Open access • Book Chapter • DOI:10.1007/978-3-319-99660-8_19

Constrained Swap Dynamics over a Social Network in Distributed Resource Reallocation — Source link

Abdallah Saffidine, Anaëlle Wilczynski

Institutions: Australian National University, Paris Dauphine University

Published on: 11 Sep 2018 - Algorithmic Game Theory

Topics: Resource allocation

Constrained Swap Dynamics over a Social Network in Distributed Resource Reallocation

Abdallah Saffidine[1] and Anaëlle Wilczynski[2](✉)

[1] College of Engineering & Computer Science,
Australian National University, Canberra, Australia
abdallah.saffidine@gmail.com
[2] PSL, CNRS, LAMSADE, Université Paris-Dauphine, Paris, France
anaelle.wilczynski@dauphine.fr

Abstract. We examine a resource allocation problem where each agent is to be assigned exactly one object. Agents are initially endowed with a resource that they can swap with one another. However, not all exchanges are plausible: we represent required connections between agents with a social network. Agents may only perform pairwise exchanges with their neighbors and only if it brings them preferred objects. We analyze this distributed process through two dual questions. *Could an agent obtain a certain object if the swaps occurred favourably? Can an agent be guaranteed a certain level of satisfaction regardless of the actual exchanges?* These questions are investigated through parameterized complexity, focusing on budget constraints such as the number of exchanges an agent may be involved in or the total duration of the process.

Keywords: Resource allocation · Distributed process
Social network · Parameterized complexity

1 Introduction

Reallocating resources among agents is a central question widely studied both in computer science and economics [1,3,10,14]. This problem refers to a particular setting of resource allocation, where the agents are initially endowed with items [6]. Resource reallocation models many real-life situations, like reallocating tasks between employees or reassigning time slots in schedules. In such examples, agents are often assigned a single task. With exactly one item per agent, the problem is known as *housing market* [1,19]. In this context, a central authority may decide how to redistribute the objects [4,19]. Alternatively, the agents may direct the reallocation by trading and negotiating among them in a distributed process [7–9]. Although largely studied in general resource reallocation [11,12,14,18], this approach has only recently been introduced in housing markets [10].

X. Deng (Ed.): SAGT 2018, LNCS 11059, pp. 213–225, 2018.
https://doi.org/10.1007/978-3-319-99660-8_19

In a distributed process, natural obstacles may inhibit the agents in the trades. Lack of trust may lead agents to adopt a greedy behavior so as to be immediately better off in their new acquisition. Logistics difficulties, e.g., communication and geographical distance, may also prevent some trades to occur. This can be modeled by restricting trades to the links of a social network [13,17], with exchanges limited to the cliques [9] or edges [16] of the graph.

We consider this latter format: a housing market with exchanges between neighbors in a social network [16]. One of the main questions is the Reachable Object (RO) problem: Given a target agent A and a target object x, is there a sequence of exchanges ensuring A is eventually allocated x? This problem is very appealing but is NP-complete even when the network is a tree [16]. We attempt to mitigate this negative result by looking at more realistic constrained settings.

We draw inspiration from the fact that an agent may not be willing to perform a large number of swaps or to wait a long time before getting the desired object. We introduce natural budget constraints, the number of exchanges agents may make and the total duration of the process, and we perform a refined complexity analysis. Moreover, we introduce Guaranteed Level of Satisfaction (GLS), a problem related to RO but more realistic. GLS asks whether an agent can be guaranteed to be eventually allocated an item at least as good as the input target item, regardless of the exchanges other agents perform, provided they are rational swaps. While RO takes an optimistic perspective, GLS adopts a more pessimistic point of view beyond "lucky" exchange sequences. These problems naturally arise when we analyze the distributed process of exchanges in reallocation but they can also model concrete issues. As an example, consider an online exchange platform where users input in the system which item they hold as well as their preference. A user may request a target object to the centralized system which would then suggest a series of intermediate swaps to bring it to her. Even in such a context, restricted rational exchanges are relevant: geographical constraints can still prevent two agents to trade and the guarantee of getting a better object is essential as otherwise an agent could be left worse off than she started, should an intermediate agent exit the system during the process.

When parameterizing the problems by the maximal number of swaps per agent, we show intractability even for highly structured graphs. However, when constraining the duration of the process, we obtain more promising results: RO and GLS are tractable in a very relevant class of networks, namely bounded degree graphs, and in general the problems are tractable when the duration does not depend on the input size. These results contrast strongly with both previous work [16] and our first parameterization, and they focus on realistic scenarios: actual social networks have indeed bounded degrees and the time that an agent is willing to wait for a target object is independent from the input size.

We start with the swap dynamics model, RO and GLS, and some complexity background. Section 3 relates RO and GLS. Our results for the max-swaps and for the total-duration parameters are in Sects. 4 and 5 respectively.

2 Formal Framework

2.1 Swap Dynamics Model

Let N be a set of n agents, and M a set of n resources (or objects). Each agent $i \in N$ has ordinal strict preferences $\succ_i$ over the objects, i.e., $a \succ_i b$ means that agent i prefers object a to object b. An allocation σ is a bijection $\sigma : N \rightarrow M$, assigning to each agent exactly one object. The object assigned to agent i in σ is denoted by σ_i. Each agent is initially endowed with an object, and we denote by σ^0 the initial allocation. The agents are embedded in a social network, represented by an undirected graph $G = (N, E)$. An instance of swap dynamics model is a tuple $\mathcal{I} = (N, M, \succ, \sigma^0, G)$.

Agents can exchange their objects so as to obtain better objects, but not all exchanges are plausible. The possibilities depend on the social network and on the preferences of the agents. We only admit *swaps*, rational trades between neighbors. Formally, a swap in an allocation σ is a trade between two adjacent agents $(i, j) \in E$ such that the exchange is rational, i.e., $\sigma_i \succ_j \sigma_j$ and $\sigma_j \succ_i \sigma_i$.

A sequence of swaps is a sequence of allocations $(\sigma^0, \ldots, \sigma^t)$ such that a swap is performed between two consecutive allocations σ^i and σ^{i+1}. An allocation σ is *reachable* if there is a sequence of swaps leading to it, i.e., there exists a sequence $(\sigma^0, \ldots, \sigma^t)$ such that $\sigma^t = \sigma$. An allocation σ is *stable* if no swap is possible from σ. An object x is reachable for agent i if there is a sequence of swaps $(\sigma^0, \ldots, \sigma^t)$ where $\sigma^t_i = x$. *Swap dynamics* refers to a distributed process where agents may rationally exchange their objects when they are neighbors in the network, until a stable allocation is reached.

Example 1. Consider an instance where $n = 4$ with the following social network and preferences. The framed objects represent the initial object of each agent.

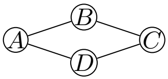

$A : b \succ c \succ \boxed{a} \succ d$ $\quad$ $C : d \succ a \succ b \succ \boxed{c}$

$B : c \succ a \succ \boxed{b} \succ d$ $\quad$ $D : a \succ b \succ \boxed{d} \succ c$

Initially, only the swaps between agents A and B, and B and C are possible. The rational swap (A, C) is not possible because the agents are not adjacent. The swap (A, D) is not possible because it is not rational for A. The sequence of exchanges (A, B), (B, C), and (C, D) gives rise to a reachable allocation where every agent gets her best object. This is stable: no further swap can be performed.

2.2 Questions

We investigate swap dynamics by analyzing two natural decision problems.

Reachable Object (RO):

Instance: $\mathcal{I} = (N, M, \succ, \sigma^0, G)$, $A \in N$, $x \in M$.

Question: Is there a sequence of swaps $(\sigma^0, \ldots, \sigma^t)$ such that $\sigma^t_A = x$?

Guaranteed Level of Satisfaction (GLS):
Instance: $\mathcal{I} = (N, M, \succ, \sigma^0, G)$, $A \in N$, $y \in M$.
Question: *For all* sequences of swaps $(\sigma^0, \ldots, \sigma^t)$ where σ^t is stable, does it hold that either $\sigma^t_A = y$ or $\sigma^t_A \succ_A y$?

When asking whether an agent can obtain some object, a large number of swaps may not be realistic. We thus study three variants of RO and GLS, referred to as {RO/GLS}-{max/sum/makespan}, where the quantity of swaps in a solution sequence is limited. In each variant, this quantity is measured differently, leading to different complexity-theoretic characterizations of the problem.

- *max*: Every agent is involved in no more than k swaps.
- *sum*: The total length of the sequence is no more than k.
- *makespan*: The makespan of the sequence is no more than k.

The *makespan* of a sequence of swaps is the minimum time that elapses from the beginning to the end, when we allow parallel swaps. This notion can be formalized as follows. Let $s = (\sigma^0, \ldots, \sigma^t)$ be a sequence of swaps. A parallel decomposition of s is a tuple of integers $\ell = (\ell_0, \ell_1, \ldots, \ell_m)$ of length $|\ell| = m$, such that $0 = \ell_0 < \ell_1 < \cdots < \ell_m = t$, and for all $0 \leq i < m$ the swaps between allocation σ^{ℓ_i} and allocation $\sigma^{\ell_{i+1}}$ do not involve the same agents. In other words, the swaps between σ^{ℓ_i} and $\sigma^{\ell_{i+1}}$ can be performed simultaneously. The makespan is the length m of the shortest parallel decomposition. Observe that the makespan of a sequence can be computed in linear time, and that the *sum* parameter is a worst case bound in case no parallel swaps take place.

2.3 Parameterized Complexity

Parameterized complexity aims at solving hard problems in time $f(k)n^{O(1)}$, called FPT time (fixed-parameter tractable), where n is the size of the instance, f is a computable function, and k is a *parameter* of the problem. Assuming the problem we are trying to solve is NP-hard, function f has to be superpolynomial, unless $\mathsf{P} = \mathsf{NP}$. However, if our parameter k is *small* compared to the size of the instance n, we achieve that the blow-up is limited to the small value k.

Some problems are highly suspected not to admit any algorithm in time $f(k)n^{\mathcal{O}(1)}$ for any computable function f, and thus not to be FPT. There are hierarchies of complexity classes beyond FPT : $\mathsf{W[1]} \subseteq \mathsf{W[2]} \subseteq \cdots \subseteq \mathsf{W[SAT]} \subseteq \ldots$, and $\mathsf{A[1]} \subseteq \mathsf{A[2]} \subseteq \cdots \subseteq \mathsf{AW[SAT]} \subseteq \ldots \subseteq \mathsf{XP}$, where $\mathsf{W[1]} = \mathsf{A[1]}, \mathsf{W[t]} \subseteq \mathsf{A[t]}$ for any $\mathsf{t} > 1$, and $\mathsf{W[SAT]} \subseteq \mathsf{AW[SAT]}$. For instance, $\mathsf{W[1]}$ is the class of parameterized decision problems that can be solved by a nondeterministic single-tape Turing machine within k steps, and XP is the class of decision problems solvable in time $\mathcal{O}(n^{f(k)})$ for some computable function f.

As a rudimentary informal intuition, FPT and $\mathsf{W[1]}$ can be thought of as corresponding to P and NP in the parameterized world. For instance, CLIQUE, the problem of finding a clique of size k in a graph, is $\mathsf{W[1]}$-complete for FPT reductions, where an FPT reduction may blow-up the instance size n only polynomially but the new parameter can be any computable function of the old parameter.

2.4 First-Order Logic

A vocabulary τ is a finite set of relation symbols. A finite structure $\mathcal{A}$ over τ consists of a finite set A, called the universe, and for each R in τ a relation over A. We assume a countably infinite set of variables. Atoms over vocabulary τ are of the form $x_1 = x_2$ or $R(x_1, \ldots, x_k)$ where $R \in \tau$ and $x_1, \ldots, x_k$ are variables. First-order (FO) formulas over τ are built from atoms over τ using standard Boolean connectives $\neg, \wedge, \vee$ and from quantifiers $\exists, \forall$ followed by a variable. Let φ be an FO formula. The variables of φ that are not in scope of a quantifier are its *free variables*. Let $\varphi(\mathcal{A})$ be the set of all assignments of elements of A to φ's free variables, such that φ is satisfied. $\mathcal{A}$ is a *model* of φ if $\varphi(\mathcal{A})$ is not empty. The class Σ_1 (resp. Σ_2) contains all FO formulas of the form $\exists x_1, \ldots, \exists x_k \varphi$ (resp. $\exists x_1, \ldots, \exists x_k \forall y_1, \ldots, \forall y_k \varphi$) where φ is a quantifier free FO formula.

Let Φ be a class of formulas. The model checking problem inputs a finite structure $\mathcal{A}$ and a formula $\varphi \in \Phi$ and asks whether the formula satisfies the model, $\varphi(\mathcal{A}) \neq \emptyset$. A natural parameter is the size of (a reasonable encoding of) φ. We will use one result bridging model checking and parameterized complexity.

Theorem 1. [15] *Model checking the existential fragment of first-order logic,* $\mathrm{MC}(\Sigma_1)$*, is* W[1]*-complete. Model checking the second level of the hierarchy,* $\mathrm{MC}(\Sigma_2)$*, is* A[2]*-complete.*

3 Relation Between RO and GLS

Reachable Object (RO) asks whether an agent A can obtain an object x by a sequence of swaps. RO is known to be NP-complete even for trees [16]. Guaranteed Level of Satisfaction (GLS) asks whether agent A is guaranteed to obtain object y or an object preferred to y in any stable reachable allocation. GLS is even more natural than RO since it offers guarantees for the agent and does not only focus on lucky configurations. It is close to the complementary of RO, and thus the study of RO also contributes to the understanding of GLS.

Proposition 2. *co-*RO *is linearly reducible to* GLS*.*

Proof (sketch). We reduce from the co-RO problem asking whether object x is unreachable for agent A. Let $\mathcal{I} = \{(N, M, \succ, \sigma^0, G), A, x\}$ be an instance of co-RO. An instance $\mathcal{I}' = \{(N', M', \succ', \sigma^{0\prime}, G'), A, y\}$ of GLS is constructed by adding an agent Y and an object y. The initial allocation $\sigma^{0\prime}$ is the same as σ^0 for all agents in N and assigns y to Y. The network G' has the same structure as G, with one more edge (Y, A). Denote by a the object of agent A in σ^0. If A prefers x to a, then denote by P_a the set containing a and the objects that are preferred to a and less preferred than x in $\succ_A$. Otherwise, P_a contains a and the objects that A prefers to a. Denote by P_x the set of objects that A prefers to x. Ranking $\succ'_A$ is constructed from $\succ_A$, by moving the objects in P_x to the end of $\succ'_A$, and by putting y at the top of $\succ'_A$. The agents in $N \setminus \{A\}$ keep the same preferences as in $\succ$ but rank y last, and Y only prefers the objects of P_a to y.

We claim that x is not reachable for A in $\mathcal{I}$ iff A obtains y or an object preferred to y in any reachable stable allocation in $\mathcal{I}'$, i.e., iff there is no stable reachable allocation where A gets z such that $y \succ_A z$. This part is omitted. □

Since RO is NP-complete even on trees [16] and the previous reduction adds only one agent and possibly one swap to an instance of RO, GLS is co-NP-hard on trees. Observe that GLS is in co-NP: after guessing a reachable stable allocation, checking whether agent A owns y or an object preferred to y is direct.

Corollary 3. GLS *is* co-NP *-complete even for trees.*

We refine the complexity of the problems using natural parameters: the number of swaps per agent and the length of the sequence. Although RO and GLS are close to be dual problems, they are indeed not complementary. GLS focuses on stable allocations. A k-bound on the sequence of swaps introduces a dependency on k on the notion of stability: a stable allocation is either stable in the standard meaning, or is reached after k swaps. Stability is not necessary in RO because for an assignment solution σ where agent A gets object x, all the stable allocations reachable from σ assign to A an object preferred to x or x itself.

4 Maximum Number of Swaps per Agent

Consider that the agents are not willing to perform an important number of swaps in the whole swap process. Surprisingly in this context, our two problems, RO and GLS, remain difficult even for a very small maximum number of swaps.

Theorem 4. *For fixed $k \geq 2$,* RO*-max is* NP*-complete, even on degree 4 graphs.*

Proof. Membership in NP is straightforward, as it is a special case of the unconstrained RO problem, known to be in NP.

For hardness, we start with $k = 2$ and reduce from (3, B2)-SAT — the restriction of SAT to instances where each clause contains three literals and each variable occurs exactly twice as a positive literal and twice as a negative literal. This variant of the propositional satisfiability problem is NP-complete [5].

We are given an instance of (3, B2)-SAT with n variables $\{x_1, \ldots, x_n\}$ and m clauses $\{C_1, \ldots, C_m\}$. We create a literal-agent Y_j^ℓ (resp., $\overline{Y_j^\ell}$) for each ℓ^{th} ($\ell \in \{1, 2\}$) occurrence of literal x_j (resp., $\overline{x_j}$), and a variable-agent Y_j for each variable x_j. Two clause-agents K_i and K_i' are created for each $0 < i < m$. Three other agents Y_0, K_0' and K_m are added. Each agent initially owns an object denoted by the lower-case version of her name, e.g., agent K_i gets object k_i.

In the network, we have the paths $[Y_{j-1}, Y_j^1, Y_j^2, Y_j]$ and $[Y_{j-1}, \overline{Y_j^1}, \overline{Y_j^2}, Y_j]$ for each $1 \leq j \leq n$, and the edge (K_i, K_i') for each $1 \leq i < m$. If the ℓ^{th} literal x_j (resp., $\overline{x_j}$) belongs to clause C_i, then we have the path $[K_{i-1}', Y_j^\ell, K_i]$ (resp., $[K_{i-1}', \overline{Y_j^\ell}, K_i]$). We connect K_m and Y_n. See for an example Fig. 1.

The preferences of the agents are given below. Notation $\{\ell_i\}$ stands for the literal-objects of clause C_i ranked in arbitrary order and $k(x_i^j)$ (resp. $k(\overline{x}_i^j)$)

for the object related to the clause in which the j^{th} occurrence of x_i (resp. $\overline{x_i}$) appears. The objects that are not mentioned in the preferences are ranked in arbitrary order after the initial endowment.

$$\begin{array}{ll}
K'_0: \ \{\ell_1\} \succ \boxed{b} \succ [\ldots] & \\
K_i: \ k'_i \succ b \succ \boxed{k_i} \succ [\ldots] & Y_j^1: k(x_j^1) \succ b \succ y_j^2 \succ a \succ \boxed{y_j^1} \succ [\ldots] \\
K'_i: \ \{\ell_{i+1}\} \succ b \succ \boxed{k'_i} \succ [\ldots] & Y_j^2: k(x_j^2) \succ b \succ y_j \succ a \succ \boxed{y_j^2} \succ [\ldots] \\
K_m: a \succ b \succ \boxed{k_m} \succ [\ldots] & \overline{Y}_j^1: k(\overline{x}_j^1) \succ b \succ \overline{y}_j^2 \succ a \succ \boxed{\overline{y}_j^1} \succ [\ldots] \\
Y_0: \ \overline{y}_1^1 \succ y_1^1 \succ \boxed{a} \succ [\ldots] & \overline{Y}_j^2: k(\overline{x}_j^2) \succ b \succ y_j \succ a \succ \boxed{\overline{y}_j^2} \succ [\ldots] \\
Y_n: \ b \succ a \succ \boxed{y_n} \succ [\ldots] & Y_j: \ \overline{y}_{j+1}^1 \succ y_{j+1}^1 \succ a \succ \boxed{y_j} \succ [\ldots]
\end{array}$$

We claim that all clauses are satisfiable iff object b reaches agent Y_n. The only way for Y_n to get hold of b is by swapping a with K_m. Object b can only reach K_m via clause-agents and literal-agents, while a can only reach Y_n via variable-agents and literal-agents. Agents perform at most two swaps, so no literal-agents can be involved in the move of both a and b.

Suppose that truth assignment ϕ satisfies all clauses. Let T_i be a literal-agent of clause C_i related to a true literal in ϕ. Since all clauses are satisfiable, object b can reach K_m via the path $[K'_0, T_1, K_1, K'_1, T_2, \ldots, T_{m-1}, K_{m-1}, K'_{m-1}, T_m, K_m]$. For variable x_j, let Z_j^1 and Z_j^2 be the literal-agents related to the literal of x_j that is false in ϕ. Clearly, these agents are not an agent T_i. It suffices for a to reach Y_n via the path $[Y_0, Z_1^1, Z_1^2, Y_1, \ldots, Y_{n-1}, Z_n^1, Z_n^2, Y_n]$.

Suppose now that object b is reachable for agent Y_n. By construction, the path of b to K_m goes through exactly one literal-agent per clause, while the path of a to Y_n goes through exactly two literal-agents associated with the same literal for each variable. Thus, the truth assignment of variables that sets to true the literals related to literal-agents in the path of object b, satisfies all the clauses.

If $k > 2$, we adapt the reduction via a delay gadget added to each agent. □

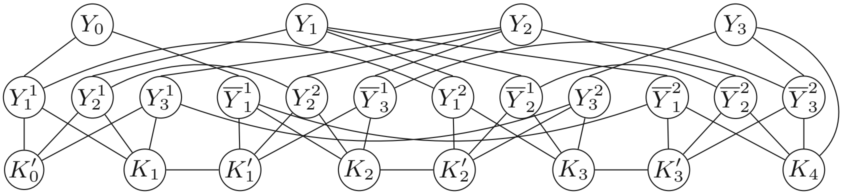

Fig. 1. Graph construction for an instance of (3, B2)-SAT with four clauses where $C_1 = (x_1 \vee x_2 \vee x_3)$, $C_2 = (\overline{x_1} \vee x_2 \vee \overline{x_3})$, $C_3 = (x_1 \vee \overline{x_2} \vee x_3)$, and $C_4 = (\overline{x_1} \vee \overline{x_2} \vee \overline{x_3})$.

From Proposition 2 and its proof, the same hardness exists for GLS, with an additional swap and one more neighbor for agent A who must obtain the object.

Corollary 5. *For $k \geq 3$,* GLS*-max is* co-NP*-complete, even on degree 5 graphs.*

One could think that the problem is easier when the structure of the network is restricted to trees. Yet, it is possible to prove that RO-max on trees is W[SAT]-hard. We leave out the lengthy formal proof but we state the main idea. We reduce from Monotone Weighted Satisfiability [2]. *Can an input propositional formula φ with no negations be satisfied with a truth assignment of weight k?* We build an instance of RO-max with a graph based on the syntax tree of φ where k chosen variable-objects must move to the occurrences of their corresponding variable as a prerequisite to given object x reaching given agent A. Since the variable-objects make up almost all the swaps, $O(k)$ swaps per agent suffice.

Globally, the problems remain difficult in very simple graphs even when the number of swaps per agent is limited. Fortunately, the parameters on the length of the sequence let us circumvent the general difficulty of the two problems.

5 Length of the Sequence of Swaps

Two parameters are used to bound the length of the sequence of swaps: the total number of exchanges and the makespan. Contrary to the previous parameter, they lead to circumscribe the problems into parameterized complexity classes that are not so high in the hierarchy, allowing tractability results when the parameters are bounded by a constant. Moreover, for bounded degree graphs, relevant in the context of social networks, we obtain fixed parameter tractability.

Theorem 6. RO-*sum and* RO-*makespan are* W[1]-*hard even for trees.*

Proof (sketch). We perform a reduction from CLIQUE, the problem of deciding whether there exists a clique of size k in a graph $G = (V, E)$ such that $V = \{1, ..., n\}$ and $|E| = m$. Assume that each edge in E is written (v, w) such that $v < w$, and consider the lexicographical order over E. Let us denote by e_i^1 and e_i^2 the first and second vertex of the i^{th} edge. Let d_v be the degree of vertex v and $\delta_v(d)$, for $1 \leq d \leq d_v$, the d^{th} edge incident to v. We construct an instance $\mathcal{I}'$ of RO (see Fig. 2 for an example) by creating:

- two connected agents X and Y, and two vertex-agents U_v^{vw} and U_w^{vw} for each edge $(v, w) \in E$, connected via a path $[Y, U_v^{vw}, U_w^{vw}]$.
- agents T and T^ℓ, for $1 \leq \ell \leq k$, representing the k vertices of the clique that we must choose. They are connected via a path to Y: $[Y, T^1, \ldots, T^k, T]$.
- agents A_v and A_v^ℓ, for $v \in V$ and $1 \leq \ell < k$, representing the choice of the $k-1$ edges of the clique that are incident to v if v belongs to the clique. They are connected via a path to Y for each v: $[Y, A_v^1, \ldots, A_v^{k-1}, A_v]$.
- agents $T^{\ell *}$ adjacent to T^ℓ, for $1 \leq \ell \leq k$, and agents $A_v^{\ell *}$ adjacent to A_v^ℓ, for $1 \leq \ell < k$ and $v \in V$. They are used to "validate" their associated agent by giving to her their initial object once they own an expected object.
- auxiliary agents used to facilitate the passage of some objects: if an agent B has a connected auxiliary agent $B^{[z]}$, then the swap with $B^{[z]}$ must precede a swap for getting an object associated with z. The auxiliary agents we use are agents $Y^{[vw]}$ corresponding to edge (v, w) and connected to Y, agents $Y^{[v]}$

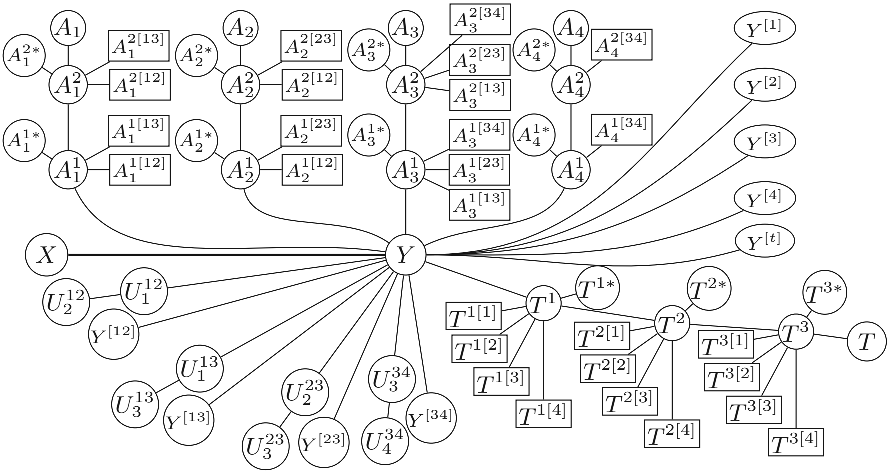

Fig. 2. Graph construction for an instance of CLIQUE with vertices $\{V_1, V_2, V_3, V_4\}$ and $k = 3$. The edges are: $\{V_1, V_2\}$, $\{V_1, V_3\}$, $\{V_2, V_3\}$ and $\{V_3, V_4\}$.

corresponding to vertex v and connected to Y, agent $Y^{[t]}$ corresponding to object t and connected to Y, agents $A_v^{\ell[\delta_v(d)]}$ corresponding to edge $\delta_v(d)$, for $1 \leq d \leq d_i$, and connected to agent A_v^ℓ, for $v \in V$ and $1 \leq \ell < k$, and agents $T^{\ell[v]}$ corresponding to vertex v and connected to agent T^ℓ for $1 \leq \ell \leq k$.

The initial object of an agent is denoted by the lower-case version of her name, e.g., agent $Y^{[v]}$ gets object $y^{[v]}$. The preferences of the agents are as follows (objects in brackets may not exist for all indices).

$$X: \quad t \succ \boxed{x} \succ [\dots] \qquad A_v : a_v^{k-1*} \succ \boxed{a_v} \succ [\dots]$$

$$U_v^{vw}: \quad a_v^{1[vw]} \succ u_w^{vw} \succ y^{[vw]} \succ \boxed{u_v^{vw}} \succ [\dots]$$

$$A_v^{\ell[\delta_v(d)]} : (a_v^{\ell+1[\delta_v(d_v)]}) \succ \cdots \succ (a_v^{\ell+1[\delta_v(1)]}) \succ a_v^\ell \succ \boxed{a_v^{\ell[\delta_v(d)]}} \succ [\dots]$$

$$T^{\ell[v]}: \quad (t^{\ell+1[v-1]}) \succ \cdots \succ (t^{\ell+1[1]}) \succ t^\ell \succ \boxed{t^{\ell[v]}} \succ [\dots]$$

$$Y^{[v]}: \quad (t^{1[v-1]}) \succ \cdots \succ t^{1[1]} \succ a_{e_m^2}^{1[e_m]} \succ \cdots \succ a_{e_1^2}^{1[e_1]} \succ \boxed{y^{[v]}} \succ [\dots]$$

$$Y^{[vw]}: \quad a_{e_m^2}^{1[e_m]} \succ \cdots \succ a_{e_1^2}^{1[e_1]} \succ y \succ \boxed{y^{[vw]}} \succ [\dots]$$

$$\begin{aligned} A_v^\ell: \quad & y^{[v]} \succ (a_v^{\ell-1*}) \succ a_v \succ a_v^{\ell*} \succ u_v^{\delta_v(d_v)} \succ a_v^{\ell[\delta_v(d_v)]} \succ \cdots \succ \\ & (a_v^{\ell+1[\delta_v(2)]}) \succ u_v^{\delta_v(2)} \succ a_v^{\ell[\delta_v(2)]} \succ (a_v^{\ell+1[\delta_v(1)]}) \succ u_v^{\delta_v(1)} \succ a_v^{\ell[\delta_v(1)]} \succ \boxed{a_v^\ell} \succ [\dots] \end{aligned}$$

$$\begin{aligned} T^\ell: \quad & y^{[t]} \succ (t^{\ell-1*}) \succ t \succ t^{\ell*} \succ a_n \succ t^{\ell[n]} \succ \cdots \succ \\ & (t^{\ell+1[2]}) \succ a_2 \succ t^{\ell[2]} \succ (t^{\ell+1[1]}) \succ a_1 \succ t^{\ell[1]} \succ \boxed{t^\ell} \succ [\dots] \end{aligned}$$

$$\begin{aligned} Y: \quad & x \succ t \succ y^{[t]} \succ t^{1[n]} \succ a_n \succ y^{[n]} \succ \cdots \succ t^{1[1]} \succ a_1 \succ y^{[1]} \succ \\ & a_{e_m^2}^{1[e_m]} \succ u_{e_m^2}^{e_m} \succ a_{e_m^1}^{1[e_m]} \succ u_{e_m^1}^{e_m} \succ y^{[e_m]} \succ \cdots \succ a_{e_1^2}^{1[e_1]} \succ u_{e_1^2}^{e_1} \succ a_{e_1^1}^{1[e_1]} \succ u_{e_1^1}^{e_1} \succ y^{[e_1]} \succ \boxed{y} \succ [\dots] \end{aligned}$$

$$T: \quad t^{k*} \succ \boxed{t} \succ [\dots]$$

$$U_w^{vw} : y^{[vw]} \succ \boxed{u_w^{vw}} \succ [\dots]$$

$$A_v^{\ell*}: \ u_v^{\delta_v(d_v)} \succ \cdots \succ u_v^{\delta_v(1)} \succ \boxed{a_v^{\ell*}} \succ [\dots]$$

$$T^{\ell*}: \ a_1 \succ a_2 \succ \cdots \succ a_n \succ \boxed{t^{\ell*}} \succ [\dots]$$

We claim that there exists a clique of size k in graph G iff object x can reach agent Y within a total of $k^3 + 4k^2 + k + 2$ swaps or a makespan of $5k(k-1)/2 + 3k + 4$. An agent A_v^ℓ (or T^ℓ) is said to be "validated" if she obtains at a moment object $a_v^{\ell*}$ (or $t^{\ell*}$). We omit the details of the proof but the idea is that, let object x reach agent Y, all the k agents T^ℓ and all the $k-1$ agents

A_v^ℓ of k branches A_v need to be validated. The associated clique in graph G is given by the vertices v for which all the $k-1$ agents A_v^ℓ have been validated. All the agents A_v^ℓ, for $1 \leq \ell < k$, are validated if we can bring in the branch $k-1$ objects u_v^{vw} (or u_v^{wv}, following the order) representing an edge incident to v. Observe that the given budget allows bringing in the branches only $k(k-1)$ objects u_v^{vw} and the construction forces to choose u_w^{vw} if u_v^{vw} has been chosen. □

Combining the proofs of Theorem 6 and Proposition 2 leads to hardness for GLS.

Corollary 7. GLS-*sum and* GLS-*makespan are* co-W[1]-*hard even for trees.*

This W[1] -hardness for RO and GLS parameterized by the length of the sequence rules out the existence of FPT algorithms even in trees under standard complexity assumptions. However, the following results on the membership to respectively W[1] and co-A[2] show that the problems are not so hard. They are notably in XP for any graph, thus tractable when the parameter is a constant.

Theorem 8. RO-*sum is in* W[1]*.*

Proof. An instance $\mathcal{I}$ of RO with a swap dynamics model $(N, M, \succ, \sigma^0, G)$, agent A and object x, and k as a total number of swaps, is transformed into an instance $\mathcal{I}' = (\mathcal{A}, \varphi)$ of $\mathrm{MC}(\Sigma_1)$, known to be W[1]-complete (Theorem 1). Structure $\mathcal{A}$ is an $(E, \succ, \sigma^0, A, X)$-structure with variables in $N \cup M$, where relations E, $\succ$, σ^0, A and X are defined as follows. The binary relation E over N^2 represents the edge set E. The ternary relation $\succ$ over $N \times M^2$ represents the preferences of the agents, i.e. $\succ(i, a, b)$ means that agent i prefers object a to b. For the sake of clarity, we write $a \succ_i b$ instead of $\succ(i, a, b)$. The binary relation σ^0 over $N \times M$ represents the initial allocation σ^0, i.e. $\sigma^0(i, z)$ means that i is initially endowed with z. Finally, the unary relations A and X respectively represent agent A and object x, i.e. $A(y)$ means that y is agent A and $X(y)$ means that y is object x.

The Σ_1-formula φ is defined as $\varphi = \exists x_0 \exists b_0 \exists x_1 \exists y_1 \exists a_1 \exists b_1 \ldots \exists x_k \exists y_k \exists a_k \exists b_k \Big(\sigma^0(x_0, b_0) \wedge \bigvee_{0 \leq k' \leq k} \psi^{k'}\Big)$ with

$$\psi^{k'} \equiv A(x_{k'}) \wedge X(b_{k'}) \wedge \bigwedge_{i=1}^{k'} \Big(E(x_i, y_i) \wedge b_i \underset{x_i}{\succ} a_i \wedge \; a_i \underset{y_i}{\succ} b_i \wedge \; \mathrm{o}_i(x_i, a_i) \wedge \mathrm{o}_i(y_i, b_i)\Big)$$

where for all i, $\mathrm{o}_i(q, r)$ stands for $\Big(\sigma^0(q, r) \wedge \bigwedge_{j=1}^{i-1} x_j \neq q \wedge y_j \neq q\Big) \vee \bigvee_{j=1}^{i-1} \Big(\bigwedge_{p=j+1}^{i-1} x_p \neq q \wedge q_p \neq q\Big) \wedge \Big((x_j = q \wedge \mathrm{o}_j(y_j, r)) \vee \Big(y_j = q \wedge \mathrm{o}_j(x_j, r)\Big)\Big)$.

One can prove by induction over i that formula $\mathrm{o}_i(q, r)$ is true iff object r is owned by agent q before i^{th} swap. Globally, formula $\psi^{k'}$ is true iff the sequence of exchanges between the agents (x_i, y_i) exchanging the objects (a_i, b_i), for $i \in \{1, \ldots, k'\}$, is a sequence of swaps leading to give object x to agent A.□

The same idea and a slightly different FO formula work for RO-makespan.

Proposition 9. RO-*makespan is in* W[1]*.*

Proof (sketch). We reduce to MC($\varSigma_1$) but face a new difficulty. We cannot quantify over all potential exchanges within makespan k: it would lead to a formula of size $\varOmega(n)$. The crux of this proof is to observe that not all exchanges are relevant to decide the problem. Assume we process independent swaps in parallel for up to k time steps. Looking at it from the end, the only relevant swap in the last step k involves agent A, so we quantify over a single swap and ignore all concurrent ones. In the one-before-last, only swaps involving A or A's partner at step k may be relevant. So considering two swaps happening at step $k-1$ and ignoring all other concurrent ones suffice. All in all, we need to quantify over no more than 2^{k+1} exchanges. The rest is similar to that of Theorem 8. □

A similar reasoning is applied to GLS. We reduce GLS to model-checking FO formula using more sophisticated $\varSigma_2$ formulas.

Proposition 10. GLS-*sum/-makespan is in* co-A[2]*.*

Proof (sketch). We reduce co-GLS to MC($\varSigma_2$), known to be A[2]-complete (Theorem 1), following an approach similar to that of Theorem 8 and Proposition 9. □

The previous results show that RO and GLS are not "so hard" considering the length of the sequence as a parameter. Furthermore, for some natural classes of graphs, the problems are even tractable with respect to these parameters.

Proposition 11. RO/GLS-*sum/makespan are* FPT *on bounded degree graphs.*

Proof. The proof follows the idea developed for Proposition 9. Let $\varDelta$ be the degree of G and consider the RO problem. At the k^{th} step, the only relevant exchange involves agent A and a neighbor, so there are $\mathcal{O}(\varDelta)$ possible swaps. The one-before-last step can only involve A or one her neighbor, therefore there are at most $2\varDelta$ possible swaps for RO-sum and at most $\varDelta + \varDelta^2$ for RO-makespan. This argument applies at any of the k steps, hence there are $\mathcal{O}(\varDelta^k.k!)$ sequences of swaps for RO-sum and $\mathcal{O}(\varDelta^{k^2})$ for RO-makespan, and it suffices to verify if one sequence assigns x to A. Concerning GLS, it suffices to test the reachability to A of any object x such that $y \succ_A x$, and so it just adds a factor of n. □

6 Conclusion and Perspectives

This article studies the distributed process of swap dynamics along a network for reallocating objects among agents. Two related problems are investigated: Reachable Object (RO), "can a given agent obtain a given object?", and Guaranteed Level of Satisfaction (GLS), "is a given agent guaranteed to get a given object or better?". Both problems are hard but the parameterized approach allows us to escape this difficulty for a relevant class of graphs.

We consider natural parameters constraining the number of swaps per agent or the duration of the sequence. Assuming that they remain small is reasonable in practice as the patience of the agents typically does not increase with the instance size. In the case of few swaps *per agent*, RO and GLS remain hard even on bounded degree graphs. So, this parameterization, although natural, does not help us to grasp the problems. However, considering the *total* length of the sequence, although both problems are intractable even for trees, this hardness is circumscribed to not "so hard" parameterized complexity classes, leading to the possibility of handling the problems when the parameters do not depend on the instance size, very natural assumption. Furthermore, unlike the first parameter, the length of the sequence permits to obtain fixed parameter tractability on bounded degree graphs, which typically model real social networks.

The parameterized approach allows progress in the understanding of the problems and leads to significant and realistic positive results. So far, we have considered restrictions on the network as well as on the solution size. A natural extension is to investigate the influence of a third dimension: constraints on the preference profile, e.g., single-peaked domains. Furthermore, assuming the full knowledge of the preferences and the network is not relevant in all the contexts. Relaxing this assumption could be a challenging future work.

References

1. Abdulkadiroğlu, A., Sönmez, T.: House allocation with existing tenants. J. Econ. Theory **88**(2), 233–260 (1999)
2. Abrahamson, K.A., Downey, R.G., Fellows, M.R.: Fixed-parameter tractability and completeness IV: on completeness for W[P] and PSPACE analogues. Ann. Pure Appl. Log. **73**(3), 235–276 (1995)
3. Aziz, H., Biró, P., Lang, J., Lesca, J., Monnot, J.: Optimal reallocation under additive and ordinal preferences. In: AAMAS. pp. 402–410 (2016)
4. Aziz, H., De Keijzer, B.: Housing markets with indifferences: a tale of two mechanisms. In: AAAI. pp. 1249–1255 (2012)
5. Berman, P., Karpinski, M., Scott, A.D.: Approximation hardness of short symmetric instances of max-3sat. Technical report (2004)
6. Chevaleyre, Y., et al.: Issues in multiagent resource allocation. Informatica **30**(1), 3–31 (2006)
7. Chevaleyre, Y., Endriss, U., Estivie, S., Maudet, N.: Multiagent resource allocation in k-additive domains: preference representation and complexity. Ann. Oper. Res. **163**(1), 49–62 (2008)
8. Chevaleyre, Y., Endriss, U., Lang, J., Maudet, N.: Negotiating over small bundles of resources. In: AAMAS. pp. 296–302 (2005)
9. Chevaleyre, Y., Endriss, U., Maudet, N.: Allocating goods on a graph to eliminate envy. In: AAAI. pp. 700–705 (2007)
10. Damamme, A., Beynier, A., Chevaleyre, Y., Maudet, N.: The power of swap deals in distributed resource allocation. In: AAMAS. pp. 625–633 (2015)
11. Dunne, P.E., Chevaleyre, Y.: The complexity of deciding reachability properties of distributed negotiation schemes. Theor. Comput. Sci. **396**(1–3), 113–144 (2008)
12. Dunne, P.E., Wooldridge, M., Laurence, M.: The complexity of contract negotiation. Artif. Intell. **164**(1–2), 23–46 (2005)
13. Easley, D., Kleinberg, J.: Networks, Crowds, and Markets: Reasoning about a Highly Connected World. Cambridge University Press, New York (2010)

14. Endriss, U., Maudet, N., Sadri, F., Toni, F.: Negotiating socially optimal allocations of resources. J. Artif. Intell. Res. **25**, 315–348 (2006)
15. Flum, J., Grohe, M.: Parameterized Complexity Theory. Springer, Berlin (2006)
16. Gourvès, L., Lesca, J., Wilczynski, A.: Object allocation via swaps along a social network. In: IJCAI. pp. 213–219 (2017)
17. Jackson, M.O.: Social and Economic Networks. Princeton University Press, Princeton (2008)
18. Sandholm, T.W.: Contract types for satisficing task allocation. In: AAAI Spring Symposium. pp. 23–25 (1998)
19. Shapley, L., Scarf, H.: On cores and indivisibility. J. Math. Econ. **1**(1), 23–37 (1974)