

The Power of Swap Deals in Distributed Resource Allocation

Anastasia Damamme, Aurélie Beynier, Yann Chevaleyre, Nicolas Maudet

► **To cite this version:**

Anastasia Damamme, Aurélie Beynier, Yann Chevaleyre, Nicolas Maudet. The Power of Swap Deals in Distributed Resource Allocation. The 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), May 2015, Istanbul, Turkey. pp.625-633, 2015. <hal-01212603>

HAL Id: hal-01212603

<https://hal.archives-ouvertes.fr/hal-01212603>

Submitted on 13 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Power of Swap Deals in Distributed Resource Allocation

Anastasia Damamme
Sorbonne Universités, UPMC
Univ Paris 06
CNRS, UMR 7606, LIP6
F-75005, Paris, France
anastasia.damamme@lip6.fr

Aurélie Beynier
Sorbonne Universités, UPMC
Univ Paris 06
CNRS, UMR 7606, LIP6
F-75005, Paris, France
aurelie.beynier@lip6.fr

Yann Chevaleyre
Université Paris 13, Sorbonne
Paris Cité
LIPN, CNRS, UMR 7030
F-93430, Villetaneuse, France
yann.chevaleyre@lipn.univ-
paris13.fr

Nicolas Maudet
Sorbonne Universités, UPMC
Univ Paris 06
CNRS, UMR 7606, LIP6
F-75005, Paris, France
nicolas.maudet@lip6.fr

ABSTRACT

In the simple resource allocation setting consisting in assigning exactly one resource per agent, the *top trading cycle* procedure stands out as being the undisputed method of choice. It remains however a centralized procedure which may not well suited in the context of multiagent systems, where distributed coordination may be problematic. In this paper, we investigate the power of dynamics based on rational bilateral deals (*swaps*) in such settings. While they may induce a high efficiency loss, we provide several new elements that temper this fact: (i) we identify a natural domain where convergence to a Pareto-optimal allocation can be guaranteed, (ii) we show that the worst-case loss of welfare is as good as it can be under the assumption of individual rationality, (iii) we provide a number of experimental results, showing that such dynamics often provide good outcomes, especially in light of their simplicity, and (iv) we prove the NP-hardness of deciding whether an allocation maximizing utilitarian or egalitarian welfare is reachable.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—Coherence and coordination, intelligent agents, Multiagent systems

General Terms

Algorithms, Economics

Keywords

Resource allocation; Negotiation

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.

Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1. INTRODUCTION

One of the most basic resource allocation setting involves n agents and n resources, ordered by each agent as a full ranking of preferences. It is known in economics as a *house allocation problem*, or a *house market* when agents initially hold resources as is the case in this paper. In this setting, the celebrated *top trading cycle* procedure [20] is known to satisfy a number of key desirable properties.

However, this procedure is centralized and involves potentially long cycles of resources reallocation among agents. When a system is distributed, such long cycles may not materialize easily: coordination among several agents is challenging [15], and thus agreeing on such exchanges may be unaffordable most of the times, already in terms of communication. Also implementing long cycles of exchanges may be problematic in practice, since it significantly increases the likelihood of failure.

In such circumstances, and if there is little or no cost in performing an exchange, a natural approach is to consider myopic dynamics based on sequences of local deals. The simplest version are bilateral deals, or *swaps*: agents randomly meet in a pairwise fashion, and contract a deal with their partner if exchanging their resources proves to be mutually beneficial. The process iterates until a stable state (an equilibrium) is reached. The same resource can thus be held successively by several agents over the sequence. The simplicity of the model discussed in this paper (a single resource per agent, preferences as strict linear orders) allows to directly answer some questions which are sometimes problematic: convergence is obviously guaranteed, and sequences of swaps have a length at most quadratic in n .

While the topic has been heavily studied, we provide some new insights which may suggest to reassess the power of swaps deals in such settings. Theoretical results hold for any specific dynamics based on swap deals. Experimental results are obtained with the least committing random dynamics (encounters between agents are drawn uniformly at random), under several cultures of preferences.

The remainder of this paper is as follows. In Section 2, we

provide the necessary background. Section 3 tackles the notion of Pareto-optimality of the allocations that are reached by such dynamics. In Section 4, we focus on the classical utilitarian social welfare (by using ranks as a measure of the utility of the resources obtained by agents). Section 5 investigates how well the approach does in terms of fairness, by studying the basic notion of (maxmin) egalitarian social welfare. Finally, Section 6 proves complexity results on the reachability problem in our context.

2. BACKGROUND

We start by giving the basic components of our model, and the specificities of the domains studied. Then, we give the details of the dynamics based on local deals, as well as the procedures used for comparison.

2.1 The Model

We start with a set $\mathcal{N} = \{a_1, \dots, a_n\}$ of agents, and a set $\mathcal{R} = \{r_1, \dots, r_n\}$ of resources. An allocation A is a mapping from agents to resources, where $A(a_i)$ stands for the object held by agent a_i , *i.e.* each agent holds exactly one resource at all time. In our model, agents are initially endowed with some resource, and A_0 is used by convention to denote the initial allocation.

Agents have preferences regarding the resources they may hold, expressed as linear orders (a complete ranking of all the resources, as in the classical house market setting). We note $r_1 \succ_1 r_2$ (sometimes $1 : r_1 \succ r_2$) the fact that agent a_1 prefers r_1 over r_2 , and $\succ_{\mathcal{N}}$ will stand for the collection of preferences of the group of agents \mathcal{N} . We also denote by $top(a_i)$ the top-ranked resource for agent a_i . Finally, we sometimes use the notation $A = [A(a_1)A(a_2) \dots A(a_n)]$ as a shortcut for an allocation, just abbreviated by the index of the resource held by each agent, *e.g.* $A = [25431]$.

An instance I of this resource allocation problem is thus a tuple $\langle \mathcal{N}, \mathcal{R}, \succ_{\mathcal{N}}, A_0 \rangle$. By an abuse of notation, we write $A \in I$ to simply denote an arbitrary allocation from the instance. We also denote by \mathcal{I} the set of all the instances.

2.2 Quality of Allocations

How good an allocation is can be given different interpretations. The first basic requirement is to seek for allocations that are Pareto-optimal, *i.e.* an allocation A^* such that there does not exist an allocation A'

$$\forall a_i \in \mathcal{N}, A'(a_i) \succeq_i A^*(a_i) \text{ and } \exists a_j \in \mathcal{N} : A'(a_j) \succ_j A^*(a_j)$$

Beyond Pareto-optimality, one may want to strengthen the efficiency and/or fairness requirement. Often, this will involve to interpret cardinally the ordinal preferential information provided by agents. We shall simply use the *rank* as a measure of satisfaction, and thus assign —using a Borda count— some utility to ranks, *i.e.* $u_i(A(i)) = n$ when a_i gets her preferred object, $n - 1$ when she gets her next preferred object, and so on. In this paper, we shall study:

- *utilitarian* social welfare:

$$sw_u(A) = \sum_{i \in \mathcal{N}} u_i(A(a_i))$$

- *egalitarian* social welfare:

$$sw_e(A) = \min_{i \in \mathcal{N}} (u_i(A(a_i)))$$

We note that it will be a strict constraint for the procedures studied in this work to respect *individual rationality* (IR), that is, no agent should be assigned an object less preferred than the one currently held.

Example 1. An example involving 5 agents, resources in boxes are initially held, *i.e.* $A_0 = [52143]$

1	:	r_1	\succ	r_3	\succ	r_2	\succ	r_4	\succ	r_5
2	:	r_2	\succ	r_1	\succ	r_3	\succ	r_4	\succ	r_5
3	:	r_1	\succ	r_5	\succ	r_2	\succ	r_3	\succ	r_4
4	:	r_3	\succ	r_5	\succ	r_2	\succ	r_1	\succ	r_4
5	:	r_4	\succ	r_5	\succ	r_2	\succ	r_3	\succ	r_1

It can be seen that this allocation is not Pareto-optimal, since $A = [42153]$ (for instance), dominates it. We have $sw_u(A) = 18$, and $sw_e(A) = 2$. Observe that we could obtain an allocation $A' = [12534]$ which would yield $sw_u(A') = 24$, but violates individual rationality for agent a_3 .

2.3 Preferences: Domains and Cultures

In this paper we deal with ordinal preferences, represented as linear orders (so, without ties). We will be interested in a domain restriction, *single-peaked* domains, well-known in voting. We argue that it is also a perfectly sensible domain in some resource allocation applications: think for instance of different items (like tee-shirts, devices for memory storage, etc.), which can be ranked according to a notion of size. Each agent may have a different, ideal, size which fits its purposes. Now, depending on their situation or personal taste, they also may have different attitudes as to whether they would rather get items of larger or smaller size if their ideal resource is not available (some agents may prefer to wear slightly oversized tee-shirts than slightly undersized, or agents may be more or less reluctant to buy additional storage resources, etc). As another example, the allocation of slots on a line is studied in [11].

In our experiments, preferences are randomly drawn from different *cultures*:

- *impartial culture* (IC): each linear ordering is drawn with uniform probability ($1/n!$). This is the classical culture, assuming no correlation between preferences.
- *uniform peak single peaked* (UP-SP): we assume there is a common axis $r_1 \triangleright r_2 \triangleright \dots \triangleright r_n$. Each preference ordering is drawn by first selecting a peak uniformly at random. The rest of the order is then completed by picking with $1/2$ probability either a resource on the left or on the right, until an extreme resource is reached.
- *real-world instances* (PL): we make use of the PrefLib [13] repository of real-world instances. Specifically, in our case, we exploit the sushi and tee-shirt datasets, with respectively $n = 7$ and $n = 9$ resources.

2.4 Dynamics based on Local Deals

The approach we take in this paper is *dynamic*: agents make local improving exchanges (or *deals*), until they reach a stable allocation, *i.e.* an allocation where no improving contract is possible [18]. Deals take the form of exchange cycles between agents. We classically denote a deal involving k agents as $(e_{(1)}e_{(2)} \dots e_{(k)})$, where by convention agent $e_{(1)}$

—the first agent involved in the cycle— gives her resource to agent $e_{(2)}$, agent $e_{(2)}$ gives her resource to agent $e_{(3)}$, and so on, concluding with agent $e_{(k)}$ giving her resource to agent $e_{(1)}$. Such a deal is rational if *all* agents are better off after the deal has been implemented:

Definition 1. A deal $(e_{(1)}e_{(2)} \dots e_{(l)})$ is rational when for all $i \in \{1, \dots, l-1\}$:
 $A(e_{(i)}) \succ_{e_{(i+1)}} A(e_{(i+1)})$ and $A(e_l) \succ_{e_1} A(e_1)$

We denote by C_k the class of deals involving at most k agents. In this paper, we focus on bilateral deals: C_2 (or *swaps*), but also use C_3 and C_n for comparison purpose.

Definition 2. An allocation is *k-stable* when there is no C_k rational deals possible.

We study dynamics consisting of allowing sequences of rational C_k deals until a *k-stable* allocation is reached. By extension, and when clear from the context, we call C_k such dynamics. Note that in a sequence, the same item can be successively held by different agents. From the same initial allocation, it is clear that many different sequences can occur, leading to different final allocations.

Example 2. In Example 1, the exchange (a_4a_5) for instance leads to a 2-stable allocation [52134] (in that case, also Pareto optimal). However, from the same initial allocation, the exchange (a_1a_4) is also possible, followed by (a_1a_5) , leading to the 2-stable allocation [32154].


Thus, to fully define a dynamics, we need to specify how a given deal is chosen when several are eligible. In this paper, the exact dynamics we study is often irrelevant, in the sense that the results are stated for any sequence of deals. However in experiments we need to make a specific choice. We use a *random dynamics*: pairs (resp. pairs or triple) of agents are selected uniformly at random, and agents check whether they can make a C_2 (resp. C_3) rational deal. This is arguably the least committing dynamics.

2.5 Top Trading Cycle

When each agent initially holds a resource, the top-trading cycle procedure [20] is *Pareto-efficient*, *Individually Rational*, and *Strategy Proof* (no agent has an incentive to misrepresent her preferences). In fact, it is the *only* mechanism presenting such guarantees [12].

The procedure constructs a bipartite graph $G = \langle E, V \rangle$ where $V = \mathcal{N} \cup \mathcal{R}$. Now, for each agent a_i , put an edge pointing to $top(a_i)$. Reciprocally, set an edge from r_j to the agent possessing it. There must be a cycle in the resulting graph. Pick one such cycle, and implement the exchange of resources. Then remove the agents and resources involved in the cycle, and reiterate the procedure on the restricted graph, until exhaustion of the graph.

Example 3. Take the profile of Example 1. The graph initially constructed is as follows:



There are 3 cycles. In this case, TTC returns (at the end of the process) allocation [52134].

The returned allocation A_c is *unique*, it is in fact the only allocation in the core, in the sense that there is no coalition of agents $X \subseteq \mathcal{N}$ and allocation A' such that $\cup_{a_i \in X} A_0(a_i) = \cup_{a_i \in X} A'(a_i)$, and $A'(a_i) \succeq_i A_c(a_i)$ for all $a_i \in X$, with $A'(a_i) \succ_i A_c(a_i)$ for at least one of them. On our running example, observe why $A = [32154]$ (also Pareto-optimal and better than $A_c = [52134]$ in terms of utilitarian and egalitarian social welfare) is *not* the core allocation: $\{a_4, a_5\}$ constitutes a group of agents who could have been better off in A_c .

2.6 Computing Optimal Allocations

It is worth noticing that computing centrally an optimal utilitarian (or indeed egalitarian) allocation is not difficult algorithmically.

To obtain an allocation with maximal utilitarian social welfare but respecting individual rationality, one can simply translate the problem to a (weighted) matching problem in a bipartite graph. Agents are only matched to objects they prefer to their current assignment (disregard this to relax IR), with weights corresponding to their utility for each resource. This can be solved by standard techniques in $O(n^3)$.

To return an optimal egalitarian allocation, the same technique can be iterated [9]. First construct the matching problem involving only the first rank: if no perfect matching exists, augment the instance by adding edges corresponding to the second rank (with weights corresponding to utility). To make sure that IR is respected, only add edges until you have reached the current assignment for this agent. The first returned matching maximizes the minimal rank among agents and if we require it to maximize weight, it is also Pareto-optimal among the allocations maximizing the minimal rank. It runs in $O(n^4)$. This is similar in spirit to the descending demand procedure of Herreiner and Puppe [10].

3. PARETO OPTIMALITY

We first investigate the efficiency of the obtained allocations, in terms of Pareto-optimality. Of course, it is obvious that a 2-stable allocation may not be Pareto-optimal. But a natural question is to ask how often C_2 reaches a Pareto-optimal allocation. We first make the following remark: in any Pareto-optimal allocation, at least one agent has to hold her preferred resource. This provides a $O(n)$ test to disprove Pareto-optimality (just check the top rank). Fortunately, proving that an allocation is Pareto-optimal is not very difficult either: remember that TTC returns a Pareto-optimal allocation and respects IR. It then suffices to run TTC over the allocation reached by C_2 (more sophisticated approaches exist, we refer to [3] for developments on this issue).

We ran experiments to test how often the final allocation obtained with C_2 was Pareto-optimal. We vary the size of the instances and compute percentages of Pareto-optimal outcomes for each size. We report below (Fig.1) results under IC and PL. It can be seen that under IC, with $n = 10$, there is about 75% chance to reach a Pareto-optimal allocation, while with $n = 14$ there is still 50% chance to reach a Pareto-optimal allocation. From $n = 30$, it becomes almost impossible. For comparison, we note that the real world instances of PrefLib provide better results. This seems to suggest that correlation between preferences is favorable to




Figure 1: Percentage of Pareto-optimal outcomes

convergence to Pareto-optimal allocations (think of the extreme case where all agents have the same preference ordering: all allocations are stable and Pareto-optimal).

The question is whether there are natural domains guaranteeing convergence to Pareto-optimal allocations for C_2 . We will see that this is the case for single-peaked domains.

3.1 Single-peaked Domains

In our analysis of single-peaked domains, it will be convenient to make use of the notion *worst-restrictedness* [19].

Definition 3. (Worst-restrictedness, WR) A profile is worst-restricted if, for any triple of resources r_x, r_y, r_z , there always exists a resource which is not ranked last when we restrict our attention to these three resources.

It is known that *worst-restrictedness* is a necessary condition (but not sufficient) for a profile to be single-peaked [6]. Our proof shows that being stuck in an allocation not Pareto-optimal while no swap deal is possible must lead to a violation of this condition.

PROPOSITION 1. *In a single-peaked domain, any sequence of rational swap deals reaches a Pareto-optimal allocation.*

PROOF. Suppose for the sake of contradiction that no swap deal is possible, but that the obtained allocation A is not Pareto-optimal. This means that there must exist a cycle deal (μ) involving $k > 2$ resources r_1, \dots, r_k such that $r_1 \succ_2 r_2, r_2 \succ_3 r_3, \dots, r_{k-1} \succ_k r_k, r_k \succ_1 r_1$ (assuming $r_i = A(i)$ without loss of generality). If agents implemented this contract, they would get to an allocation dominating A . Note that resources not involved in μ are irrelevant in this proof, thus they will be ignored in examples and figures. We now proceed by induction on the length k of the cycle.

Base case ($k = 3$): Consider agents a_1 and a_2 : because no bilateral deal is possible, it must be the case that $r_1 \succ_1 r_2$. The same holds for all the agents involved in the cycle, yielding the preferences:

$$\begin{array}{l} 1 : r_3 \succ \boxed{r_1} \succ r_2 \\ 2 : r_1 \succ \boxed{r_2} \succ r_3 \\ 3 : r_2 \succ \boxed{r_3} \succ r_1 \end{array}$$

We note that this directly violates the WR condition, since the three resources appear in last position.

Induction step: We now assume that no cycle of length $k - 1$ is possible, and show that no cycle of length k can

occur. For illustration, for a cycle involving 4 agents we would have (with resources in brackets in no specific order):

$$\begin{array}{l} 1 : r_4 \succ \boxed{r_1} \succ \{r_2, r_3\} \\ 2 : r_1 \succ \boxed{r_2} \succ \{r_3, r_4\} \\ 3 : r_2 \succ \boxed{r_3} \succ \{r_4, r_1\} \\ 4 : r_3 \succ \boxed{r_4} \succ \{r_1, r_2\} \end{array}$$

We know, for instance, that a_1 must have $r_1 \succ_1 r_3$ otherwise the cycle $(a_1 a_2 a_3)$ of length $k - 1 = 3$ would be rational, and similarly for the other agents.

Generalizing this, it can be seen that no other resource from μ can appear between the resources held and the resources from the cycle. Now, to respect the condition WR, it must be the case that among all the resources in the cycle, at most two can be ranked last (among those of the cycle). Let us then denote r_w an arbitrary such “worst” resource, and note a_l an agent putting this resource in last position.

But now, regardless of the chosen resource r_w , we can pick agent a_w such that $A(a_w) = r_w$, and a_{w+1} , noting that $top(a_{w+1}) = r_w$. Note that $a_w \neq a_l$ and $a_{w+1} \neq a_l$.

The three agents a_l, a_w, a_{w+1} constitute a witness of impossibility for the single-peaked ordering if we focus on resources r_{w-1}, r_w, r_{w+1} . Indeed a_l puts r_w in last position, but a_w must put r_{w+1} in last position (since $top(a_w) = r_{w-1}$), and finally a_{w+1} must put r_{w-1} in last position (since $top(a_{w+1}) = r_w$). The condition WR is violated, leading to a contradiction. \square

4. UTILITARIAN

In the previous section, we have seen that, in the general case, the hope to reach a Pareto-optimal allocation fades as we approach 30 agents. This is however a very crude view of the situation: how bad are really the allocations reached?

4.1 Worst-case Analysis

A standard measure is to compute the Price of Anarchy (PoA), that is, the deficit of social welfare induced by a decentralized procedure, compared to centralized optimization. More formally, for an instance I with an initial allocation A_0 , denoting respectively $C_k(I)$ the set of k -stable allocations reachable from A_0 , we are ultimately interested in the ratio:

$$PoA = \max_{I \in \mathcal{I}} \frac{\max_{A \in ISW_u(A)} sw_u(A)}{\min_{A \in C_k(I)} sw_u(A)}$$

But we will first see that the IR constraint alone induces a significant PoA.

LEMMA 1. *Any procedure respecting IR have $PoA \geq 2$.*

PROOF. Consider the following instance and allocations:

$$\begin{array}{l} 1 : \boxed{r_1} \succ \boxed{r_n} \succ \dots \succ \dots \succ \dots \\ 2 : \boxed{r_1} \succ \boxed{r_2} \succ \dots \succ \dots \succ \dots \\ 3 : \boxed{r_2} \succ r_1 \succ \boxed{r_3} \succ \dots \succ \dots \\ 4 : \boxed{r_3} \succ r_1 \succ r_2 \succ \boxed{r_4} \succ \dots \\ \vdots \\ n : \boxed{r_{n-1}} \succ \dots \succ \dots \succ \dots \succ \boxed{r_n} \end{array}$$

Observe that any procedure respecting IR must leave the ‘white box’ allocation A unchanged if it is the initial allocation: agent 1 holds her preferred resource, agent 2 would

only swap to r_1 which is held by agent 1, and so on. The $sw_u(A) = n(n+1)/2$. Now, A' allocation ('yellow box' allocation) would provide $sw_u(A') = n^2 - 1$. This yields (asymptotically) a PoA ≥ 2 . \square

PROPOSITION 2. All C_k procedures have PoA = 2.

PROOF. Take C_2 . For an allocation A to be 2-stable, it must be that, for each pair of agents (x, y) , at least one agent ranks the resource held by the other below her current resource, otherwise a bilateral exchange would be possible. This means that, there must be overall (by all the agents) at least $n(n-1)/2$ resources ranked below the resources they currently hold, yielding a social welfare of $n(n+1)/2$, hence (asymptotically) a PoA ≤ 2 . As C_2 is a procedure respecting IR, the claim follows from Lemma 1, and holds a fortiori for any C_k procedure. \square

These results show that no procedure guaranteeing IR can provide better guarantees, and, incidentally, that the size of cycles allowed in cycle-based procedures does not make any difference (in the worst case).

At this point the reader may be confused, as there are clearly situations where a certain gap will exist between the *best* allocations which may be reached with "short" cycles, and the *worst* allocations which may be reached with cycles of arbitrary length. A different measure, which we may call "the price of short cycles", captures this notion:

$$PoSC = \max_{I \in \mathcal{I}} \frac{\min_{A \in C_n(I)} sw_u(A)}{\max_{A \in C_k(I)} sw_u(A)}$$

PROPOSITION 3. The C_2 procedure has PoSC = 2.

PROOF. Consider the following instance and allocations, with the middle ranking $m = (n+1)/2$, n being odd:

1	:	r_1	\succ	r_2	\succ	\dots	\succ	r_m	\succ	\dots	\succ	r_{n-1}	\succ	r_n
2	:	r_2	\succ	r_3	\succ	\dots	\succ	r_{m+1}	\succ	\dots	\succ	r_n	\succ	r_1
\vdots	\vdots													
n	:	r_n	\succ	r_1	\succ	\dots	\succ	r_{m-1}	\succ	\dots	\succ	r_{n-2}	\succ	r_{n-1}

The only Pareto-optimal allocation is the one in which each agent holds her preferred resource ('yellow box' allocation). Note that the 'white box' allocation is 2-stable, and yields asymptotically a social welfare equals to half of the social welfare of the Pareto-optimal allocation. By taking precisely 'white box' as the initial allocation, it is the best C_2 allocation on this instance. Thus, the C_2 procedure has a PoSC of at least 2. By the same argument used in Prop. 2, it holds that C_2 must have a PoSC of at most 2. \square

4.2 Average-case Analysis

PoA and PoSC give theoretical results in the worst-case. We ran experiments to study utilitarian social welfare in the average case. Unless stated otherwise, our experimental results regarding average social welfare are given as a ratio of the maximal sw_u , computed as explained in Section 2.6. Figures 2 and 3 present average social welfare obtained for different sizes of instances under Impartial Culture and Single Peaked preferences respectively. For each instance size, a run is an instance (including an initial allocation) on which we apply the different methods mentioned, *i.e.* for C_2 deals

are performed until a stable allocation is reached. Average values are obtained on 2000 runs.

We first report on the performance of C_2 as far as the average utilitarian social welfare is concerned. As a means of comparison, we also provide *maxRat*, that is, the optimal value which can be obtained (centrally) but still respecting IR, TTC (bearing in mind of course that this procedure is not designed to optimize utilitarian social welfare), and C_3 , to appreciate the gain induced by slightly larger cycles.




Figure 2: Mean value of sw_u under IC




Figure 3: Mean value of sw_u under UP-SP

A first global observation is that the outcomes provide fairly high values of social welfare (above 90% of the theoretically max value even disregarding IR) under IC and UP-SP. Of course this is relative to an increasing max value —so the absolute loss augments, but very moderately.

Under IC, it can also be noticed that C_3 provides no improvement over C_2 for small size instances, and then from $n = 10$ this improvement is rather small: an almost constant 1%. Under UP-SP, there is no significant difference.

The results under UP-SP also show that the obtained values of sw_u are on average higher (above 95%) than under IC. It is also important thing to notice that these experiments also support the very good behaviour of C_2 under UP-SP relatively to other procedures: for instance, we see that it slightly outperforms TTC (remember both are guaranteed to return Pareto-optimal allocations under this culture).

For completeness, we also note with the datasets sushi and tee-shirts from PrefLib, we observe respectively $\simeq 82\%$ and $\simeq 92\%$ (and slightly higher with C_3), hence a lower social welfare than under IC and a fortiori UP-SP. This may look

contradictory with the fact that the likelihood to reach a Pareto-optimal allocation is higher. However, in that case this is the role of IR which is crucial. Indeed the *maxRat* value culminates at 87% (sushi) and 97% (tee-shirts) of the optimal allocation disregarding IR.

5. EGALITARIAN

One might wonder whether C_2 favors egalitarian stable allocations. Egalitarian social welfare (sw_e) is interested in the worst served agents. As defined in Section 2.2, maximizing the egalitarian social welfare consists in maximizing the minimum of the individual utilities.

5.1 Worst-case Analysis


The PoA definition can be readily adapted to the case of egalitarian social welfare. In this case, it is easy to see that no guarantee on the gap can be given. In fact, the construction of Lemma 1 shows that any procedure respecting IR can exhibit a $sw_e(A) = 1$ and get stuck, while $sw_e(A') = n - 1$ would be possible. Since this ratio cannot be larger than n , we get that:

PROPOSITION 4. *For egalitarian social welfare, all C_k procedures have a PoA of $\Theta(n)$.*

Similarly, the PoSC can be shown to have the same ratio, by exhibiting a slightly different instance.

PROPOSITION 5. *For egalitarian social welfare, the C_2 procedure has a PoSC of $\Theta(n)$.*

PROOF. Consider the following instance and allocation:



Take the initial allocation as the ‘white box’ allocation A : it is 2-stable with $sw_e(A) = 2$, while the ‘yellow box’ allocation A' is the only Pareto-optimal, with $sw_e(A') = n$. \square

5.2 Average-case Analysis

Experiments are performed using the same experimental setting as the one described in Section 4. Single-peaked domains and impartial culture domains are considered. The main observation is that C_2 gives very good results (for instance, significantly better than TTC), especially in IC. On average, the poorest agent will receive a resource ranked in the top-third of her preferences.

The performance of TTC with respect to this notion of welfare is no surprise: as explained, the procedure implements the ‘best’ cycles and then discards the resources: this reduces a lot the range of possible cycles for other agents. Some agents may thus keep low-ranked (or even their initial) resources, leading to low individual utility. On the other hand, C_2 does not favor any agents and allows each agent to perform rational bilateral deals. Moreover, each agent can trade several times. Agents with low-ranked initial resources




Figure 4: Mean value of sw_e under IC

have then more opportunities to exchange their resource and improve their individual utilitarian welfare. Consequently, we obtain better egalitarian social welfare with C_2 or C_3 .

A second important observation is that C_3 also improves little over C_2 for this notion of welfare, but more than for utilitarian welfare (note that the scale of y -axis differs). Moreover, the gap between both procedures augments very moderately as larger sizes of instances are considered, reaching about 4% when $n = 30$.

The results under UP-SP (not pictured here for lack of space) essentially show the same trends, but with significantly lower values. C_2 and C_3 go slightly below 50% for $n = 30$, while TTC reaches a small 30% for the same size of instances. Clearly, the structure of preferences makes it difficult to satisfy equally all agents.

For completeness, with the datasets from PrefLib, the worst served agent will respectively get on average her median resource (for sushi) and the resource ranked between rank 3 and 4 (for tee-shirts). These results are comparable to those obtained under UP-SP for $n = 30$. We see that, as in the utilitarian case, in these instances the IR constraint is much more demanding due to the correlation of preferences.

6. REACHABILITY

The ultimate question we address is the following: even if the procedure is to be limited to cycles of restricted length, in the presence of a central authority, an alternative approach would be to let the center plan ahead the sequences of rational swap deals. In this section we show that this may not be a viable solution, because of the complexity of this reachability problem.

More specifically, we address the problem of *Minimization of the Utilitarian Social Welfare with Swap Deals* (MIN-USW-C2) and *Minimization of the Egalitarian Social Welfare with Swap Deals* (MIN-ESW-C2).

MIN-USW-C2	
Instance:	An instance I , with initial allocation A_0 , an integer k
Question:	YES if and only if there exists a sequence of rational swap deals starting from A_0 and leading to an allocation A' such that $sw_u(A') \geq k$.

Instance: An instance I , with initial allocation A_0 , an integer k

Question: YES if and only if there exists a sequence of rational swap deals starting from A_0 and leading to an allocation A' such that $sw_e(A') \geq k$.

PROPOSITION 6. *The MIN-ESW-C2 and MIN-USW-C2 problems are NP-complete.*

PROOF. First, both problems are in NP. In fact, note that the length of any sequence of rational swap deals is $O(n^2)$. Thus, if a sequence of rational swap deals is given as a guess, the social welfare of the resulting allocation can be computed in a polynomial time, which proves NP-membership. To prove NP-hardness, we will show that both problems reduce to the hamiltonian circuit problem in directed graphs with an in-degree and an out-degree of at most 2 [14].

Let $G = (V, E)$ be a directed graph with an in-degree and an out-degree of at most 2 with $V = \{1 \dots n\}$. Assume without loss of generality that each vertex has an out-degree of at least one (if a vertex has a null out-degree, then no hamiltonian circuit exists). From G , let us build a graph G' with an additional vertex $n+1$ such that all arcs in G pointing towards vertex 1 now point towards vertex $n+1$. Formally, $G' = (V \cup \{n+1\}, \{(i, j) \in E : j \neq 1\} \cup \{(i, n+1) : (i, 1) \in E\})$.

There exists a hamiltonian circuit in G iff there exists a path in G' starting from 1, ending on $n+1$ and visiting all vertices. Let us now build the reduction.

The set of resources will be $\mathcal{R} \cup \{u\} \cup B \cup C$ where $\mathcal{R} = \{r_1 \dots r_{n+1}\}$, $B = \{b_1 \dots b_{n+1}\}$ and $C = \{c_1 \dots c_{n+2}\}$. This sums up to $M = 3n + 5$ resources (and agents). Intuitively, resource u will be used as a token moving from agent to agent: the sequence of agents visited by u corresponds to a path in G' . Resources B will be used to unlock some resources r_j which will be then exchanged against u . Finally, resources C do not play any role other than amplifying the social welfare when some r_j resource reaches the top rank.

In an abuse of notation, the symbol $allC$ in a linear order will refer to the lexicographic ordering among resources of C . For example, $r_1 \succ allC \succ r_2$ is the linear ordering $r_1 \succ c_1 \succ c_2 \dots \succ c_{n+2} \succ r_2$. Also, in each linear order, all resources ranked below the originally held resource do not play any role, so they will be omitted here.


Now we build our preference profile. For each vertex i of G' , let $P = \{p_1 \dots p_k\}$ be the set of predecessors of i in G' (the size of P is at most 2). Simply add to the profile the linear order:

$$r_i \succ allC \succ u \succ r_{p_1} \succ \dots \succ r_{p_k} \succ \boxed{b_i}$$

Then, add $u \succ \boxed{r_{n+1}}$ and $b_1 \succ \boxed{u}$.

Now for each $i \in \{1 \dots n\}$, add $b_{s_1} \succ \dots \succ b_{s_l} \succ \boxed{r_i}$ where $S = \{s_1 \dots s_l\}$ is the set of successors of i in G' . Finally, for each $i \in \{1 \dots n+2\}$, add $\boxed{c_i}$.

To make this construction easier to grasp, consider the graph $G = (\{1, 2, 3\}, \{(1, 2), (2, 3), (3, 1), (2, 1)\})$. From G , we build the graph G' shown below:




The profile (and initial 'box' allocation) build from G' is the following:

1	: r_1	\succ	$allC$	\succ	u	\succ	$\boxed{b_1}$	
2	: r_2	\succ	$allC$	\succ	u	\succ	r_1	\succ
3	: r_3	\succ	$allC$	\succ	u	\succ	r_2	\succ
4	: r_4	\succ	$allC$	\succ	u	\succ	r_2	\succ
5	: u	\succ	$\boxed{r_4}$					
6	: b_1	\succ	\boxed{u}					
7	: b_2	\succ	$\boxed{r_1}$					
8	: b_3	\succ	b_4	\succ	$\boxed{r_2}$			
9	: b_4	\succ	$\boxed{r_3}$					
10	: $\boxed{c_1}$							
	:							
14	: $\boxed{c_5}$							

Each sequence of swap deals induces a path in G' . Let us see how to extract such a path. Consider *any* sequence of swap deals. Assume the t^{th} deal involves exchanging resource r_i with resource u , for some $i > 1$. Consider the subsequence of deals ranging from the first deal to the $(t-1)^{th}$. Take the last deal in this subsequence involving u and some other resource r_j . Then, $(j, i) \in E$. Note in addition that a deal (r_i, u) cannot appear twice in a sequence of rational deals. Let $r_{j_1} \dots r_{j_k}$ be the list of resources exchanged with u in the sequence of deals, in this order. Then, $j_1 \dots j_k$ is a path in G' .

On our example, there are two possible sequences of swap deals (modulo possible permutations in these sequence). The longest one is depicted below (each vertex is an agent, and arcs show which resources are exchanged between agents). We clearly see that resource u is exchanged against resources r_1, r_2, r_3, r_4 (in that order). So this sequence induces the path 1, 2, 3, 4 in G' .



Assume now that there exists a path $i_1, i_2 \dots i_{n+1}$ passing through each vertex once in G' and such that $i_1 = 1$ and $i_{n+1} = n+1$. Then the following sequence of swap deals is rational:

$$(u, b_1), (r_{i_1}, b_{i_2}), (r_{i_1}, u), (r_{i_2}, b_{i_3}), (r_{i_2}, u), (r_{i_3}, b_{i_4}), \dots, (r_{n+1}, u)$$

Moreover, this sequence leads to a final allocation where the first and the last $n+2$ agents have their preferred resources, and the remaining $n+1$ have their best or second best resource. The final social welfares will thus be bounded as follows:

$$\begin{aligned} sw_u &\geq 2M \cdot (n+2) + (M-1)(n+1) \\ &= 9n^2 + 29n + 24 \\ sw_e &\geq M-1 = 3n+4 \end{aligned}$$

Conversely, assume that no such path exists. Then, no sequence of IR deals can contain all deals involving r_i, u for all $i \in \{1 \dots n+1\}$. Thus, at least one of the $n+1$ first agents will have a resource positioned below the rank $n+2$. Thus, the social welfares will be:

$$\begin{aligned} sw_u &\leq M^2 - (n+2) \\ &= 9n^2 + 29n + 23 \\ sw_e &\leq M - (n+2) = 2n+3 \end{aligned}$$

To summarize, there exists a sequence of rational swap deals leading to a utilitarian social welfare of at least $9n^2 + 29n + 24$ and an egalitarian social welfare of at least $3n+4$ if and only if G admits an hamiltonian circuit. This concludes the NP-hardness proof. \square

7. RELATED WORK AND CONCLUSION

While TTC is a celebrated method to allocate indivisible items among agents without money balance, it is difficult to apply in distributed settings. In this paper, we investigated swap deals for distributed resource allocation in house market settings, under a very simple dynamics that allows the agents to improve their satisfaction without requiring complex coordination. We provide several new insights to assess the power of this approach in such settings. Pareto optimality is thus guaranteed under single-peaked preferences. In a larger context, we showed that no other IR mechanism (even those involving longer cycles) can ever provide better guarantee about utilitarian social welfare loss in the worst case. While the ‘price of short cycles’ may be high in principle, our experimental findings show that in average, performances are fairly good in terms of social welfare. Finally, we proved the NP-completeness of deciding whether an allocation maximizing utilitarian or egalitarian social welfare can be reached. This complexity results reinforce the case for such approaches based on local myopic improvements.

Being based on cycle reallocation of resources, it is useful to recall that a given permutation can be decomposed in products of transpositions (cycles of length 2). This means that we could always decompose a permutation (hence reach any allocation), if it was not for the rationality constraint. Our dynamics based on local improving requires indeed each swap to be rational. Such results could however prove useful to identify easily verifiable sufficient conditions for the reachability problem (which, we saw, is difficult in general).

Recently, [11] studied a domain where the single-peaked restriction is natural: assigning agents to slots on a line, where each agent prefers to be allocated as close as possible from her preferred slot (actually, this induces by default a more restrictive domain since the notion of distance is symmetric). Interestingly, they prove a convergence result

implying Pareto-optimality, but under a different rationality criteria (“aggregate gap-reduction”).

Similar issues occur in related but different problems.

The kidney exchange problem has recently stimulated a large body of work. This is a case where the prospects of failures induced by long exchange cycles cannot be accepted, since they may have dramatic consequences. As a result, most programs restrict their implementation to pairwise exchanges, sometimes cycles of length 3 [16]. The approach is of course centralized. In this context, the clearing problem (consisting in finding a collection of disjoint cycles maximizing social welfare) is known to be NP-complete [2] when the bound on cycle length is ≥ 3 . Here we saw that for cycles of length 2, the *reachability* problem in our setting is hard already. On the other hand, the reader should keep in mind that such ‘planning-like’ problems have typically high complexity (in our case we remain in NP because the sequences of cycles are polynomially bounded). An open question is whether hardness also holds for reachability of Pareto-optimal allocations. We strongly suspect this to be the case, but could not prove it so far.

We may also point the reader to studies of online exchange markets (allowing users to exchange books, dvd, etc.), where the stakes are certainly less life threatening. In such settings the dynamics is also challenging: preferences, *i.e.* “wish lists” are not necessarily complete and may get updated. As a result, agents may have to stall their activity for a while, perhaps discarding other rational proposals. Abassi et al. [1] observed this in datasets from existing barter exchange models (“in a dynamic network where item and wish lists may get updated, it may take a long time before such exchange opportunities materialize”).

The kind of uncoordinated approaches studied in this paper have been investigated in the context of two-sided matching. Roth and Vade Vate [17] show that better response dynamics will always convergence in expectation, but Ackermann et al. [4] exhibit an exponential lower bound on the sequence length. Even closer in spirit to our approach is the work of [5], which investigates the price of anarchy and price of stability of stable matchings, providing both theoretical bounds and experimental findings. Even though their notion of stability is different, they notice like us a big discrepancy between the worst-case predictions and the experimental findings about the quality of matchings.

Bilateral deals have been investigated in a context of distributed resource allocation involving *payments* among agents [7], and in a context where agents may hold several resources simultaneously, and have arbitrary preferences about such bundles of resources. Convergence in the domain of additive preferences can be proven, and bounds on the length of sequences can be derived [8]. One result is that no larger domain than additive preferences can guarantee convergence. A similar question could be asked for our result in single-peaked domains: is there a larger (ideally, natural) domain still guaranteeing convergence?

Acknowledgments

We would like to thank Sylvain Bouveret, Anisse Ismaili, Julien Lesca and the anonymous reviewers, for their very helpful feedback on this paper. This work is partially supported by the ANR project 14-CE24-0007-01 - CoCoRiCo-CoDec.

REFERENCES

- [1] Z. Abbassi, L. V. Lakshmanan, and M. Xie. Fair recommendations for online barter exchange networks. In *Proc. 16th International Workshop on the Web and Databases (WebDB)*, 2013.
- [2] D. J. Abraham, A. Blum, and T. Sandholm. Clearing algorithms for barter exchange markets: Enabling nationwide kidney exchanges. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07*, pages 295–304, New York, NY, USA, 2007. ACM.
- [3] D. J. Abraham, K. Cechlárová, D. F. Manlove, and K. Mehlhorn. Pareto optimality in house allocation problems. In R. Fleischer and G. Trippen, editors, *Algorithms and Computation*, volume 3341 of *Lecture Notes in Computer Science*, pages 3–15. Springer Berlin Heidelberg, 2005.
- [4] H. Ackermann, P. W. Goldberg, V. S. Mirrokni, H. Röglin, and B. Vöcking. Uncoordinated two-sided matching markets. *SIAM Journal on Computing*, 40(1):92–106, July 2011.
- [5] E. Anshelevich, S. Das, and Y. Naamad. Anarchy, stability, and utopia: Creating better matchings. *Journal of Autonomous Agents and Multi-Agent Systems*, 26(1):120–140, 2013.
- [6] M. A. Ballester and G. Haeringer. A characterization of the single-peaked domain. *Social Choice and Welfare*, 36(2):305–322, 2011.
- [7] 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*, 20(2):234–259, 2010.
- [8] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research*, 23:41–78, 2005.
- [9] R. S. Garfinkel. An improved algorithm for the bottleneck assignment problem. *Operations Research*, pages 1747–1751, 1971.
- [10] D. Herreiner and C. Puppe. A simple procedure for finding equitable allocations of indivisible goods. *Social Choice and Welfare*, pages 415–430, 2002.
- [11] J. L. Hougaard, J. D. Moreno-Ternero, and L. P. Østerdal. Assigning agents to a line. *Games and Economic Behavior*, 87:539–553, 2014.
- [12] J. Ma. Strategy-proofness and the strict core in a market with indivisibilities. *International Journal of Game Theory*, 23:75–83, 1994.
- [13] N. Mattei and T. Walsh. Preflib: A library of preference data. In *Proceedings of Third International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer, 2013.
- [14] J. Plesnik. The NP-completeness of the Hamiltonian cycle problem in planar digraphs with degree bound two. *SIAM Journal on Computing*, 8(4):199–201, July 1979.
- [15] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [16] A. Roth, T. Sonmez, and U. Unver. Pairwise kidney exchange. *Journal of Economic Theory*, 125(2):151–188, 2005.
- [17] A. E. Roth and J. H. V. Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, 1990.
- [18] T. W. Sandholm. Contract types for satisficing task allocation: I Theoretical results. In *Proc. AAAI Spring Symposium: Satisficing Models*, 1998.
- [19] A. K. Sen. A possibility theorem on majority decisions. *Econometrica*, 34(2):491–499, 1966.
- [20] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of mathematical economics*, 1(1):23–37, 1974.