Multiagent Resource Allocation with k-additive Utility Functions

Yann Chevaleyre[1], Ulle Endriss[2], Sylvia Estivie[1], and Nicolas Maudet[1]

[1]LAMSADE, Université Paris-Dauphine
Email: {chevaley,maudet}@lamsade.dauphine.fr

[2]Department of Computing, Imperial College London
Email: ue@doc.ic.ac.uk

Abstract

In this extended abstract, we briefly review previous work on the *welfare engineering* framework in which autonomous software agents negotiate on the allocation of a number of discrete resources, and we point out several connections to combinatorial optimisation problems that shed light on the computational complexity of the framework. We give particular consideration to scenarios where the preferences of agents are modelled in terms of k-additive utility functions.

1 Introduction

Distributed systems in which autonomous software agents interact with each other, in either cooperative or competitive ways, can often be usefully interpreted as *societies of agents*; and we can employ formal tools from microeconomics to analyse such systems. If we model the interests of individual agents in terms of a notion of *individual welfare*, then the overall performance of the system provides us with a measure of *social welfare*.

Individual welfare may be measured either *quantitatively*, typically by defining a utility function mapping "states of affairs" (outcomes of an election, allocations of resources, agreements on a joint plan of action, etc.) to numeric values; or *qualitatively*, by defining a preference relation over alternative states. The concept of social welfare, as studied in welfare economics, is an attempt to characterise the well-being of a society in relation to the welfare enjoyed by its individual members [1, 2, 14, 19]. The best known examples (both relying on quantitative measures of individual welfare) are the *utilitarian* programme, according to which social welfare should be interpreted as the sum of individual utilities, and the *egalitarian* programme, which identifies the welfare of society with the welfare of its "poorest" member.

For instance, in an electronic commerce application where users pay a fee to the provider of the infrastructure depending on the personal benefits incurred by using the system, the increase in utilitarian social welfare correctly reflects the profit generated by the provider. The application discussed by Lemaître et al. [13], on the other hand, where agents representing different stake-holders repeatedly negotiate over the access to an earth observation satellite (which has been jointly funded by the stake-holders), requires a *fair* treatment of all agents. Here, the respective values of different access schedules may be better modelled by an egalitarian social welfare ordering.

We are particularly interested in applications where negotiation between autonomous agents serves as a means of addressing a resource allocation problem. Recent results in this framework concern the feasibility of reaching an allocation of resources that is optimal from a social point of view [7, 17], as well as (certain aspects of) the complexity of doing so, in terms of both computational costs and the amount of communication required [4, 5, 6]. Other applications include automatic contracting [17], selfish routing in shared networks [9], distributed reinforcement learning [20], and data mining [12]. This area of activity, which we may term *computational microeconomics*, brings together theoretical computer science and microeconomics in new and fruitful ways, benefiting not only these disciplines themselves but also "hot" research topics such as multiagent systems and electronic commerce.

In previous work, we have put forward the framework of *welfare engineering* [7], which addresses the design of suitable rationality criteria for autonomous software agents participating in negotiations over resources in view of different notions of social welfare, as well as the development of such notions of social welfare themselves. In this extended abstract, we briefly review the underlying multiagent resource allocation system and recall two previous results on the feasibility of reaching a socially optimal allocation of resources from a utilitarian point of view. As we shall see, in cases where the utility functions used by agents to model their preferences over alternative bundles of resources are *additive*, it is sufficient to use very simple negotiation protocols that only cater for deals involving a single resource at a time. This result suggests to investigate generalisations of the notion of additivity, and hence we consider the case of *k-additive* functions, as studied, for instance, in the context of fuzzy measure theory [11]. It turns out that the positive result obtained for additive functions *cannot* be generalised in the expected manner. However, the notion of k-additivity suggests an alternative representation of utility functions that can be usefully exploited in other ways. It does, for instance, often allow for a more concise representation of utility functions.

In the final part of this abstract, we discuss connections to some well-known NP-complete *combinatorial optimisation* problems (namely, weighted set packing and the independent set problem). These can be used to prove *NP-hardness* results for the problem of finding a socially optimal resource allocation. We indicate these results with respect to both the standard representation of utility functions and the representation based on k-additivity. While *NP-completeness* can also be shown for an important subclass of the general problem, we briefly discuss the difficulty of establishing a general completeness result along the same lines. We also briefly discuss connections to *combinatorial auctions*.

2 Resource Allocation by Negotiation

An instance of our negotiation framework consists of a finite set of (at least two) *agents* $\mathcal{A}$ and a finite set of non-divisible *resources* $\mathcal{R}$. A resource *allocation* A is a partitioning of the set $\mathcal{R}$ amongst the agents in $\mathcal{A}$. For instance, given an allocation A with $A(i) = \{r_3, r_7\}$, agent i would own resources r_3 and r_7. Given a particular allocation of resources, agents may agree on a (multilateral) *deal* to exchange some of the resources they currently hold. In general, a single deal may involve any number of resources and any number of agents. It transforms an allocation of resources A into a new allocation A'; that is, we can define a deal as a pair $\delta = (A, A')$ of allocations (with $A \neq A'$).

Each agent $i \in \mathcal{A}$ is equipped with a *utility function* u_i mapping bundles of resources (subsets of $\mathcal{R}$) to real numbers. We abbreviate $u_i(A) = u_i(A(i))$ for the utility value assigned by agent i to the set of resources it holds for allocation A. While individual agents may have their own interests, as a system designer, we are interested in the *social welfare* associated with a given allocation. According to the aforementioned *utilitarian* programme, the social welfare of an allocation A is given by the sum of utilities exhibited by all the agents in the system:

$$sw(A) \quad = \quad \sum_{i \in \mathcal{A}} u_i(A)$$

One of the main questions we are interested in in the welfare engineering framework is under what circumstances negotiation between agents will result in an improvement, and eventually an optimisation, with respect to such a notion of social welfare.

A deal may be coupled with a number of monetary side payments to compensate some of the agents involved for an otherwise disadvantageous deal. We call a deal *rational* iff it results in a gain in utility (or money) that strictly outweighs a possible loss in money (or utility) for each of the agent involved in that deal. As shown in previous work [8], a deal is rational iff it results in an increase in utilitarian social welfare. Given this connection between the "local" notion of rationality and the "global" notion of social welfare, we can prove the following result on the sufficiency of rational deals to negotiate socially optimal allocations [8, 17]:

> *Any sequence of rational deals (with side payments) will eventually result in an allocation of resources with maximal utilitarian social welfare.*

This means that (i) there can be no infinite sequence of deals all of which are rational, and (ii) once no more rational deals are possible the agent society must have reached an allocation that has maximal social welfare. The crucial aspect of this result is that *any* sequence of deals satisfying the rationality condition will cause the system to converge to an optimal allocation. That is, whatever deals are agreed on in the early stages of the negotiation, the system will never get stuck in a local optimum and finding an optimal allocation remains an option throughout.

A drawback of the general framework is that the above result only hold if deals involving any number of resources and agents are admissible [8, 17]. In some cases this problem can be alleviated by putting suitable restrictions on the utility functions agents may use to model their preferences. Interesting special classes of utility functions to consider include *non-negative* functions (where an agent may not assign a negative utility to any bundle), *monotonic* functions (where the utility of a set of resources cannot be lower than the utility assigned to any of its subsets), or additive functions (described in a further section).

3 Representations of Utility Functions

Agents' utility functions may be represented in different ways. Maybe the most intuitive form is the *bundle* one which amounts to listing all bundles of resources to which the agent assigns a non-zero value. Clearly, this technique can soon become hard to handle, as they may be an $2^n - 1$ bundles to value in the worst case.

An alternative representation, introduced in the context of fuzzy measure theory [11], is based on the notion of *k-additive* functions. A utility function is called k-additive iff the utility assigned to a bundle of resources R can be represented as the sum of basic utilities ascribed to subsets of R with cardinality $\leq k$. More formally, a k-additive utility can be written as follows:

$$u_i = \sum_{T \subseteq \mathcal{R}, |T| \leq k} \left[\alpha_T^i \prod_{r \in T} I_j(r) \right]$$

Where $I_j(r)$ equals to 1 iff agent j owns resource r, 0 otherwise. For the sake of simplicity, the I_j will be omitted, and the utilities will be written $u_i = \sum_{T \subseteq \mathcal{R}, |T| \leq k} \left[\alpha_T^i \prod_{r \in T} r \right]$. Clearly, agent i enjoys an utility increased by α_T^i when owning items $r \in T$ *together* (that is, α_T^i represents the synergetic value of this items held together).

Both representations are equivalent in term of expressive power, in the sense that they both can represent all utility functions[1]. As both representations are equivalent regarding expressiveness, one should consider whether one is stricly better in term of simplicity. This proves not be the case, *i.e* there are cases where concise utilities in the bundle form corresponds to huge k-additive formulas, and the other way around. The k-additive form appears to be more concise in cases where there are synergies between items.

[1] Converting the k-additive form into the bundle form is trivial, whereas converting the bundle form into a k-additive form involves the Moebius transform.

In the field of *combinatorial auctions*, several bidding languages have also been introduced and studied [15]. There are links with the representations discussed above, which the details are currently under investigation.

4 Complexity of Deals with k-additive Utilities

It was shown in an earlier paper that in scenarios where utility functions may be assumed to be 1-additive (also called additive functions), it is possible to guarantee optimal outcomes even when agents only negotiate deals involving a single resource and a pair of agents at a time (so-called *one-resource-at-a-time deals*) [8]:

> *If all utility functions are additive, then any sequence of rational one-resource-at-a-time deals (with side payments) will eventually result in an allocation of resources with maximal utilitarian social welfare.*

This result is of great practical relevance, because it shows that it is sufficient to design negotiation protocols for pairs of agents (rather than larger groups) and single resources (rather than sets) for applications in which preferences can be modelled in terms of additive utility functions.

Intuitively, we could have expected a similar result for k-additive utilities with $k \geq 2$. However, we will see that the deals required to reach maximal social welfare in the k-additive case are much more complex:

> *If all utility functions are k-additive, then rational n-resource-at-a-time deals may be needed to reach maximal utilitarian social welfare.*

To prove this result, let us build an example with 2-additive utility functions in which n-resource-at-a-time deals are needed. Consider 2 agents sharing n resources $\{r_1, r_2, ...r_n\}$, with the following 2-additive utility functions : $u_1 = 0$ and $u_2 = r_1 - r_1.r_2 - r_1.r_3 - r_1.r_4 - \ldots - r_1.r_n$. Let A_{init} be the initial allocation describing which agent owns which resource at time 0, and let A_{opt} be the allocation maximizing the utilitarian social welfare.

	A_{init}	A_{opt}
a_1	$\{r_1\}$	$\{r_2, r_3, \ldots, r_n\}$
a_2	$\{r_2, r_3, \ldots, r_n\}$	$\{r_1\}$

Here, $sw(A_{init}) = 0$ and $sw(A_{opt}) = 1$. In fact, the *only* allocation which has a social welfare greater than $sw(A_{init})$ is A_{opt}. Thus, the only rational deal here is $\delta(A_{init}, A_{opt})$, which is a bilateral deal of n resources at a time. Furthermore, we showed that adding constraints on the utility functions did not drastically decrease the complexity of the required deals:

> *If all utility functions are k-additive, monotone and super-additive, k-resources-at-a-time deals will not be sufficient in all cases.*

5 Connections to Combinatorial Optimisation

If we view the problem of finding an allocation with maximal social welfare as an algorithmic problem faced by a central authority (rather than as a problem of designing suitable negotiation mechanisms), then we can observe an immediate relation to the so-called *winner determination problem* in combinatorial auctions [15, 16, 18]. In a combinatorial auction, bidders can put in bids for different *bundles* of items (rather than just single items). After all bids have been received, the auctioneer has to find an allocation for the items on auction amongst the bidders in a way that maximises his revenue. If we interpret the price offered for a particular bundle of items as the utility the agent in question assigns to that set, then maximising revenue (i.e. the sum of prices associated with winning bids) is equivalent to finding an allocation with maximal utilitarian social welfare. This equivalence holds, at least, in cases where the optimal allocation of items in an auction is such that *all* of the items on auction are in fact being sold (so-called *free disposal*).

Winner determination in combinatorial auctions is known to be NP-complete [16]. The quoted result applies to the case of the "standard" bidding language, which allows bidders to specify prices for particular bundles and makes the implicit assumption that they are prepared to obtain any number of disjoint bundles for which they have submitted a bid (Nisan [15] calls this the "OR language"). Our languages for expressing utilities are more general than this. Still, the correspondence to combinatorial auctions suggests that the problem of finding an allocation with maximal utilitarian social welfare is at least NP-hard. We can make this observation more precise by showing how our problem relates to well-known NP-complete "reference problems".

We use a reduction to *weighted set packing* [3] to establish NP-hardness with respect to the standard (tabular) representation of utility functions:[2]

> *The problem of finding an allocation with maximal utilitarian social welfare whith utilites represented in bundle form is NP-hard.*

By further exploiting the correspondence to weighted set packing, we can also show that our problem is NP-complete provided the following conditions are met: (i) all utility functions are non-negative; (ii) $u_i(\{\,\}) = 0$ for all agents $i \in \mathcal{A}$; and (iii) we have free disposal, i.e. for any allocation not covering all resources there is always a full allocation that is not worse (this is the case, for instance, if the society includes at least one agent with a monotonic utility function). The first restriction comes from the fact that weights are required to be positive. The second reflects the fact that agents not included in any of the sets in the optimal allocation will not contribute to the measure either. Finally, the third restriction is due to the fact that an optimal solution to the weighted set packing problem is not required to cover all the items present in any of the sets. This latter point appears to be the main issue that makes our resource allocation problems potentially harder than optimisation problems without such a constraint. We are currently investigating this issue further.

A related combinatorial optimisation problem, namely *maximal independent set* can be used to establish also NP-hardness with respect to our representation based on k-additive functions:

> *The problem of finding an allocation with maximal utilitarian social welfare with utilities in k-additive form is also NP-hard.*

The reason why it is important to have these two variants of the NP-hardness result is that neither one of our two alternative representations for utility functions is the more concise in all cases, as previously discussed.

6 Conclusion

In this extended abstract, we have given a brief overview of recent work on multiagent resource allocation in the context of the welfare engineering framework and we have hinted

[2]Proofs are omitted in this abstract; we only sketch the basic idea of the reduction. An instance of the weighted set packing problem is a collection of sets, each of which is associated with a positive weight. A solution to the problem is a collection of disjoint sets out of the full collection and the measure with respect which such a solution is evaluated is the sum of the associated weights. The problem of finding a solution that maximises this measure is known to be NP-complete. Broadly speaking, we can reduce our problem of finding a socially optimal resource allocation to weighted set packing by introducing a set for each pair of agents i and bundles R such that $u_i(R) \neq 0$, where the associated weight is $u_i(R)$. Then the measure used in weighted set packing corresponds to the utilitarian social welfare.

at some of the connections to other fields, in particular combinatorial optimisation, we are currently exploring. We see this work as part of a wider research trend, that brings together ideas from different areas including microeconomics, game theory, complexity theory, and algorithm design.

Finally, we would like to stress that the high complexity of of our negotiation framework does not, at least not necessarily, mean that it cannot be usefully applied in practice. This view is supported by the fact that, in recent years, several algorithms for winner determination in combinatorial auctions (a problem of comparable complexity to the problems arising in the context of welfare engineering) have been proposed and applied successfully [10, 16, 18].

Acknowledgements. We would like to thank Jérôme Lang for many fruitful discussions on the topics addressed in this paper, and in particular for suggesting to consider k-additive utility functions.

References

[1] K. J. Arrow. *Social Choice and Individual Values.* John Wiley and Sons, 2nd edition, 1963.

[2] K. J. Arrow, A. K. Sen, and K. Suzumura, editors. *Handbook of Social Choice and Welfare*, volume 1. North-Holland, 2002.

[3] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation.* Springer-Verlag, 1999.

[4] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. Technical Report ULCS-04-001, Department of Computer Science, University of Liverpool, 2004.

[5] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. Technical Report ULCS-03-002, Department of Computer Science, University of Liverpool, 2003.

[6] U. Endriss and N. Maudet. On the communication complexity of multilateral trading. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2004).* IEEE Computer Society Press, 2004. To appear.

[7] U. Endriss and N. Maudet. Welfare engineering in multiagent systems. In *Engineering Societies in the Agents World IV*, volume 3071 of *LNAI.* Springer-Verlag, 2004.

[8] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On optimal outcomes of negotiations over resources. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2003).* ACM Press, 2003.

[9] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish routing in non-cooperative networks: A survey. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS-2003)*, volume 2747 of *LNCS*. Springer-Verlag, 2003.

[10] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-1999).* Morgan Kaufman, 1999.

[11] M. Grabisch. k-order additive discrete fuzzy measures and their representation. *Fuzzy Sets and Systems*, 92:167–189, 1997.

[12] J. M. Kleinberg, C. H. Papadimitriou, and P. Raghavan. Segmentation problems. In *Proceedings on 30th Annual ACM Symposium on Theory of Computing (STOC-1998).* ACM, 1998.

[13] M. Lemaître, G. Verfaillie, and N. Bataille. Exploiting a common property resource under a fairness constraint: A case study. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-1999).* Morgan Kaufmann, 1999.

[14] H. Moulin. *Axioms of Cooperative Decision Making.* Cambridge University Press, 1988.

[15] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC-2000)*, 2000.

[16] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinational auctions. *Management Science*, 44(8):1131–1147, 1998.

[17] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proceedings of the AAAI Spring Symposium: Satisficing Models*, 1998.

[18] T. W. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, 2002.

[19] A. K. Sen. *Collective Choice and Social Welfare*. Holden Day, 1970.

[20] D. H. Wolpert and K. Tumer. An introduction to collective intelligence. Technical Report NASA-ARC-IC-99-63, NASA Ames Research Center, 1999.