assumed complete, which restrict a lot the field of applicability once more. Compensatory payments are usually allowed during negotiations. However, even if the use of money is constrained (no money creation during a deal), there is often no limit on agent budgets, which is not a plausible assumption. Questions related to compensatory payments are beyond the scope of our study and thus are not considered in the sequel.

In this work, our objective is not only to determine a socially optimal resource allocation, but also to emphasize the way to achieve it. We introduce the notion of contact network that can correspond to any kind of connected graph, ranging from complete networks, over structured ones, to free-scale ones [7]. We seek to define the simplest and most efficient agent behavior, according to Occam's razor principle, in order to favor scalability and dynamicity. In our work, only bilateral transactions without compensatory payments are considered, but the main social welfare notions are considered. For each case, we compare our results with a centralized approach.

Section 2 presents the distributed framework and the features of negotiation processes. Section 3 details the experimental protocol. Finally, Sections 4, 5 and 6 present approaches and results respectively for the utilitarian welfare, the egalitarian welfare and the Nash product.

2. Distributed negotiations

2.1. Definitions and notations

Let $\mathcal{P} = \{a_1, \ldots, a_n\}$ be a population of agents and $\mathcal{R} = \{r_1, \ldots, r_m\}$ a set of resources. A resource allocation is a partition of all resources in $\mathcal{R}$ among agents of $\mathcal{P}$. Each agent a has a set of neighbors $\mathcal{N}_a$, and owns a bundle of m_a resources denoted by $\mathcal{R}_a$. The preferences of the agents are described by means of an additive and normalized utility function [8]: $u_a : \mathcal{R}^{m_a} \to \mathbb{R}$. The set of deal kinds that are allowed during negotiation processes is denoted by $\mathcal{T}$.

2.2. Social welfare

The social welfare theory is used to evaluate the multi-agent system, considering the welfare of each agent [9]. In this study, the three main notions of the social welfare theory are considered.

Definition 1 (Social welfare notion)**.**

- The utilitarian welfare considers the welfare of the whole society: $sw_u(A) = \sum_{a \in \mathcal{P}} u_a(\mathcal{R}_a)$;
- The egalitarian welfare focuses on the fairness among agents: $sw_e(A) = \min_{a \in \mathcal{P}} u_a(\mathcal{R}_a)$;
- The Nash product is a compromise between utilitarian and egalitarian welfares: $sw_N(A) = \prod_{a \in \mathcal{P}} u_a(\mathcal{R}_a)$.

Of course the final resource distribution depends on the used notion. The utilitarian welfare may lead to allocations in which agents have no resource. The egalitarian welfare has not this drawback, but an agent may obtain most of resources if it has "low preferences". The resulting allocation may be very unbalanced. The Nash product maximizes the global utility and decreases inequalities. It leads to more balanced allocations by avoiding the draining phenomenon, and ensure that no agent is neglected.

2.3. Interactions between agent

The problematic is the definition of an agent behavior that ensures the end of negotiation process, with a decision making based on local information only. However, negotiations can be managed in many different ways. If the participant a' rejects the offer of the agent a, then: (i) the initiator ends the negotiation, (ii) it involves another neighbor, or (iii) it offers another resource.

Various behaviors have been designed, implemented and evaluated, but only the best behavior (described in Figure 1) is presented in the sequel. This agent behavior is flexible and volatile.

2.4. Transaction kinds

As we said before, only bilateral transactions are considered, i.e., transactions involving simultaneously two agents.

```
Sorts its resource bundle R_a
for all r ∈ R_a do
    for all a' ∈ N_a do
        for all δ ∈ T do
            if TEST then
                PERFORM δ
                Ends the negotiation
            end if
        end for
    end for
end for
```

Figure 1. Flexible and volatile behavior of an agent initiator a

Definition 2 (Bilateral transaction)**.** A bilateral transaction between a and a' is a couple $\delta_a^{a'}\langle u, v\rangle = (\rho_a, \rho'_a)$ where a and a' provide respectively a set ρ_a of u resources and ρ'_a of v resources.

Any kind of bilateral deal can be modeled thanks to the cardinality parameters. For instance, a gift from a to a' is written: $\delta_a^{a'}\langle 1, 0\rangle = (\{r\}, \emptyset)$. Of course, the number of possible deals grows exponentially with u and v. Thus, we restrict it. Our experiments show that increasing the size of u and v does not improve significantly the obtained solution.

In the experiments of Section 3, four negotiation policies are compared: the gift $\langle 1, 0\rangle$, the swap $\langle 1, 1\rangle$, the swap+gift where $u, v \leq 1$ and finally the cluster-swap with $u \leq m_a$ and $v \leq m_{a'}$.

2.5. The sociability criterion

In order to ensure that a negotiation process is finite, an acceptability criterion for a deal must be used. In this study, the individually rational criterion [3] has been evaluated, but we have shown that it always lead negotiations to socially weaker allocations. The lone acceptability criterion used and presented in the sequel is the social criterion. A transaction is social when the associated social value is increased by this transaction:

Definition 3 (Social transaction)**.** A deal, δ, which changes the initial allocation A into a new one A', is social when: $sw(A') > sw(A)$.

The important point, is that in practice, it is not essential to determine the welfare value that can only be done by means of a central entity. The evolution of the welfare value is sufficient, and this can be determined on a local basis.

3. Experimental setup

The evaluation of negotiation processes is not an obvious issue. In this study, the optimal solution is determined using a centralized approach, and then used as a reference in order

Table 1. Utilitarian efficiency for different kinds of contact network with $n = 100$ and $m = 500$.

Contact network kind	Gift	Swap	S+G	CS
Full	100	97.6	100	100
Erdős-Rényi	98.8	95.9	99.1	98.9
Grid	71.7	68.5	76.3	66.1
Small world WS	96.1	92.7	97.2	96.1
Small world PA	88.8	87.8	91.0	89.2

Table 2. Egalitarian efficiency for different kinds of contact network with $n = 100$ and $m = 500$.

Contact network kind	Gift	Swap	S+G	CS
Full	66.6	15.8	99.9	99.9
Erdős-Rényi	66.1	14.1	86.2	89.9
Grid	61.0	12.9	80.9	81.0
Small world WS	65.3	18.7	84.4	86.1
Small world PA	49.1	13.9	55.6	54.4

to evaluate the efficiency of negotiation processes. Such an approach does not consider the deal sequence that is required to achieve an optimal allocation, and implicitly assumes that the contact network is complete.

During the experiments, various contact networks have been generated, from complete networks, over grids, to Erdös-Rényi graphs. Two kinds of small worlds are also considered: the ones generated using the preferential attachment, and others generated by circle rewiring [7].

The initial resource allocation is randomly generated, as well as the preferences of the agents. The speech turn is uniformly distributed. For the experiments of this paper, we have fixed the population size to 100, and the number of available resources to 500. We have then generated 20 graphs of each kind, 10 agent preferences sets for each graph, and 100 different initial allocations for each set of agent preferences.

4. Utilitarian Welfare

When the utilitarian welfare is considered, an optimal allocation can be generated quite easily by allocating each resource to the agent that estimates it the most.

4.1. Utilitarian agent's approach

When the utilitarian welfare is considered, the acceptability condition of a deal can be expressed as follows. In Figure 1, a gift transaction $\delta_a^{a'}$ is acceptable if $\{\text{TEST}\}$= $u_{a'}(r) > u_a(r)$, while a swap must satisfy $\{\text{TEST}\}$= $u_a(r') + u_{a'}(r) > u_a(r) + u_{a'}(r')$.

Table 1 describes the utilitarian efficiency obtained according to the different kinds of contact network.

Negotiation processes based on a gift are always sufficient (100%) for a complete graph. Even if some other policies obtain similar results, their use lead to useless additional costs.

Distributed negotiations can handle restricted contact networks, but it is not possible to guarantee that an optimal allocation can be achieved, because it depends strongly on the network topology.

Experiments shows that less the mean connectivity is, less efficient are negotiations. On Erdős-Rényi graphs, negotiations achieve 99% of the optimum when only gifts are used while on grids, the use of gifts and swaps is required to achieve 76% for the best case.

Utilitarian solution: When the mean connectivity is high, social gifts are sufficient to achieve efficient allocations. In figure 1, $\mathcal{T} = \{gift\}$. However, when the mean connectivity decreases a lot, the combination with swaps are required to achieve efficient allocations, $\mathcal{T} = \{gift, swap\}$. Both of them are very scalable and can be applied to large instances.

5. Egalitarian Welfare

The determination of an egalitarian optimum can be formulated using an equationnal model, which can be solved using any mathematical programming optimizer. However, such a method is not scalable for large instances. It cannot exhibit an acceptable transaction sequence, and assume that the contact network is complete.

5.1. Egalitarian agent's approach

When the egalitarian welfare is considered, an agent cannot accept a transaction $\delta_a^{a'}$ if it becomes poorer than the poorest agent before the transaction: in Figure 1, $\{\text{TEST}\}$=$\min[u_a(\mathcal{R}'_a), u_{a'}(\mathcal{R}'_{a'})] > \min[u_a(\mathcal{R}_a), u_{a'}(\mathcal{R}_{a'})]$, where $\mathcal{R}'_a, \mathcal{R}'_{a'}$ are the resource bundles after the deal $\delta_a^{a'}$.

Table 2 describes the social efficiency that is obtained according to the different kinds of contact network.

Even on complete graphs, it is not possible to guarantee the achievement of an optimal allocation. Policy based either on gifts or on swaps are not efficient at all. Both are essential to achieve socially efficient allocations. The connectivity has a more important impact on egalitarian negotiations. A weak shortest path between agents helps the resource traffic, but some agents with a small neighborhood can penalize the whole system. The standard deviation increases in opposition of the mean connectivity: the initial allocation is then important.

Egalitarian solution: Hence, the social swap+gift policy is the best alternative and leads to socially more interesting alternatives in a scalable way. Then, in Figure, 1 $\mathcal{T} = \{gift, swap\}$. While maximizing the utilitarian welfare can be done on networks with a weak connectivity, maximizing the egalitarian welfare requires a larger mean connectivity to be efficient.

Table 3. Nash efficiency for different kinds of contact network with $n = 100$ and $m = 500$.

Contact network kind	Gift	Swap	S+G	CS
Full	80.2	169.2	394.1	277.2
Erdős-Rényi	49.6	79.9	270.0	206.5
Grid	50.9	6.4	353.2	106.7
Small world WS	178.8	22.8	1149.1	673.6
Small world PA	77.4	2.11	1119.2	337.5

6. Nash Product Welfare

When the Nash product is considered, an optimal allocation cannot be determined using a classic way since the problem is not linear. Optimization methods exists to overcome such a problem, but only a time-consuming estimation can only be obtained. However, we choose to develop a heuristic in that purpose. It focuses on the resource value, by first allocating each resource to an agent that estimates it the most. Then, the algorithm checks if every agent has at least one resource, and if the need arises, it looks for picking up the resource maximizing the social value to an agent that has at least two resources.

6.1. Nash agent's approach

When the Nash product is considered, a transaction $\delta_a^{a'}$ is acceptable if, in Figure 1, $\{\text{TEST}\} = u_a(\mathcal{R}'_a) \times u_{a'}(\mathcal{R}'_{a'}) > u_a(\mathcal{R}_a) \times u_{a'}(\mathcal{R}'_{a'})$.

Table 3 describes the social efficiency that is obtained according to the different kinds of contact network. A social efficiency greater than 100% is possible since the value used as reference is provided by a heuristic. In such case, the distributed negotiation process lead to greater results than the heuristic.

The social swap+gift strategy leads to greatest results, associated with the lowest standard deviation, on all the kinds of considered contact networks. This strategy always leads to better results than the heuristic. The social gift policy and the social swap one are not efficient, since they lead to allocations that are far. The social swap strategy is especially sensitive to the mean connectivity: on grids or on preferential attachment small worlds, processes lead to allocations representing around 5% of the value given by the heuristic. The strategy based on social cluster swap transactions leads to better results than the heuristic but remains too expensive to be considered as interesting.

Some value may appear really large (e.g. 1119% of the value given by the heuristic). However, since the objective function is a product, a simple exchange of two resources can increase the objective value by a factor more than 100. The standard deviation is also sensitive to the mean connectivity, like in the case of the egalitarian welfare but, in the case of the swap+gift strategy, it never exceeds 5%.

Nash product solution: Thus, the social swap+gift policy is enough flexible to lead to socially efficient allocations, in a scalable time, independently of the contact network kind that is considered. Then, in Figure 1, $\mathcal{T} = \{gift, swap\}$.

7. Conclusion

This paper offers a scalable method to solve efficiently the resource allocation problem in a decentralized way. Based on local negotiations among agents, this approach considers the relationships among them thanks to the contact network notion which can be built on any connected graph. This method is an adaptive process where the addition of new agents is possible during the negotiation process. It is also an anytime approach: the quality of the solution increases gradually and the negotiation process can be interrupted anytime. At the opposite of centralized approaches, our distributed algorithms can exhibit a transaction sequence leading to the provided allocation, with respect to the contact network topology. Such a facet of the problem cannot be handle by centralized approaches in a scalable way. Although we can only guarantee an optimal solution on a complete contact network when the utilitarian welfare is considered, we are able to provide efficient and practical algorithms leading to allocations in all other cases, for many kinds of contact network.

References

[1] T. Sandholm, "Algorithm for optimal winner determination in combinatorial auctions," *Artificial Intelligence*, pp. 1–54, 2002.

[2] J. Ferber, *Multi-Agent System: An Introduction to Distributed Artificial Intelligence*. Addison Wesley Longman, 1999.

[3] T. Sandholm, "Contract Types for Satisficing Task Allocation: I Theoretical Results," in *AAAI Spring Symposium: Satisficing Models*, 1998, pp. 68–75.

[4] Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet, "Negotiating over Small Bundles of Resources," in *AAMAS'05*, 2005, pp. 296–302.

[5] U. Endriss, N. Maudet, F. Sadri, and F. Toni, "Negotiating Socially Optimal Allocations of Resources," *Journal of Artificial Intelligence Research*, pp. 315–348, 2006.

[6] S. Saha and S. Sen, "An Efficient Protocol for Negotiation over Multiple Indivisible Resources," in *IJCAI'07*, 2007, pp. 1494–1499.

[7] R. Albert and A. Barabási, "Statistical Mechanics of Complex Networks," *Reviews of Modern Physics*, pp. 47–97, 2002.

[8] P. Miranda, M. Grabisch, and P. Gil, "Axiomatic Structure of k-Additive Capacities," *Mathematical Social Sciences*, pp. 153–178, 2005.

[9] K. Arrow, A. Sen, and K. Suzumura, *Handbook of Social Choice and Welfare*. Elsevier, 2002.