A Realistic Approach to Solve the Nash Welfare

A. Nongaillard*†, P. Mathieu† and B. Jaumard‡

Abstract The multi-agent resource allocation problem is the negotiation of a set of resources among a population of agents, in order to maximize a social welfare function. The purpose of this study is the definition of the agent behavior which leads, if possible, to an optimal resource allocation at the end of the negotiation process as an emergent phenomenon. This process can be based on any kind of contact networks. Our study focuses on a specific notion: the Nash product, which has not the drawbacks of the other widely used notions. However, centralized approaches cannot handle large instances, since the social function is not linear. After a study of different bilateral transaction types, we underline the most efficient negotiation policy in order to solve the multi-agent resource allocation problem with the Nash product and provide an adaptive, scalable and anytime algorithm.

1 Introduction

The multi-agent resource allocation problem has been studied for a long time, either within a centralized framework or a distributed one. In studies with a centralized approach, agents report their preferences on the resources to a specific agent, e.g., an auctioneer, who then determines the final resource allocation. Within this context, authors [17] have suggested different transaction models for given types of auctions. At the opposite, in studies based on a distributed framework, an initial resource allocation evolves by means of local negotiations among agents. The convergence of such a negotiation process can be viewed as an emergent phenomenon, due to local negotiations among agents. The advantages are adaptability and dynamicity of the system, while keeping privacy for all users. However, assumptions have been implicitly made. Indeed, an agent is able to communicate and negotiate with all other

† LIFL, University of Lille, Villeneuve d'Ascq, France e-mail: forename.name@lifl.fr
* CSE, ‡ CIISE, Concordia University, Montreal, Canada e-mail: bjaumard@ciise.concordia.ca

agents. This assumption is not always plausible as soon as real world applications are considered.

The evaluation of a resource allocation is usually made by means of notions from the social welfare theory [2]. The most widely used notions such as the utilitarian welfare or the egalitarian welfare may have some undesirable effects on the resource allocation which is obtained. In this work, we choose to focus on the Nash product, which has not these drawbacks (see Section 2.3). Considering this welfare, an approach based on a centralized framework is not efficient. The determination of the optimal resource allocation is a complex and time-consuming problem which can only be solved in a very large amout of time.

In this study, our purpose is the definition of the best agent behavior, in order to ensure the convergence of the negotiation process towards a socially optimal resource allocation, or when the need arises, towards a socially close allocation. We provide a scalable and anytime algorithm which can be based on any kind of contact network, when the Nash product is considered as social welfare measure. The solution which is provided in this paper is adaptive, new agents can be added with new resources during the negotiation process, while with a centralized approach, such a thing is not possible without restarting the whole solving process with the new data.

After a presentation of related studies in Section 1.1, we define basic notions in Section 2. Section 3 presents a centralized approach, and Section 4 presents our distributed approach. Section 5 presents the experiment protocol, the evaluation criteria, and an analysis of the results.

1.1 Related works

Lots of studies focus on mathematical properties of the multi-agent resource allocation problem. In [16], the author handles the properties of the allowed transactions and establishes a classification of the basic transactions along with theorems on the existence or the non-existence of a transaction sequence leading from any initial resource allocation to an optimal one. However, no process are provided to reach an optimal resource allocation. Along the same lines, mathematical properties on some classes of utility functions and payment functions are studied in [7] in order to design negotiation processes, which terminate after a finite number of iterations. [8] presents the impact of the acceptability criterion, the utility function and the transaction properties on the society welfare, without regard for the agent behavior which leads the negotiation process to such a socially optimal resource allocation. In other studies, authors define criterion which favors equitable deals [9] and others study the envy-freeness in the resource allocation process [4, 6]. In [14, 15], agent behaviors are studied, but only in the case of the utilitarian welfare, for which an obvious centralized solution exists. The notion of neighborhood is seldom considered, whereas it is one of the most important points for real world applications.

2 Multi-Agent Resource Allocation Problem

2.1 Definitions and notations

The multi-agent resource allocation problem is based on a population $\mathscr{P} = \{a_1, \ldots, a_n\}$ of agents, and on a set $\mathscr{R} = \{r_1, \ldots, r_m\}$ of available resources, which are assumed indivisible and static.

This set of resources $\mathscr{R}$ is initially distributed over the population of agents $\mathscr{P}$. Each agent a owns a bundle of resources, $\mathscr{R}_a$. A resource allocation A is a partitioning of the resources in $\mathscr{R}$ among the agents of $\mathscr{P}$, $A = \{\mathscr{R}_1, \ldots, \mathscr{R}_n\}$. $\mathscr{A}$ is the set of all the possible allocations. The preferences of the agents are expressed by means of 1-additive utility function [5, 12]. $u_a : \mathscr{R}^{m_a} \rightarrow \mathbb{R}$ and $u'_a : \mathscr{R} \rightarrow \mathbb{R}$ with the following relationship: $u_a(\mathscr{R}_a) = \sum\limits_{r \in \mathscr{R}_a} u'_a(r)$. Even if their mathematical definitions are different, since they are used in the same purpose, u_a will be used equally in order to simplify the notations.

2.2 Contact network

The relationships among the agents can be represented by means of a graph: the contact network. A link between two agents means that they are able to communicate between them. Most of the studies rely on the hypothesis of a complete contact network. Any agent is able to negotiate with all other agent in the population. Such a hypothesis has a strong impact on the negotiation process, and it is not realistic as soon as real world applications are considered. For instance, in the case of social networks, a person only knows a subset of the overall set of actors in the network. Thus, the neighborhood of an agent a, denoted by $\mathscr{N}_a$, define the set of agents with who he is able to talk. In this study, we consider that the contact network can be any connected graph, ranging from complete graphs to small-world graph [1], including structured graphs like rings, trees, or grids.

2.3 Social welfare

In order to evaluate a resource allocation, notions which come from the social welfare theory are considered [2, 13]. Several notions exist, and each has advantages and drawbacks.

Definition 1. The **utilitarian welfare** of a resource allocation A, denoted by $sw_u(A)$, corresponds to the summation of the agent utilities: $sw_u(A) = \sum\limits_{a \in \mathscr{P}} u_a(\mathscr{R}_a)$.

Definition 2. The **egalitarian welfare** of a resource allocation A, denoted by $sw_e(A)$, corresponds to the utility of the poorest agent: $sw_e(A) = \min\limits_{a \in \mathscr{P}} u_a(\mathscr{R}_a)$.

Definition 3. The **Nash product** of a resource allocation A, denoted by $sw_N(A)$, corresponds to the product of the agent utilities: $sw_N(A) = \prod_{a \in \mathscr{P}} u_a(\mathscr{R}_a)$.

In order to illustrate the difference among these notions, let us consider a population of 3 agents, $\mathscr{P} = \{a_1, a_2, a_3\}$, and a set of 6 available resources, $\mathscr{R} = \{r_1, r_2, r_3, r_4, r_5, r_6\}$. Their preferences are expressed by means of a utility function, as described in Table 1. Optimal social values are gathered in Table 2 with a corresponding resource allocation. The utilitarian welfare considers the welfare of the

Table 1 Agent preferences

Agents	Resources					
	r_1	r_2	r_3	r_4	r_5	r_6
a_1	10	7	10	9	2	1
a_2	6	10	3	4	8	6
a_3	1	2	1	2	1	3

Table 2 Optimal social values

Social welfare	Value	Resource allocation
sw_u	53	$[\{r_1, r_3, r_4\}\{r_2, r_5, r_6\}\{\}]$
sw_e	6	$[\{r_1\}\{r_5\}\{r_2, r_3, r_4, r_6\}]$
sw_N	1800	$[\{r_1, r_3\}\{r_2, r_5\}\{r_4, r_6\}]$

whole agent community, without concern about the individual welfare, and then can lead to resource allocations where one agent, a_3, does not get any resource. Some agents can be neglected, especially if, for each resource, there exists another agent who estimates more this resource. At the opposite, the egalitarian welfare considers only the individual welfare, and then leads to resource allocations where every agent owns at least one resource. No agent is neglected, but an agent with low preferences, like a_3, drains the resources, and the resulting allocation may be very unbalanced. In between, the Nash product is a compromise which leads to more balanced resource allocations, avoiding such a draining phenomenon, and where no agent is neglected. This notion can only be used when utility values are positive. Moreover, small variations in an allocation lead to very large variations of the welfare: for instance, a simple exchange of r_1 with r_6 leads to a decrease of the social value from 1800 to 594.

3 A centralized approach

Of course, this resource allocation process can be solved using a centralized framework. The optimal social value, and a corresponding resource allocation, can be determined by means of the following mathematical model. The boolean variables x_{ra} represent the ownership of a resource r by an agent a, with $r \in \mathscr{R}, a \in \mathscr{P}$. Then, the optimal value of the Nash product can be found by solving this equation system:

$$sw_N^\star = \begin{cases} \max \prod_{a \in \mathscr{P}} \sum_{r \in \mathscr{R}} u_a(r) x_{ra} \\ \text{subject to: } \sum_{a \in \mathscr{P}} x_{ra} = 1 \qquad r \in \mathscr{R} \\ \qquad\qquad x_{ra} \in \{0,1\} \qquad r \in \mathscr{R}, a \in \mathscr{P}. \end{cases}$$

Such a system cannot be handled in a classic way since the objective function is not linear. However, an estimation could be made with the following method. First, a *Lagrangian* relaxation is used [10]. This method can solve non linear equation system if the objective function is convex. However, it is not the case when the Nash product is considered, and a multi-start algorithm has to be combined with the relaxation using a sampling small enough for initial solutions. Starting from multiple initial solutions may avoid local optima when non convex function are considered [11].

Moreover, since resources are not divisible, an integer solution has still to be found. Indeed, the relaxation changes the discrete value set $\{0,1\}$ into a continuous value set $[0,1]$. In order to obtain an integer solution, a *branch-and-bound* algorithm is then used, guided by the values that have been provided by the relaxed solution.

Such a method cannot certify the optimality of the found solution. This centralized approach is not really scalable, consequently of the nonlinearity of the objective, and of the exponential solution space. A resource allocation problem with n agents and m resources leads to a solution space of size n^m ($n \ll m$). Moreover, the implicit assumption of a complete contact network has been made with such a model. Large sets of constraints have to be added to this model in order to prohibit exchanges between agents who are not related.

Since such a method is not scalable, we developed two scalable heuristics in order to determine an estimation of the optimal social value. The first heuristic is focused on the resource value. The first step of this algorithm is to allocate each resource to the agent who estimates it the most. The second step is to be sure that all agents own at least one resource, otherwise, it looks for picking up the resource maximizing the social value to an agent who has at least two resources. The second heuristic is focused on the resource distribution uniformity, by allocating successively the best remaining resource to each agent.

4 A Distributed approach

Our proposition is then to use a distributed approach. The purpose of such an approach is to define the agent behavior which leads to a socially optimal resource allocation, as an emergent phenomenon at the end of the negotiation process. At the opposite of centralized approach, our distributed approach can be based on any kind of communication network, as discussed in Section 2.2. The question is: “which behavior must we give to the agent in order to obtain a good Nash welfare as emergent phenomenon ?”

4.1 Acceptability Criteria

Such criteria have a strong impact on negotiation processes. Indeed, if an agent can accept any kind of deal, then the negotiation process will not be able to stop.

Even if the resource allocation process reaches an optimal state, the agents will continue to negotiate among them and leave the optimum. Moreover, there is not guarantee that such a process reaches one time a suitable resource allocation. The acceptability criteria help the agent to determine whether a transaction is profitable or not. An agent has to based his decision on an acceptability criteria, with respect to the agent behavior. Such criteria restrict a lot the set of possible transactions among the agents. A negotiation process ends when no agent in the population is able to find an acceptable deal.

Let two agents, a and a', illustrate the considered criteria. The agent a initiates a transaction $\delta(A,A')$ with an agent a': the initial resource allocation $A = [\ldots,\mathscr{R}_i,\ldots,\mathscr{R}_j,\ldots]$ evolves towards a new one A'.

Definition 4. A **rational agent** is an agent who only accepts transactions that increase his utility. If the agent a is rational, he accepts a transaction only if: $u_a(\mathscr{R}'_a) > u_a(\mathscr{R}_a)$.

The rationality criterion is the most widely used in the literature, especially in the case of non cooperative and selfish agents.

Definition 5. A **rational transaction** is a transaction in which all involved agents are rational. If a transaction is rational, involved agents accept it if: $u_a(\mathscr{R}'_a) > u_a(\mathscr{R}_a) \quad \text{and} \quad u_{a'}(\mathscr{R}'_{a'}) > u_{a'}(\mathscr{R}_{a'})$.

However, this criterion restricts a lot the set of possible transactions, and may lead the negotiation process to a sub-optimal resource allocation.

Another criterion that ensures the end of the negotiation process after a finite number of transactions is the sociality. This criterion is based on a local evaluation of the social welfare evolution.

Definition 6. A **social agent** is an agent who can only accept transactions that increase the considered social welfare function of the multi-agent system.

Definition 7. A **social transaction** is a transaction which causes an increase of the considered social welfare function. Such a transaction can only be accepted by the involved agents if: $sw(A') > sw(A), \quad A,A' \in \mathscr{A} \quad \text{such that } A \xrightarrow{\delta} A'$.

The determination of the social value associated to a resource allocation needs global information: Indeed, it is essential to have the value of the utility of each agent. However, it is possible to determine the variation of this social value with local information. From the agent point of view, the acceptability can be determine from information given by involved agents. It is then not necessary to determine its value.

$$sw_N(A') > sw_N(A) \Longleftrightarrow u_a(\mathscr{R}'_a) * u_{a'}(\mathscr{R}'_{a'}) > u_a(\mathscr{R}_a) * u_{a'}(\mathscr{R}_{a'}).$$

where $\mathscr{R}_a$ and $\mathscr{R}'_a$ is the bundle of a before and after the deal. Since a finite number of agents are involved in a transaction, two in the case of bilateral transactions, only their resource bundle change. Then the utility of the agents that are not involved in this transaction can be considered as a constant value.

4.2 Transaction Kinds

Our study is restricted to bilateral transactions, i.e., transactions involving simultaneously two agents. Indeed, multilateral transactions are too much complex and time-consuming to be optimally determined, especially when the Nash product is considered. Moreover, our aim is to define the simplest agent behavior in order to favor the scalability of the algorithm. Three kinds of bilateral transactions can be distinguished. Others are combinations of these basic transactions. In each case, the transaction is initiated by a, in which is involved one of his neighbor a'. They own respectively m_a and $m_{a'}$ resources in their bundle

First, the *gift* transaction. The initiator a can only give one resource to a'. Only m_a gifts are possible. The gift transaction cannot be rational for the initiator and is always rational for the agent participant (since utilities are positive).

Then, the *swap* transaction. Each agent provides a unique resource. This deal is symmetric: the number of resources per agent cannot vary. Hence, an optimal solution can be reached only if the initial allocation has the same resource distribution as one of the optimal allocation. $m_a \times m_{a'}$ swaps are possible.

Finally, the *cluster-swap* (CS). Each agent can involve a subset of their resources. At the opposite of the swap, it can be asymmetric. The cluster-swap contains the gift and the swap transactions. $2^{m_a+m_{a'}}$ are possible

In the experiments of Section 5, besides "pure" negotiation policies which use only one transaction kind, a "mixed" policy is defined: the swap+gift policy (S+G) in which the initiator tries first to find an acceptable swap, and a gift if the need arises. Agents use these policies according a specific behavior, which is defined in the next Section.

4.3 Agent Behavior

A negotiation can be managed in many different ways. Indeed, during a negotiation, if the participant rejects the offer, three alternatives arise: (i) the initiator gives up and stops the negotiation, (ii) he selects another neighbor, or (iii) he changes the offered resource. Based on these, various behaviors have been designed, implemented and evaluated. Nevertheless, we always assumed that each agent tries to give first his resource associated with the lowest utility. Only the behavior which leads to the best results is presented in the sequel.

This agent behavior is flexible and volatile, which means that the initiator can change either the selected neighbor or the offered resource. Such a behavior is "complete", meaning that according to kind of allowed transactions, if an acceptable transaction exists in the neighborhood, it will be identified. This completeness, which leads to greater results, has a cost. Costless behaviors can be designed, according to the application and its quality requirements.

Algorithm 1: Behavior of the initiator a

```
Sorts his resource bundle ℛ_a ;
forall r ∈ ℛ_a do
    forall a′ ∈ 𝒩_a do
        if δ_a is acceptable then
            Performs the transaction δ_a ;
            Ends the negotiation ;
        end
    end
end
```

5 Experiments

5.1 Experimental setup and evaluation

During the experiments, various contact networks have been generated, some complete and some Erdos-Renyi networks [3]. The mean connectivity degree of such networks is $\frac{n}{4}$, which means that an agent can talk at most to 25% of the population. The resources are initially distributed randomly. The preferences are also generated randomly with values in 1..100. During the negotiation processes, the speech turn is uniformly distributed over the agents. For the different population's sizes, the different allocation kinds, and the different networks, 100 instances are run each time.

The evaluation of negotiation processes is not an obvious issue. It is always possible to find a metric which makes a process the best. Various metrics can be considered like the number of performed deals, the number of exchanged resources, the number of speech turns or the number of attempted transactions. The relative standard deviation among the social value of emergent allocations are also considered: a large value means that the considered negotiation process is very sensitive to the initial resource allocation, and thus the quality of the emergent allocation quite varies. Finally, a comparison with the estimation that is obtained with centralized heuristics is made over 100000 instances. A ratio is computed to determine the gap between them: the distributed social value over the centralized social value.

5.2 Result Analysis

First, results related to complete networks. The social swap+gift policy is more robust than others, with a relative standard deviation of 3.03% among the social values for instances with 50 agents and 300 resources. This policy is less sensitive to the initial allocation. At the opposite, the swap policy is not reliable, because strongly sensitive to the initial allocation, with a social deviation of 114.09%. The swap policy is not enough flexible to avoid a local optimum. The swap+gift policy leads to socially greater allocation than than ones reached by other policies, but is more expensive in terms of attempted transactions(until 6 times). The cluster-swap policy is the most expensive in time and in attempted transactions, the social deviation is large enough to not compensate the additional costs, and thus is not interesting. Computation times, shown in Table 3, and the corresponding number of performed deals, shown in Table 4, are obtained for the swap+gift policy on complete contact networks. Even an instance involving 50 agents and 2500 resources remains scalable in a reasonable time.

Table 3 Computation time

Mean number of	Number of agents		
resources per agent	5	25	50
5	50ms	250ms	600ms
25	150ms	45s	4min
50	6s	5min	25min

Table 4 Number of performed deals

Mean number of	Number of agents		
resources per agent	5	25	50
5	35	350	900
25	150	1800	4600
50	400	4000	9000

Experiments on Erdos-Renyi networks bring about similar conclusions. The social swap+gift policy still obtains best results, with the greatest social value, coupled with the lowest deviation ($\simeq$17%). The value of this deviation depends strongly of the network topology. Greater is the mean connectivity, lower will be the relative standard deviation among the social values.

Finally, complex transactions such as the cluster-swaps, lead to a higher number of attempted transactions, a larger computation time, but a lower number of speech turns. At the opposite, simpler transactions such as gifts, lead to short negotiations in time, but more speech turn are required before the end of the processes. However, when the Nash product is considered, the gift policy or the swap policy are not enough flexible to leave local optima and then leads to weaker allocations. The swap+gift policy is a good compromise between scalability and complexity.

Since the swap+gift policy leads to the greatest social value, it has been compared to the value obtained by the centralized heuristics, defined in Section 3. First, the comparison to the heuristic which focused on the resource value. The distributed negotiation process leads to better results on 99.97% of the instances, with a social improvement of 140.86%, whereas when it leads to worst results, the gap is only of 1.13%. During our experiments, the second heuristic, which is focused on the resource distribution, never reaches a better allocation than our distributed negotiation process. The gap between the social value are huge, more than 10000%.

Thus, the social swap+gift policy is a flexible policy which leads to socially efficient allocations as well on complete networks as on Erdos-Renyi networks. Negotiation processes end in scalable time, for a small additional costs in terms of attempted transactions.

6 Conclusion

A centralized approach makes the solution of the resource allocation problem very complex and time-consuming, as soon as the Nash product is considered as the social welfare function. In this study, we have designed a negotiation process among agents which leads to the emergence of a suitable resource allocation, by means of local negotiations. This solving method is scalable, robust in terms of solution quality, and adaptable. At the opposite of a centralized approach, ours can be based on any type of contact network, and the addition of new agents (or new resources) is possible and does not need to restart of the negotiation process. Moreover, it is also an "anytime" algorithm: The quality of the solution increases as transactions go along, and the solving process can be interrupted anytime.

References

1. R. Albert and A. Barabási. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74(1):47–97, 2002.
2. K. Arrow, A. Sen, and K. Suzumura. *Handbook of Social Choice and Welfare*. Elsevier, 2002.
3. B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.
4. S. Bouveret and J. Lang. Efficiency and envy-freeness in fair division of indivisible goods: logical representation and complexity. In *IJCAI'05*, pages 935–940, 2005.
5. Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation with k-additive utility functions. In *Proc. DIMACS-LAMSADE Workshop on Computer Science and Decision Theory, Annales du LAMSADE*, volume 3, pages 83–100, 2004.
6. Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Reaching Envy-free States in Distributed Negotiation Settings. In *IJCAI'07*, pages 1239–1244, 2007.
7. Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. Negotiating over Small Bundles of Resources. In *AAMAS'05*, pages 296–302, 2005.
8. U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.
9. S. Estivie, Y. Chevaleyre, U. Endriss, and N. Maudet. How Equitable is Rational Negotiation? In *AAMAS'06*, pages 866–873, 2006.
10. M. Fisher. The Lagrangian Relaxation Method for Solving Integer Programming Problems. *Management Science*, 50(13):1861–1871, 2004.
11. F. Hickernell and Y. Yuan. A Simple Multistart Algorithm for Global Optimization. *OR Transactions*, 1(2), 1997.
12. P. Miranda, M. Grabisch, and P. Gil. Axiomatic structure of k-additive capacities. *Mathematical Social Sciences*, 49:153–178, 2005.
13. H. Moulin. Choosing from a Tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
14. A. Nongaillard, P. Mathieu, and B. Jaumard. A Multi-Agent Resource Negotiation for the Utilitarian Welfare. In *ESAW'08*, 2008.
15. A. Nongaillard, P. Mathieu, and B. Jaumard. La négociation du bien-être social utilitaire. In *JFSMA'08*, pages 55–64, EU, France, Brest, 2008.
16. T. Sandholm. Contract Types for Satisficing Task Allocation: I Theoretical Results. In *AAAI Spring Symposium: Satisficing Models*, volume 99, pages 68–75, 1998.
17. T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.