

Multiagent Resource Allocation with Sharable Items: Simple Protocols and Nash Equilibria

Stéphane Airiau Ulle Endriss

Institute for Logic, Language and Computation
University of Amsterdam

ABSTRACT

We study a particular multiagent resource allocation problem with indivisible, but sharable resources. In our model, the utility of an agent for using a bundle of resources is the difference between the valuation of that bundle and a congestion cost (or delay), a figure formed by adding up the individual congestion costs of each resource in the bundle. The valuation and the delay can be agent-dependent. When the agents that share a resource also share the resource's control, the current users of a resource will require some compensation when a new agent wants to use the resource. We study the existence of distributed protocols that lead to a social optimum. Depending on constraints on the valuation functions (mainly modularity), on the delay functions (e.g., convexity), and the structural complexity of the deals between agents, we prove either the existence of some sequences of deals or the convergence of all sequences of deals to a social optimum. When the agents do not have joint control over the resources (i.e., they can use any resource they want), we study the existence of pure Nash equilibria. We provide results for modular valuation functions and relate them to results from the literature on congestion games.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Algorithms, Economics, Theory

Keywords

Multiagent Resource Allocation, Congestion Games

1. INTRODUCTION

The generic problem of allocating a set of resources to a group of agents is a key problem in multiagent systems [3]. Some independent dimensions of a resource allocation problem are the type of resources (e.g., resources can be divisible or not, sharable or not), the allocation procedure (e.g., auctions or distributed mechanisms), and the criteria used to

Cite as: Multiagent Resource Allocation with Sharable Items: Simple Protocols and Nash Equilibria, S. Airiau and U. Endriss, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 167-174
Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

specify what makes an allocation optimal (e.g., maximising utilitarian or egalitarian social welfare, obtaining an envy-free division). Hence, there are many classes of problems to study in *multiagent resource allocation* (MARA).

Thus far, a great effort has been spent on distributed allocations of indivisible *nonsharable* resources [4, 5, 6, 7, 8]. In this setting, a resource is owned by an agent, and an agent receives a valuation for owning a set of resources. Synergies between resources can be taken into account, which plays an important role in the results. To improve an allocation, agents can exchange resources (sometimes in combination with a monetary transfer). This line of research has identified protocols that converge to optimal allocations for certain classes of valuation functions. In some cases, simple protocols (e.g., involving only two agents and one resource at a time) are sufficient, while more complex protocols are required in others. Simple protocols exist, for instance, for maximising utilitarian social welfare [5] and for finding envy-free divisions [4]. The goal of this paper is to extend this line of research to the case of indivisible resources that are *sharable*, as many resources are by their very nature sharable (for example, roads, supercomputers).

The problem of sharing a set of resources is not new, and has received much attention in the game theory literature. In particular, *congestion games* [11] feature agents that share a set of resources and obtain utility for each resource they use. For a resource, the utility obtained is a function of the number of agents using that resource. This class of games is of particular interest, as congestion games have the property of possessing pure-strategy Nash equilibria. In the seminal paper by Rosenthal [11], agents using the same resource receive the same payoff, and this payoff depends only on the number of agents that use that resource. Milchtaich [9] extended the model by allowing the payoffs to be player-dependent, but the existence of pure-strategy Nash equilibria is guaranteed only when each of the agents uses a single resource.

In a congestion game, the utility of an agent is the sum of the utilities received for each resource. In particular, these models do not take into account the synergies between resources. Our goal is to study distributed resource allocation problems where synergies between resources may exist and where resources can be shared. We introduce a model where the utility of an agent is the difference between a benefit from the set of resources it uses (a valuation function, as in the MARA framework for nonsharable items) and a cost that depends on the congestion of each resource (a delay function, as in congestion games). Our goal is first to define

simple protocols that lead to a socially optimal outcome, and we will restrict ourselves to utilitarian social welfare to define what constitutes a social optimum. Our second goal is to identify instances of our model in which the existence of a pure-strategy Nash equilibrium can be guaranteed.

Section 2 defines the model of MARA we shall be working with and recalls a number of relevant results from the literature. In Section 3, we present our results regarding simple protocols that permit (or even guarantee) convergence to a socially optimal allocation when individually rational agents negotiate; and in Section 4, we present our results concerning the existence of Nash equilibria.

2. THE MODEL

In this section, we introduce our model of MARA with sharable items; but first, we recall some details regarding the MARA framework with nonsharable items. Then we briefly discuss one specific issue arising in the context of sharable resources: the question of control of the resources.

2.1 MARA with Nonsharable Items

A MARA problem with indivisible nonsharable items [5, 8] is defined as a triplet $(\mathcal{N}, \mathcal{R}, \mathcal{V})$, where $\mathcal{N} = \{1, 2, \dots, n\}$ is a finite set of *agents*, \mathcal{R} is a finite set of *resources* (or *items*), and $\mathcal{V} = \langle v_1, \dots, v_n \rangle$ is a profile of *valuation functions* where $\forall i \in \mathcal{N}, v_i : 2^{\mathcal{R}} \rightarrow \mathbb{R}$. An *allocation* is a partition of the set of resources between the agents. A solution to a MARA problem is an allocation that satisfies certain properties. For example, we may want a solution to maximise utilitarian social welfare (i.e., the sum of the valuations of all agents), or be envy-free (no agent wants to swap resources with any other agent). In a framework where agents can use money, a *payment function* is a vector $p = \langle p_1, \dots, p_n \rangle$ such that $\sum_{i \in \mathcal{N}} p_i = 0$. When $p_i > 0$, agent i must make a payment. When $p_i < 0$, agent i receives a payment. A deal $\delta = (\sigma, \sigma')$ is a transformation from an allocation σ to an allocation σ' . A 1-deal is a deal involving the exchange of a single resource between two agents. A deal $\delta = (\sigma, \sigma')$ is *individually rational (IR)* if there exists a payment function p such that $\forall i \in \mathcal{N}, v_i(\sigma') - v_i(\sigma) > p_i$, except for agents i with $\sigma(i) = \sigma'(i)$ for whom $p_i = 0$ is also permitted. The following theorem applies to problems of this kind [8, 12]:

THEOREM 1. *For allocation problems with nonsharable items, any sequence of IR deals will eventually result in an allocation with maximal utilitarian social welfare.*

One drawback is that IR deals may be complex and involve many agents and resources. For *modular* valuation functions, satisfying $v(S \cup S') = v(S) + v(S') - v(S \cap S')$ for any sets $S, S' \subseteq \mathcal{R}$, however, simple deals involving only two agents and one resource are sufficient to reach a social optimum, as shown by the following theorem [5, 8]:

THEOREM 2. *For allocation problems with nonsharable items, if all valuation functions are modular, then any sequence of IR 1-deals will eventually result in an allocation with maximal utilitarian social welfare.*

2.2 MARA with Sharable Items

We now introduce a variant of the above MARA framework where resources are sharable. This is the framework we shall be working with for the remainder of this paper. A MARA

problem with indivisible sharable items is defined as a tuple $(\mathcal{N}, \mathcal{R}, (\Sigma_i)_{i \in \mathcal{N}}, (d_{i,r})_{i \in \mathcal{N}, r \in \mathcal{R}}, (v_i)_{i \in \mathcal{N}})$, and we present each term of the tuple in the following:

- $\mathcal{N} = \{1, 2, \dots, n\}$ is a finite set of n *agents*.
- \mathcal{R} is a finite set of m *resources* (or *items*).
- Σ_i is the set of *strategies* of agent i : it is a subset of the power-set of \mathcal{R} , i.e., $\forall i \in \mathcal{N}, \Sigma_i \subseteq 2^{\mathcal{R}}$. We will also refer to a strategy as a *bundle*. We denote by Σ the joint strategy space or the set of all allocations, i.e., $\Sigma = \prod_{i \in \mathcal{N}} \Sigma_i$. We will call $\sigma \in \Sigma$ a *strategy profile* or an *allocation*, i.e., $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle$ where $\forall i \in \mathcal{N}, \sigma_i \in \Sigma_i$. We shall use the terms strategy profile (from the congestion game literature) and allocation (from the MARA literature) interchangeably.
- $d_{i,r} : \{1, \dots, n\} \rightarrow \mathbb{R}$ is the *delay* perceived, or *cost* experienced by agent i when using resource r . The delay depends only on the number of agents that use r . Let $n_r(\sigma)$ be the number of agents that use resource r in allocation σ , i.e., $n_r(\sigma) = |\{i \in \mathcal{N} \mid r \in \sigma_i\}|$. The delay of r experienced by agent i in allocation σ is then $d_{i,r}(n_r(\sigma))$. We shall assume that the delay is a *nondecreasing* function in the number of agents using the resource, i.e., $d_{i,r}(k+1) \geq d_{i,r}(k)$ for all $k \in \mathbb{N}$. This models situations where an agent always prefers not to share the resource. When other agents also use the resource, the delay increases or remains the same. A (delay) function d is *convex* if $d(k+2) - d(k+1) \geq d(k+1) - d(k)$ for all $k \in \mathbb{N}$; it is *concave* if $d(k+2) - d(k+1) \leq d(k+1) - d(k)$ for all $k \in \mathbb{N}$; and it is *linear* if it is both convex and concave. That is, if d is linear, then there exists an $\alpha \in \mathbb{R}$ such that $d(k) = k \cdot \alpha$ for all $k \in \mathbb{N}$.
- $v_i : \Sigma_i \rightarrow \mathbb{R}$ is the *valuation function* for agent i : for each strategy $\sigma_i \in \Sigma_i$, i.e. for a bundle σ_i of resources used by agent i , $v_i(\sigma_i)$ is the utility of using the set σ_i , irrespective of the congestion. To simplify presentation, we shall assume that all valuation functions are *normalised*, i.e., $v_i(\emptyset) = 0$ for all agents $i \in \mathcal{N}$ (this assumption does not affect any of our results). For *modular* valuation functions (as defined in Section 2.1), we sometimes write $v_i(r)$ for $v_i(\{r\})$.

Now, the *utility* of agent i in profile σ is defined as

$$u_i(\sigma) = v_i(\sigma_i) - \sum_{r \in \sigma_i} d_{i,r}(n_r(\sigma)).$$

That is, agent i receives a benefit from using the resources of the bundle σ_i , but this benefit is reduced by the effects of the congestion.

We say that a MARA problem with sharable resources is *symmetric* when, for any given resource, all agents experience the same delay, i.e., $\forall i \in \mathcal{N} d_{i,r} = d_r$.

Observe that the original MARA framework can be simulated within the MARA framework with sharable items: using a delay function with a very high delay for $n \geq 2$, it will not be rational for an agent to share any item (for example, $d_{i,r}(1) = 0$ and $\forall k \geq 2, d_{i,r}(k) > \max_{\sigma_i \in \Sigma_i} v_i(\sigma_i)$, assuming positive valuation functions).

When the valuation functions are modular, we get a congestion game with the delay function of resource r for agent i being $d_{i,r}^*(k) = v_i(\{r\}) - d_{i,r}(k)$, i.e., a player-specific game as in the work of Milchtaich [9] and Ackermann et al. [1].

We are interested in distributed mechanisms, i.e., there is no central entity that knows the valuation and delay function

of each agent, and that decides on the allocation. Agents need to communicate and are allowed to exchange resources. As before, a *deal* $\delta = (\sigma, \sigma')$ is a pair of allocations. We consider the following types of *simple deals*:

- **ADD**(i, r): agent i adds to its bundle a single resource it is not currently using. For $r \notin \sigma_i$, agent i will have $\sigma_i \cup \{r\}$ after the **ADD**(i, r) action.
- **DROP**(i, r): agent i drops a resource it currently uses. i.e., after the drop, agent i will use $\sigma_i \setminus \{r\}$.
- **SWAP**(i, j, r): agent i swaps the use of resource r with agent j , i.e., agent i drops the use of r and agent j adds the resource.
- **1-deal**: a deal that concerns a single item, but possibly multiple agents.

With nonsharable resources, the utility of agents not taking part in the deal does not change. With sharable resources, the utility of agents currently using a resource that is part of the deal can be affected.

Deals may be coupled with monetary side payments. As in Section 2.1, a *payment function* is a vector $p = \langle p_1, \dots, p_n \rangle$ such that $\sum_{i \in \mathcal{N}} p_i = 0$. A deal $\delta = (\sigma, \sigma')$ is *individually rational (IR)* if there exists a payment function p such that $\forall i \in \mathcal{N}, u_i(\sigma') - u_i(\sigma) > p_i$, except for agents i unaffected by δ for whom $p_i = 0$ is also permitted. Here, an agent i is *unaffected* by a deal $\delta = (\sigma, \sigma')$ if $\sigma(i) = \sigma'(i)$ and $|\{j \in \mathcal{N} \mid r \in \sigma(j)\}| = |\{j \in \mathcal{N} \mid r \in \sigma'(j)\}|$ for all $r \in \sigma(i)$. Note that an agent i that does not change its bundle may still receive a payment (from agents starting to use resources i uses) or may make a payment (to agents that stop using resources i uses). Side payments are important as they make it possible for a single agent to start using a resource when the bundles of the other agents remain the same. Otherwise, with a congestion increase, the agents already using that resource would suffer an unacceptable loss in utility.

2.3 Resource Control

With nonsharable resources, agents have complete control over the resources they own. For example, if agent i wants to use a particular resource owned by agent j , j must agree to give up the item to i . With sharable resources, the notion of control is less clear. We can differentiate two variants.

In the first variant, agents are free to use any resource they wish. This means that there is no mechanism to prevent an agent from starting to use a resource. This relates to strategic games with self-interested agents.

In the second variant, agents must receive the consent of the agents using a resource before starting to use that resource. If the agents are rational, they will not accept that a new agent uses the resource if the delay function is strictly increasing. The only way to get access to a resource is either to compensate the current users with a side payment, or to perform a swap: to free other resource(s) that is (are) also used by the current users.

In this paper, we study both variants. In Section 3, we assume that agents accept and allow deals only when they are beneficial, in particular, all agents owning a resource must agree before allowing another agent to use that resource. We study mechanisms that lead to allocations maximising utilitarian social welfare. In Section 4, we assume that agents are non-cooperative and are free to use any resource they want. In that context, we investigate the problem of the existence of a pure-strategy Nash equilibrium.

Note that we could also assume that each resource has a

single owner, who permits other agents to use the resource. This is the case studied in the work of Bachrach and Rosenschein [2], in which an owner knows the private production function of the resource and other agents can bid to use it. The goal of that work is to find protocols where no agent has an incentive to lie (e.g., the owner of a resource should not lie about the production function). In our work, we assume that the resources are initially allocated to the agents, and they have to find an optimal allocation to use them.

3. OPTIMISING SOCIAL WELFARE

We now investigate a MARA problem with the following properties: (1) the resources are indivisible and sharable; (2) agents using a resource also share the control of that resource; and (3) side payments between agents are allowed. In the following, we shall seek to identify protocols that lead to an allocation that maximises *utilitarian social welfare*, i.e., that maximise the function $sw(\sigma) = \sum_{i \in \mathcal{N}} u_i(\sigma)$ (in the remainder of the paper, we will mostly just write “social welfare”). We will show that we can guarantee convergence of sequences of arbitrarily complex deals. We will also show that a sequence of deals involving a single resource is sufficient when the valuation functions are modular. Then, we will present results in which we allow only certain types of simple deals. We will prove the existence of some path in some cases, and convergence of all paths in others.

3.1 General Convergence Results

We first show that Theorems 1 and 2 from Section 2.1 also apply to the framework with sharable resources. Closely following the approach familiar from the framework with nonsharable resources [8], we first establish an important lemma showing that side payments can be arranged in such a way that a given deal is beneficial for all the agents involved if and only if that deal increases social welfare.

LEMMA 3. *A deal $\delta = (\sigma, \sigma')$ is IR iff $sw(\sigma) < sw(\sigma')$.*

PROOF. The proof of Lemma 1 in [8] goes through: That an IR deal necessarily increases social welfare is shown by summing the inequalities $u_i(\sigma') - u_i(\sigma) > p_i$ over all agents and noting that the sum of the payments must be zero. To prove that a deal is IR when social welfare increases, one can check that the following function is a valid payment function: $p_i = u_i(\sigma') - u_i(\sigma) - (sw(\sigma') - sw(\sigma))/n$. \square

It is now easy to prove the counterpart of Theorem 1:

THEOREM 4. *Any sequence of IR deals will eventually result in an allocation of resources with maximal social welfare.*

PROOF. The proof is very close to the corresponding proof in [8]: the number of allocations is finite and, by Lemma 3, any IR deal increases social welfare and any improvement in social welfare corresponds to an IR deal; so we must eventually reach an allocation maximising social welfare. \square

The significance of the theorem is that agents have no need to consider anything but their individual interests. Every single deal is bound to increase social welfare and there are no local optima the system could get stuck in. However, an IR deal may be quite complex as it may involve many agents and many resources at the same time. Finding such complex deals may turn out to be a difficult task. Under some constraints, it is possible to consider simpler deals. Indeed, we can prove a counterpart of Theorem 2:

THEOREM 5. *If all valuation functions are modular, then any sequence of IR 1-deals will eventually result in an allocation with maximal social welfare.*

PROOF. The set of resources and agents is finite, hence there is a finite number of allocations. Moreover, any IR deal strictly increases social welfare (see Lemma 3). Hence, the search for an allocation with maximal social welfare must terminate in a finite number of deals. As termination is guaranteed, we now must ensure there always exists an IR 1-deal from a suboptimal allocation.

Let σ and σ^* be two allocations such that σ^* maximise social welfare and $sw(\sigma) < sw(\sigma^*)$. We denote by $\delta_{\sigma_i, r}$ the characteristic function that returns 1 when resource r is in σ_i and 0 otherwise. We can write the social welfare of σ as: $sw(\sigma) = \sum_{i \in \mathcal{N}} (v_i(\sigma_i) - \sum_{r \in \sigma_i} d_{i,r}(n_r(\sigma))) = \sum_{r \in \mathcal{R}} \sum_{i \in \mathcal{N}} \delta_{\sigma_i, r} (v_i(\{r\}) - d_{i,r}(n_r(\sigma)))$. This expression shows that the utility generated by at least one resource must increase for social welfare to increase. Hence, a deal involving that single resource must exist for improving social welfare. In addition, by Lemma 1, this deal will be IR, which proves the theorem. \square

Theorem 5 is independent from any assumptions regarding the delay functions; only the valuation functions are required to be modular. Under this condition, by means of a sequence of deals concerning a *single resource* each, it is possible to reach an allocation that maximises social welfare. However, each deal may involve *many agents* at the same time.

It is not always possible to decompose a complex deal into a sequence of only ADD- or only DROP-deals: SWAP-deals are sometimes needed. For example, consider the following resource allocation problem with two agents i and j and one resource r : the valuation functions are $v_i(r) = 4$ and $v_j(r) = 6$ and both agents have the same delay function defined by $d_r(1) = 2$ and $d_r(2) = 5$. Let us assume that agent i uses r , obtaining a utility of $4 - 2 = 2$. The action $\text{ADD}(j, r)$ is not rational as the utility of agent i would drop to $4 - 5 = -1$ and agent j would receive $6 - 5 = 1$, which is not sufficient to compensate the drop of agent i . Only $\text{SWAP}(i, j, r)$ would be rational: agent j would get a utility of $6 - 2 = 4$, which is enough to compensate the drop of utility of agent i (who loses 2 units of utility).

We now may ask whether Theorem 5 can be strengthened by only allowing certain types of 1-deals, in particular ADD-, DROP-, and SWAP-deals.

3.2 ADD-Deals only from Empty Allocation

Let us first consider the case of protocols that only permit ADD-deals. Clearly, for this case we cannot hope for a convergence theorem, even under the strongest assumptions on the delay functions, and even if the initial allocation is the empty allocation. A simple counterexample would be the case where an agent who has low (but above zero) valuation for a resource r claims that resource first, after which no sequence of ADD-deals could possibly still lead to an optimal allocation (assuming there are many slower agents who place a higher valuation on r).

For the case of MARA with nonsharable items, in the face of failure of convergence by means of simple IR deals, Dunne et al. [7] and Dunne and Chevalyere [6] have studied the problem of checking whether it is at least the case that a sequence of deals of the desired type leading to an optimal allocation *exists* for a given scenario (the cited works anal-

yse the computational complexity of this kind of problem). This is an interesting question also for our framework: For a given allocation problem, does there exist a sequence of IR ADD-deals leading from the initial allocation to an optimal allocation? Maybe somewhat surprisingly, we will be able to give a positive answer to this question whenever the initial allocation is the empty allocation and all delay functions are nondecreasing and convex (symmetry is not required). We first prove the following lemma:

LEMMA 6. *For allocation problems with a single resource r , if all delay functions are nondecreasing and convex, and if $sw(\sigma) < sw(\sigma^*)$ and $N \subset N^*$ for two allocations σ and σ^* with corresponding sets $N = \{i \in \mathcal{N} \mid r \in \sigma_i\}$ and $N^* = \{i \in \mathcal{N} \mid r \in \sigma_i^*\}$, then there exists an agent $j \in N^* \setminus N$ such that the deal $\text{ADD}(j, r)$ will be IR in allocation σ .*

PROOF. We will show that $\text{ADD}(j, r)$ is IR for any agent $j \in \text{argmax}_i \{v_i(r) - d_{i,r}(|N|) \mid i \in N^* \setminus N\}$. From $sw(\sigma^*) > sw(\sigma)$ we get:

$$\sum_{i \in N^*} v_i(r) - d_{i,r}(|N^*|) > \sum_{i \in N} v_i(r) - d_{i,r}(|N|)$$

Let $\ell = |N^* \setminus N|$. Simplifying above inequality, and dividing by ℓ yields:

$$\frac{1}{\ell} \sum_{i \in N^* \setminus N} v_i(r) - d_{i,r}(|N|) > \frac{1}{\ell} \sum_{i \in N^*} d_{i,r}(|N^*|) - d_{i,r}(|N|)$$

Given our choice of j , this entails:

$$v_j(r) - d_{j,r}(|N|) > \frac{1}{\ell} \sum_{i \in N^*} d_{i,r}(|N^*|) - d_{i,r}(|N|)$$

As each $d_{i,r}$ is *convex*, we have $\frac{1}{\ell} [d_{i,r}(|N^*|) - d_{i,r}(|N|)] \geq d_{i,r}(|N| + 1) - d_{i,r}(|N|)$ for any agent i ; and thus:

$$v_j(r) - d_{j,r}(|N|) > \sum_{i \in N^*} d_{i,r}(|N| + 1) - d_{i,r}(|N|)$$

Now we subtract $d_{j,r}(|N| + 1) - d_{j,r}(|N|)$ on either side of the inequality:

$$v_j(r) - d_{j,r}(|N| + 1) > \sum_{i \in N^* \setminus \{j\}} d_{i,r}(|N| + 1) - d_{i,r}(|N|)$$

As each $d_{i,r}$ is *nondecreasing*, the term $d_{i,r}(|N| + 1) - d_{i,r}(|N|)$ is nonnegative for all i , and we can subtract it from the righthand side any number of times. Note that $N \subseteq N^* \setminus \{j\}$. Thus:

$$v_j(r) - d_{j,r}(|N| + 1) > \sum_{i \in N} d_{i,r}(|N| + 1) - d_{i,r}(|N|)$$

The lefthand side of this inequality is the utility gain of j for adding r to her bundle in allocation σ ; the righthand side is the loss in utility of the agents already holding r . That is, above inequality expresses that the deal $\text{ADD}(j, r)$ will be IR in allocation σ . \square

The existence of an ADD-path from the empty allocation to an optimal allocation now follows almost immediately:

THEOREM 7. *If all valuation functions are modular and all delay functions are nondecreasing and convex, then there exists a sequence of IR ADD-deals leading from the empty allocation to an allocation with maximal social welfare.*

PROOF. We have seen earlier that in *modular* domains we can let agents negotiate over resources on an item-by-item basis. So it suffices to prove the claim for scenarios with just a single resource r .

Let σ^* be an optimal allocation and let $N^* = \{i \in \mathcal{N} \mid r \in \sigma_i^*\}$ be the set of agents holding r in that allocation. Now consider any suboptimal allocation σ with $N \subset N^*$ for $N = \{i \in \mathcal{N} \mid r \in \sigma_i\}$. By Lemma 6, there exists an agent $j \in N^* \setminus N$ such that the deal $\text{ADD}(j, r)$ is IR from σ . As the initial allocation (i.e., the empty allocation) satisfies the conditions required for Lemma 6 to apply and as any new allocation produced this way satisfies the same conditions, this shows that there always exists a finite sequence of IR ADD-deals leading from the initial allocation to σ^* . \square

Note that this result does *not* suggest any obvious protocol for *finding* such an optimal sequence. The reason is that it will be difficult for the agents to find out which agent j should claim r at any given stage: in the proof (of Lemma 6), j is defined as an agent belonging to the set $N^* \setminus N$, which is unknown to the agents.

The restriction to convex delay functions in Theorem 7 is *necessary*: If (some) agents have strictly concave delay functions, then we can systematically construct examples where there exists no IR ADD-path from the empty to an optimal allocation. For instance, suppose there are a single resource r and three agents with the same valuation function v with $v(r) = 5$ and $v(\emptyset) = 0$, and the same concave delay function d_r with $d_r(1) = 0$ and $d_r(k) = 3$ for $k > 1$. Then, if no agent claims r , social welfare will be 0; if one agent claims r , social welfare will be 5; if two agents do, it will be $2 \cdot (5 - 3) = 4$; and if all three claim r , it will be $3 \cdot (5 - 3) = 6$ (maximal). But the full allocation cannot be reached from the empty allocation via an IR ADD-path, since adding the second agent would result in a loss of social welfare and thus not be IR (cf. Lemma 3). This situation is reminiscent of the *maximality theorems* of Chevalyre et al. [5], who amongst other things show that no class of valuation functions strictly including the modular functions will permit convergence by means of IR 1-deals (for allocation problems with nonsharable items).

3.3 DROP-Deals only from Full Allocation

Next, we present a similar result for protocols that only allow for IR DROP-deals. Here we are able to establish a path-existence property if we start from the *full* (rather than the empty) allocation. Again, the core of the argument is in a technical lemma:

LEMMA 8. *For allocation problems with a single resource r , if all delay functions are nondecreasing and convex, and if $sw(\sigma) < sw(\sigma^*)$ and $N \supset N^*$ for two allocations σ and σ^* with corresponding sets $N = \{i \in \mathcal{N} \mid r \in \sigma_i\}$ and $N^* = \{i \in \mathcal{N} \mid r \in \sigma_i^*\}$, then there exists an agent $j \in N \setminus N^*$ such that the deal $\text{DROP}(j, r)$ will be IR in allocation σ .*

PROOF. The proof is similar to that of Lemma 6; so we only give a compressed version here. We will show that $\text{DROP}(j, r)$ is IR for any agent $j \in \text{argmin}_i \{v_i(r) - d_{i,r}(|N|) \mid i \in N \setminus N^*\}$. Let $\ell = |N \setminus N^*|$. From $sw(\sigma^*) > sw(\sigma)$, after some rewriting and dividing by ℓ , we get:

$$\frac{1}{\ell} \sum_{i \in N^*} d_{i,r}(|N|) - d_{i,r}(|N^*|) > \frac{1}{\ell} \sum_{i \in N \setminus N^*} v_i(r) - d_{i,r}(|N|)$$

Given our choice of j , this entails:

$$\frac{1}{\ell} \sum_{i \in N^*} d_{i,r}(|N|) - d_{i,r}(|N^*|) > v_j(r) - d_{j,r}(|N|)$$

As each $d_{i,r}$ is *convex*, we have $d_{i,r}(|N|) - d_{i,r}(|N| - 1) \geq \frac{1}{\ell} [d_{i,r}(|N|) - d_{i,r}(|N^*|)]$ for all i ; and as each $d_{i,r}$ is *nondecreasing*, we have $d_{i,r}(|N|) - d_{i,r}(|N| - 1) \geq 0$ and we can add this term any number of times on the lefthand side:

$$\sum_{i \in N} d_{i,r}(|N|) - d_{i,r}(|N| - 1) > v_j(r) - d_{j,r}(|N|)$$

The righthand side of this inequality is the utility lost by agent j if she drops r ; the lefthand side is the cost saved by the other agents holding r . That is, this inequality states that the deal $\text{DROP}(j, r)$ is IR in σ . \square

THEOREM 9. *If all valuation functions are modular and all delay functions are nondecreasing and convex, then there exists a sequence of IR DROP-deals leading to an allocation with maximal social welfare from the full allocation.*

PROOF. The claim follows from Lemma 8 in the same way as Theorem 7 did follow from Lemma 6. \square

3.4 Mix of ADD/DROP/SWAP-Deals

We now turn our attention to more powerful protocols, with the aim of deriving convergence rather than just path-existence theorems. An important first result shows that if we allow all of our three simple types of deals (ADD, DROP, and SWAP), then we can get convergence from *any* initial allocation, albeit under stronger restrictions on the delay functions (namely, we now require symmetry):

THEOREM 10. *If all valuation functions are modular and all delay functions are symmetric as well as nondecreasing and convex, then any sequence of IR ADD-, DROP-, and SWAP-deals will converge to an allocation with maximal social welfare.*

PROOF. As we are operating in modular domains, it suffices to prove the claim for allocation problems with a single resource r . Let σ be any suboptimal allocation and let $N = \{i \in \mathcal{N} \mid r \in \sigma_i\}$. All we need to prove is that there exists an IR ADD-, DROP-, or SWAP-deal starting from σ . This will show that even when the protocol is restricted to these deal types, we can never get stuck in a suboptimal allocation; and as social welfare improves with every IR deal (cf. Lemma 3), we must eventually reach an optimal allocation.

To simplify the presentation, we shall assume that no two agents give the same value to r , i.e., $v_i(r) \neq v_j(r)$ whenever $i \neq j$, but the proof easily extends to the general case. Define for each $k \leq n$ the allocation σ^k as follows: $r \in \sigma_i^k$ if and only if $\#\{j \in \mathcal{N} \mid v_j(r) > v_i(r)\} < k$, i.e., this is the allocation where the k top agents (in terms of valuing r) obtain r . Observe that, since the delay functions are symmetric, amongst all allocations assigning exactly k agents to r , allocation σ^k has maximal social welfare.

Now, let $k = |N|$ be the number of agents holding r in the current allocation σ . We distinguish three cases:

- (1) $\sigma \neq \sigma^k$: Then there exists an agent j with $r \notin \sigma_j$ such that $v_j(r) > v_i(r)$ for some agent i with $r \in \sigma_i$, i.e., the deal $\text{SWAP}(i, j, r)$ will be IR (here we are using the assumption that delay functions are symmetric).

- (2) $\sigma = \sigma^k$ and there exists a $k^* > k$ with $sw(\sigma^{k^*}) > sw(\sigma)$: Then we must have $N \subset N^*$ for $N^* = \{i \in \mathcal{N} \mid r \in \sigma_i^*\}$, because in both allocations the k top agents obtain r . Thus, as delay functions are nondecreasing and convex, Lemma 6 applies and we can infer that there exists an IR ADD-deal.
- (3) $\sigma = \sigma^k$ and there exists a $k^* < k$ with $sw(\sigma^{k^*}) > sw(\sigma)$: Then we must have $N \supset N^*$. Thus, as delay functions are nondecreasing and convex, Lemma 8 applies and there exists an IR DROP-deal.

There are no further cases, so we are done. \square

This result is stronger than Theorem 5 in the sense that it relies on a simpler class of deals (never involving more than two agents at a time); it is weaker in the sense that it requires stronger (but not unreasonable) restrictions to the range of admissible delay functions. Compared to Theorems 7 and 9, Theorem 10 establishes again a convergence property, rather than just the existence of a path.

The symmetry assumption in Theorem 10 is *necessary*: For example, if $v_1(r) = 10$ and $d_{1,r}(k) = 6k$, and $v_i(r) = 5$ and $d_{i,r}(k) = k$ for $i \in \{2, 3\}$, then the optimal allocation where agents 2 and 3 hold r is not reachable from the allocation where only agent 1 holds r by means of ADD-, DROP-, and SWAP-deals alone. Convexity is also a necessary condition (see the example at the end of Section 3.2).

3.5 Mix of ADD/SWAP-Deals with Control

Finally, we want to explore convergence for protocols using just two of our simple deals, namely ADD and SWAP. As we shall see, in this case we can prove convergence (from the empty allocation) if we add an additional “control component” that allows agents to avoid certain dead-ends. We shall suggest two such control mechanisms for this setting. Both results will heavily rely on the following technical lemma:

LEMMA 11. *For allocation problems with a single resource r , if all delay functions are symmetric as well as nondecreasing and convex, then any sequence of IR ADD- and SWAP-deals starting from the empty allocation will converge to an allocation with maximal social welfare, provided no ADD-deals are applied once k^* agents are holding r , where k^* is the maximum number of agents holding r in any allocation with maximal social welfare.*

PROOF. Inspection of the proof of Theorem 10 shows that as long as the number of agents currently holding r is at most k^* , either an IR ADD- or an IR SWAP-deal will be available (or an optimal allocation has already been reached). Provided we never apply an ADD-deal once k^* agents hold r , this condition will continue to be satisfied. The claim of the lemma follows. \square

The next theorem shows that there are natural protocols for which the (seemingly cumbersome) precondition for the applicability of Lemma 11 is satisfied:

THEOREM 12. *If all valuation functions are modular and all delay functions are symmetric as well as nondecreasing and convex, then any sequence of IR ADD- and SWAP-deals starting from the empty allocation will converge to an allocation with maximal social welfare, provided ADD-deals are only applied when no SWAP-deal is IR.*

PROOF. Due to modularity, we can restrict attention to allocation problems with a single resource r and Lemma 11

becomes applicable. Let k^* be the maximal number of agents holding r in an optimal allocation. All we need to show is that once k^* agents do hold r , no ADD-deal will ever be applied. But this is clearly so if ADD-deals are only applied when no more SWAP-deals are IR. \square

In allocation σ , we say that an IR deal $\delta = (\sigma, \sigma')$ is *greedy* with respect to a set Δ of deals applicable in σ , if it produces maximal social surplus of all the deals in Δ ; that is, if $sw(\sigma') \geq sw(\sigma'')$ for all $\sigma'' \in \Delta$. A sequence of greedy deals of a given type is a sequence of deals for which the next deal is always the deal maximising social surplus over all applicable deals of the given type.

THEOREM 13. *If all valuation functions are modular and all delay functions are symmetric as well as nondecreasing and convex, then any sequence of greedy IR ADD- and SWAP-deals starting from the empty allocation will converge to an allocation with maximal social welfare.*

PROOF. Restricting once again attention to scenarios with a single resource r (permissible due to modularity), let k^* be the maximal number of agents holding r in an optimal allocation. We need to show that whenever a greedy protocol chooses an ADD-deal, then the number of agents currently holding r is still less than k^* . By Theorem 12, the only critical case we need to account for is when there are both IR ADD- and SWAP-deals available.

To simplify presentation, assume $v_i(r) \neq v_j(r)$ whenever $i \neq j$ (this restriction is not crucial and the proof generalises easily). Let σ be the current allocation, let $N = \{i \in \mathcal{N} \mid r \in \sigma_i\}$, and let $k = |N|$. Let $j \in \operatorname{argmin}_i \{v_i(r) \mid r \in \sigma_i\}$ be the agent placing the lowest value on r amongst those holding r in σ ; and let $j' = \operatorname{argmax}_i \{v_i(r) \mid r \notin \sigma_i\}$ be the agent putting the highest value on r of those *not* holding r .

Then the best possible SWAP-deal is $\operatorname{SWAP}(j, j', r)$. It increases social welfare by a margin of $v_{j'}(r) - v_j(r)$. The best possible ADD-deal is $\operatorname{ADD}(j', r)$. It increases social welfare by $v_{j'}(r) - (k+1) \cdot d_r(k+1) + k \cdot d_r(k)$. Hence, under a greedy protocol, an ADD-deal will only be chosen if:

$$v_j(r) - (k+1) \cdot d_r(k+1) \geq -k \cdot d_r(k)$$

Now, let N^k be the set of the top k agents in terms of valuing r . As j valued r the least of all the $k+1$ agents in $N \cup \{j'\}$, we know that $j \notin N^k$, and we can rewrite above inequality as follows:

$$\sum_{i \in N^k \cup \{j\}} v_i(r) - d_r(k+1) \geq \sum_{i \in N^k} v_i(r) - d_r(k)$$

The lefthand side of this inequality is the social welfare generated if the $k+1$ agents in $N^k \cup \{j\}$ hold r ; the righthand side is the social welfare for the best possible allocation in which k agents hold r . That is, there are allocations in which $k+1$ agents claim r that are at least as good as the best allocation in which k agents do. Hence, $k^* > k$, which means that under a greedy protocol, an ADD-deal will only ever get applied if k^* has not yet been reached. The claim then follows from Lemma 11. \square

The control mechanism of Theorem 13 (greediness) may be more relevant in practice than that of Theorem 12 (giving SWAP precedence over ADD) because it is reasonable to assume that agents will actively search for deals giving them high profit first and thereby indirectly implement a sequence of deals that will at least be approximately greedy.

It is possible to derive corresponding results for protocols allowing for DROP- and SWAP-deals only, starting from the full allocation. We omit the details here for lack of space.

4. NON-COOPERATIVE MARA

In the MARA framework with nonsharable resources, the user of a resource had full control over it. With sharable resources, the previous section dealt with situations where the agents that share the use of a resource also share its control: the acceptable deals are IR, i.e., all the agents involved in a deal either benefit or are indifferent. In this section, we consider situations where no such control exists: an agent is free to use any resource it likes. In particular, if an agent can gain utility by using an additional resource (even if this means that other agents will suffer a decrease of their utility), the agent will add the resource.

A first question that arises concerns the existence of a Nash equilibrium: is there an allocation where each agent, assuming no other agent changes its bundle, has no incentive to use a different bundle. If the answer is negative, agents will never agree on an allocation. If the answer is positive, the problem will be for the agents to converge to such an allocation. For this study, the interesting actions are the ones performed by a *single agent*, i.e., an addition or a drop of a set of resources, or a combination of both.

Formally, a strategy profile $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle$ is called a *pure-strategy Nash equilibrium (pure NE)* if $\forall i \in \mathcal{N}$, $\nexists \sigma'_i \subseteq \mathcal{R}$ with $\sigma'_i \neq \sigma_i$ such that $u_i(\sigma') > u_i(\sigma)$, where $\sigma' = \langle \sigma, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots, \sigma_n \rangle$.

The existence of a pure NE is a key property of congestion games, as shown by the title of Rosenthal's seminal paper, "A class of games possessing pure-strategy Nash equilibria" [11]. In his model, the delay function of a given resource is the same for all agents. The agents, however, can have different strategies, i.e., they may be restricted to use certain bundles. Milchtaich [9] introduced player-specific delay functions, and proved the existence of NE when the strategies of the players are restricted to a set of singletons (i.e., an agent can only use one resource). Ackermann et al. [1] showed that the matroid property is a sufficient and maximal property on the structure of the player's strategy space for guaranteeing the existence of a pure NE.

Note that in our framework, we do not have any restriction on the strategies of the agents: each agent can access and use any of the resources. Rationality may restrict the set of strategies, e.g., it may not be rational to use some bundles of resources as this would lead to a negative utility for all possible congestions.

4.1 Examples

In the most general case for a MARA problem with sharable resources, there may not exist a pure NE. We consider the example in Table 1, adapted from Milchtaich [10], with two agents and six resources. Agents have the same delay function when they share a resource, and different when they use it on their own. For any bundle other than the ones indicated in the table the agents do not get any valuation. One can check that there is a cycle of best responses: when 2 uses $\{b, d\}$, the best response of 1 is to use f ; if 1 uses f , 2's best response is to use $\{a, c\}$; when 2 uses $\{a, c\}$, 1 should in turn use $\{a, d, e\}$; and finally, when 1 uses $\{a, d, e\}$, 2 should use $\{b, d\}$. Hence, there is no pure strategy NE.

Furthermore, when a pure NE does exist, the correspond-

resource	a	b	c	d	e	f	$v_1(\{a, d, e\}) = 100$
$d_{1,r}(1)$	20	45	48	20	16	65	$v_1(\{f\}) = 100$
$d_{2,r}(1)$	24	45	48	28	32	130	$v_2(\{b, d\}) = 100$
$d_{i,r}(2)$	28	45	48	30	48	195	$v_2(\{a, c\}) = 100$

Table 1: Example of game with no pure NE

ing allocation need not have maximal social welfare. For example, consider the following symmetric problem with two agents 1 and 2 and a single resource r : $v_1(r) = 4$, $v_2(r) = 3$, $d(1) = 0$ and $d(2) = 2$. The social optimum is when agent 1 uses the resource, with a value of $4 - 0 = 4$. However, agent 2 has an incentive to also use r as it can get $3 - 2 = 1$ instead of 0. In that situation, however, the social welfare drops to $4 - 2 + 3 - 2 = 3$, and agent 1 has no incentive to drop the resource. Hence, a suboptimal allocation is a NE.

4.2 MARA with Pure Nash Equilibria

In the following, we focus our attention on MARA problems possessing pure-strategy Nash equilibria. The first fact covers games where the cost of congestion is smaller than the valuation of a bundle; hence, the allocation where all agents use all resources is a pure NE.

FACT 14. *Every allocation game in which marginal valuation always exceeds delay, i.e., in which $v_i(\sigma \cup \{r\}) - v_i(\sigma) > d_{i,r}(k)$ for any $k \leq n$ (for all $i \in \mathcal{N}$, $\sigma \subseteq \mathcal{R}$, $r \in \mathcal{R} \setminus \sigma$), has got a pure NE.*

PROOF. The allocation where every agent claims every resource is a NE in this kind of game. (In fact, above inequality only needs to hold for $k = n$.) \square

In order to prove a result for MARA problems with modular valuation functions, we first prove a lemma regarding allocation games where the set of resources is a singleton. The result is not surprising in view of the results of Milchtaich [9] and Ackermann et al. [1].

LEMMA 15. *Every allocation game with a single resource and with nondecreasing delay functions has got a pure NE.*

PROOF. Let $A_0 = \mathcal{N}$ and for $k \geq 1$, let $A_k = \{i \in \mathcal{N} \mid v_i(r) - d_{i,r}(k) \geq 0\}$, i.e., A_k is the set of agents having nonnegative utility when k agents are using the resource. Let $k^* = \max\{k \in [0, n] \mid |A_k| \geq k\}$. Let A be a set of k^* agents such that $A_{k^*+1} \subseteq A \subseteq A_{k^*}$. Such a set A exists because $|A_{k^*}| \geq k^*$ by construction; $A_{k^*+1} \subseteq A_{k^*}$ by nondecreasingness of the delay functions and $|A_{k^*+1}| \leq k^*$ by maximality of k^* . We claim that the allocation where all agents in A use the resource is a pure NE. If $i \in A$, agent i gets a nonnegative utility, hence, i has no incentive to drop the resource. If $i \notin A$, then $i \notin A_{k^*+1}$. Consequently, agent i would not get positive utility if it added the resource. \square

We are now ready to present a theorem for allocation problems with modular valuation functions.

THEOREM 16. *Every allocation game with modular valuation functions and nondecreasing delay functions has got a pure NE.*

PROOF. For each resource $r \in \mathcal{R}$, Lemma 15 guarantees the existence of a pure NE σ_r . Let σ be the allocation where the strategy of each agent i is the union of strategies in $\sigma_r(i)$. Given that for modular valuation functions we can treat the problem item-by-item, the allocation σ is a NE. \square

Theorem	Result	Val.	Delay	Symm.	Deals	Init. Alloc.	Control
4	convergence	any	any	no	all	any	none
5	convergence	modular	any	no	1-deals	any	none
7	existence	modular	n.d.+convex	no	ADD	empty	none
9	existence	modular	n.d.+convex	no	DROP	full	none
10	convergence	modular	n.d.+convex	yes	ADD-DROP-SWAP	any	none
12	convergence	modular	n.d.+convex	yes	ADD-SWAP	empty	precedence
13	convergence	modular	n.d.+convex	yes	ADD-SWAP	empty	greedy

Table 2: Summary of convergence and path-existence results

The fact that the agents are allowed to use any resource and the modularity of the valuation functions make it possible to treat the problem issue by issue. In other words, we have a collection of independent 1-resource congestion games with player-specific valuation functions.

5. CONCLUSION

We have introduced a powerful and flexible model of multi-agent resource allocation with sharable items. The model integrates features from models developed in two different strands of the literature: the distributed approach to resource allocation in multiagent systems and congestion games studied in game theory. Most of our technical contributions focus on specific instances of the general model, particularly (but not exclusively) the case of allocation problems with agents that have modular valuation functions.

Our first set of results concerns conditions for the *convergence* to a social optimum by means of simple negotiation protocols. As for the previously studied case of nonsharable resources, we have seen that convergence can always be guaranteed when arbitrarily complex deals are available, and that deals involving just one resource suffice in modular domains. Unlike for nonsharable resources, in our scenario the latter type of deal may involve more than two agents, which calls for a finer analysis: we have been able to show that deals involving one resource and at most two agents suffice when the delay functions meet certain conditions and that the protocols can be further simplified by assuming that agents are greedy in the sense of making the most profitable deal first. We have proved the *existence* of a path to an optimum under weaker conditions. These results, summarised in Table 2, complement existing ones on convergence for different MARA scenarios and deepen our understanding of the area as a whole.

Our second set of results concerns the existence of pure-strategy Nash equilibria. In particular, we have been able to show that when valuation functions are modular (a strong condition) and when delay functions are nondecreasing (a very common and unproblematic assumption), then such an equilibrium will always exist. This ties in nicely with existing results in the literature on congestion games.

Most of our results apply to particular instances of the general model for multiagent resource allocation with sharable items, by imposing relevant restrictions on valuation functions, delay functions, or both. Future work should seek to explore further such instances. For instance, we may ask what types of protocols can guarantee convergence to a social optimum if the class of potential valuation functions is neither the class of modular functions nor the class of all set functions. We may also investigate conditions for convergence to allocations that are optimal in the sense of maximising the utility of the weakest agent (egalitarian so-

cial welfare) or in the sense of being envy-free. Regarding the existence of pure Nash equilibria, it would be interesting to see how far we can generalise the two classes of games for which we have obtained positive results without losing the guarantee of the existence of a pure Nash equilibrium.

6. REFERENCES

- [1] H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science*, 410(17):1552 – 1563, 2009.
- [2] Y. Bachrach and J. S. Rosenschein. Distributed multiagent resource allocation in diminishing marginal return domains. In *Proc. 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2008)*, 2008.
- [3] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.
- [4] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Reaching envy-free states in distributed negotiation settings. In *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI-2007)*, 2007.
- [5] Y. Chevaleyre, U. Endriss, and N. Maudet. Simple negotiation schemes for agents with simple preferences: Sufficiency, necessity and maximality. *Journal of Autonomous Agents and Multiagent Systems*, In press.
- [6] P. E. Dunne and Y. Chevaleyre. The complexity of deciding reachability properties of distributed negotiation schemes. *Theoretical Computer Science*, 396(1–3):113–144, 2008.
- [7] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1–2):23–46, 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] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
- [10] I. Milchtaich. Social optimality and cooperation in nonatomic congestion games. *Journal of Economic Theory*, 114(1):56–87, 2004.
- [11] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [12] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.