

Pareto optimal coalitions of fixed size

ÁGNES CSEH – TAMÁS FLEINER – PETRA HARJÁN

CERS-IE WP – 2020/5

January 2020

<https://www.mtaki.hu/wp-content/uploads/2020/01/CERSIEWP202005.pdf>

CERS-IE Working Papers are circulated to promote discussion and provoke comments, they have not been peer-reviewed.
Any references to discussion papers should clearly state that the paper is preliminary.
Materials published in this series may be subject to further publication.

ABSTRACT

We tackle the problem of partitioning players into groups of fixed size, such as allocating eligible students to shared dormitory rooms. Each student submits preferences over the other individual students. We study several settings, which differ in the size of the rooms to be filled, the orderedness or completeness of the preferences, and the way of calculating the value of a coalition---based on the best or worst roommate in the coalition. In all cases, we determine the complexity of deciding the existence, and then finding a Pareto optimal assignment, and the complexity of verifying Pareto optimality for a given assignment.

JEL codes: C63, C78

Keywords: Pareto-optimality, coalition, game theory

Ágnes Cseh

Centre for Economic and Regional Studies, Institute of Economics, Tóth Kálmán utca
4., Budapest, 1097, Hungary
e-mail: cseh.agnes@krtk.mta.hu

Tamás Fleiner

Centre for Economic and Regional Studies, Institute of Economics, Tóth Kálmán utca
4., Budapest, 1097, Hungary
and
Department of Computer Science and Information Theory, Budapest University of
Technology and Economics
e-mail: fleiner@cs.bme.hu

Petra Harján

Department of Computer Science and Information Theory, Budapest University of
Technology and Economics
e-mail: harjan@cs.bme.hu

Fix méretű Pareto-optimális koalíciók

ÁGNES CSEH – TAMÁS FLEINER – PETRA HARJÁN

ÖSSZEFOGLALÓ

A koalíciók tekinthetők a párosítások egyik általánosításaként, mivel egy koalícióban is egymáshoz rendeljük a játékosokat, viszont a csoportok mérete nem feltétlenül két fő, ahogy a párosításoknál. Egy koalícióbeosztást Pareto-optimálisnak nevezünk akkor, ha nincs olyan alternatív beosztás, amivel egyetlen játékos sem jár rosszabbul az előzőhöz képest, legalább egy játékos viszont jobban jár. Olyan koalícióformációs helyzeteket tanulmányozunk, ahol az egyes játékosoknak egymáson szigorúan rendezett preferencialistájuk van. Ilyen például egy kollégiumi szobabeosztás, ahol minden leendő kollégista rangsorolja a lehetséges szobatársait. A szobák előre meghatározott, fix méretűek. A végső szobabeosztásban egy hallgató elégedettségét kétféleképpen mérhetjük: a legkedvesebb és a legkevésbé kedvelt szobatárs alapján. Ez szinkronban van a játékelméletből ismert best-hedonic és worst-hedonic fogalmakkal. Cikkünkben annak a számításelméleti bonyolultságát vizsgáljuk, hogy milyen nehéz a két elégedettségi mérték szerinti Pareto-optimális szobabeosztást találni.

JEL: C63, C78

Kulcsszavak: Pareto-optimális, koalíció, játékelmélet

Pareto optimal coalitions of fixed size [★]

Ágnes Cseh¹, Tamás Fleiner^{1,2}, and Petra Harján²

¹ Centre for Economic and Regional Studies, Institute of Economics

² Department of Computer Science and Information Theory, Budapest University of Technology and Economics

Abstract. We tackle the problem of partitioning players into groups of fixed size, such as allocating eligible students to shared dormitory rooms. Each student submits preferences over the other individual students. We study several settings, which differ in the size of the rooms to be filled, the orderedness or completeness of the preferences, and the way of calculating the value of a coalition—based on the best or worst roommate in the coalition. In all cases, we determine the complexity of deciding the existence, and then finding a Pareto optimal assignment, and the complexity of verifying Pareto optimality for a given assignment.

1 Introduction

The ubiquitous nature of coalition formation has been triggering researchers from various fields to study the behavior of individuals forming groups [10, 11, 14, 16, 21]. A large portion of the game-theoretical studies focuses on the individuals having ordinal preferences over the possible outcomes [5–8, 20].

Our setting involves n players who need to be partitioned into coalitions. For convenience, we talk about assigning each player to a room. We assume the rooms to have no specific feature besides their capacity. To ensure a feasible partition in the outcome, we assume that the total capacity of rooms adds up to n , and a feasible assignment fills each room to its capacity. Having once entered the scheme, players have no option to opt out. Each player submits her preference list over the other players. Four features of the problem define various settings, each of which is realistic and will be investigated by us.

- (i) Preference lists might be *complete* or *incomplete*, this latter meaning that players have the right to declare some of the other players unacceptable as a roommate. No player can be put in the same room with an unacceptable roommate.
- (ii) Players might have *strictly* or *weakly ordered* lists.
- (iii) Players might compare two coalitions based on their *best* or *worst* assigned roommate.
- (iv) The rooms might all accommodate 3 players, or they might have a predefined *capacity* each, which we denote by r_1, r_2, \dots, r_k for each of the k rooms.

[★] Supported by the Cooperation of Excellences Grant (KEP-6/2018), by the Ministry of Human Resources under its New National Excellence Programme (ÚNKP-18-4-BME-331), the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2016), its János Bolyai Research Fellowship, and OTKA grant K128611. This work is connected to the scientific program of the “Development of quality-oriented and harmonized R+D+I strategy and functional model at BME” project, supported by the New Hungary Development Plan (Project ID:TÁMOP-4.2.1/B-09/1/KMR-2010-0002).

The optimality principle we are focusing on is *Pareto optimality*. A room assignment, or, a set of coalitions is Pareto optimal if there is no other assignment in which at least one player is better off, and no player is worse off than in the first assignment. The comparison here is defined based on point (iii) above. We shorten the term ‘Pareto optimal assignment’ to POA. Our goal is to study all 2^4 combinations of the above four features, and for each of them, determine the complexity of the following three problems:

- (1) verifying whether a given feasible assignment is a POA;
- (2) checking whether a POA exists;
- (3) finding a POA.

1.1 Related literature

Pareto optimality in coalition formation has a rich literature. We start with an overview on coalition formation viewed as a hedonic game. Then we review the settings in which the coalition size matters.

Hedonic games. Coalition formation under preferences can be seen as a hedonic game [3, 5, 6]. In such a game, players have preferences over the possible coalitions they can be part of, and coalitions can be of any size—notice that these two basic features strikingly differ from our setting. Pareto optimal coalition formation as a hedonic game is extensively studied by Aziz et al. [2]. They analyze two restricted variants of hedonic games that are closely related to our setting. Both of them operate under the assumption that the coalition size is arbitrary, but they both derive players’ preferences on coalitions from a preference list on individual players. They show that if the preference lists are incomplete, and preferences are based on the best roommate, then verifying Pareto optimality for a given assignment is **coNP**-complete, and computing a POA is **NP**-hard. For complete lists and the same preferences, the grand coalition is a trivial optimal solution, since every player has their first-choice roommate in the sole room. Aziz et al. [2] also show that if preferences are based on the worst roommate, then both verifying Pareto optimality for a given assignment, and computing a POA can be done in polynomial time.

2-person rooms. Some of the literature concentrates on each coalition being of size 1 or 2. In this setting, a player can compare two coalitions simply based on the rank of the sole roommate (if any exists), so our point (iii) does not apply here. Using Morrill’s algorithm [18], Aziz et al. [2] show that even if preferences contain ties, both verifying Pareto optimality for a given matching, and calculating a POA are solvable in polynomial time. Their results are valid for complete and incomplete lists as well. Abraham and Manlove [1] consider Pareto optimal matchings as a means of coping with instances of the stable roommates problem with strict lists, which do not admit a stable matching. They show that while a maximum size POA is easy to find, finding a minimum size POA is **NP**-hard.

3-person rooms. For the setting in which a room can accommodate up to 3 players, two versions of the problem have been studied. A *three-cyclic game* is a hedonic game in which the set of players is divided into men, women, and dogs and only kind of acceptable coalitions are man-woman-dog triplets [17, 19]. Furthermore, men only care about women, women only care about dogs and dogs only care about men. Computing a POA is known to be **NP**-hard for these games, while the corresponding verification problem is **coNP**-complete, even for strict preferences [2]. In *room-roommate games*, a set of players act as rooms, and these have no preferences whatsoever.

The ordinary players, on the other hand, have a preference list over all possible roommate-room pairs they find acceptable. A triplet is feasible if exactly one player in it plays a room. If preferences are strict, then a POA can be computed in polynomial time, but the problem becomes NP-hard as soon as ties are introduced, even if all lists are complete [2]. Just as in the previous problem, the verification version is coNP-complete, even for strict preferences [2].

Preferences depending on the room size. Anonymous hedonic games [4] are a subclass of hedonic games in which the players' preferences over coalitions only depend on coalition sizes. Both verification and finding a POA are hard in such games [2]. Darmann [9] studies a group activity selection model in which players have preferences not only on the activity, but also on the number of participants in their coalition. He provides an efficient algorithm to find a POA, if each player wants to share an activity with as many, or as few players as possible.

As we have seen, a number of papers investigate the complexity of finding a POA under various settings, such as limited coalition size or preferences over players rather than over coalitions. However, there is no work on the combination of these two. In our setting, homogeneous players *rank each other*, and form coalitions of an arbitrary, but *fixed* size.

1.2 Our contribution

We tackle the problems of verifying Pareto optimality, deciding the existence of POA, and finding a POA for a set of fixed coalition sizes. We distinguish 2^4 cases, based on the completeness and the orderedness of preferences, the way of comparing two coalitions by a player, and the room sizes. Our findings are summarized below and in Figure 1.

- Verification is coNP-complete in all cases. We show this in Section 3 by two reductions from triangle cover problems.
- If lists are incomplete, then deciding whether a feasible assignment exists is already NP-complete in all cases. On the other hand, if a feasible assignment does exist, then an optimal one exists as well. These are due to a simple NP-completeness reduction and a monotonicity argument, which can be found in Section 4.
- For complete lists, a POA is bound to exist. In 3 out of the 16 cases, serial dictatorship delivers one. In all other cases, by computing *any* POA in polynomial time, one could answer an NP-complete decision problem in polynomial time. In Section 5 we elaborate on these easy and hard cases. For the positive results, we interpret serial dictatorship in the current problem settings. Then, we utilize a tool developed by Aziz et al. [2] in the hardness proofs.

2 Preliminaries

In this section we set the solid mathematical basis for discussing our problems. Then, we introduce a selection of NP-complete problems and prove that a specific variant of them is hard as well. We do this to prepare the ground for our hardness proofs in Sections 3-5.

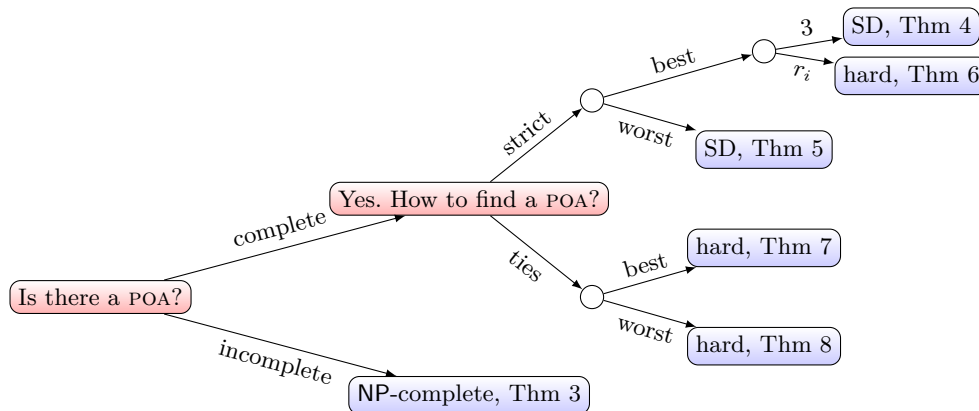


Fig. 1. The complexity of finding a POA. The arrows explain the case distinctions according to points (i)–(iv) listed in the Introduction. SD stands for serial dictatorship.

2.1 Problem definition

Our input consists of a set of players P , with $|P| = n$, a multiset r_1, r_2, \dots, r_k of room capacities with $r_1 + r_2 + \dots + r_k = n$, and strictly ordered, but not necessary complete preferences for each player $i \in P$ on other players. An assignment R is the partition of P into sets R_1, R_2, \dots, R_k . This means that players have no choice to opt out once they have entered the market.

Definition 1. An assignment $R = \{R_1, R_2, \dots, R_k\}$ is called feasible if

1. $|R_i| = r_i$ for each $1 \leq i \leq k$;
2. each player in R_i declares every other player in the same R_i acceptable.

Players evaluate their situation in an assignment solely based on the roommates they are grouped together with and do not care about the other rooms. For point 2 in Definition 1, players only need to compare outcomes in which they find all other roommates acceptable. We define two comparison principles.

When the *best* roommate counts, coalition R_i is preferred to R'_i by player i if player $j \in R$ ranked highest by i among all players in R_i is preferred by i to player $j' \in R$ ranked highest by i among all players in R'_i . Analogously, when the *worst* roommate counts, coalition R_i is preferred to R'_i by player i if player $j \in R$ ranked lowest by i among all players in R_i is preferred by i to player $j' \in R$ ranked lowest by i among all players in R'_i . Our definitions are aligned with \mathcal{W} -preferences in [8], and with \mathcal{B} -hedonic and \mathcal{W} -hedonic games in [2]. However, \mathcal{B} -preferences in [8] are different, because there, the size of the coalition serves as a tiebreaker in case the best roommate is identical in the two coalitions. We consider two such coalitions equally good.

2.2 Relevant hard problems

We first introduce two variants of a hard graph cover problem that we will later use in our hardness reductions, and then shortly mention a bin packing problem as well.

Problem 1 TRIANGLE COVER

Input: A simple graph $G = (V, E)$.

Question: Does there exist a partition of V into sets of cardinality 3 so that for each set, the three vertices span a 3-cycle?

Problem 2 DIRECTED TRIANGLE COVER

Input: A simple directed graph $D = (V, A)$.

Question: Does there exist a partition of V into sets of cardinality 3 so that for each set, the three vertices span a directed 3-cycle?

The TRIANGLE COVER problem asks for a set of vertex-disjoint 3-cliques (triangles) in the input graph G , so that these cover all vertices of G . This problem has been shown to be NP-complete by Garey and Johnson [12]. The directed version has not been proved to be NP-complete, so we give a simple hardness proof below. The hard problem we reduce to DIRECTED TRIANGLE COVER is 3D HYPERGRAPH MATCHING [13, 15].

Problem 3 3D HYPERGRAPH MATCHING

Input: A hypergraph $H = (U \cup V \cup W, E)$, where each $e \in E$ is a triple (u, v, w) so that $u \in U$, $v \in V$, and $w \in W$.

Question: Does there exist a perfect matching in H ?

Claim. DIRECTED TRIANGLE COVER is an NP-complete problem.

Proof. We start with the input graph of an arbitrary instance of 3D HYPERGRAPH MATCHING and transform it into an instance of DIRECTED TRIANGLE COVER. We keep the set of vertices, and simply replace each hyperedge (u, v, w) of H by the gadget shown in Figure 2.2. This gadget introduces 9 new vertices per hyperedge to the set of vertices, which we will call the *gadget vertices*.

Assume first that a directed triangle cover exists in the resulting directed graph D . For each gadget, two scenarios can occur. Either vertices a, b, c form a triangle and then all of u, v , and w must be covered by triangles inside the same gadget, or triangle (a, b, c) is not in the cover, but then all of u, v , and w are covered by triangles outside the gadget. The first case corresponds to the hyperedge (u, v, w) being in our perfect matching M , while the second case tells otherwise. We still need to show that M is indeed a perfect matching. Firstly, no vertex can appear in two hyperedges in M , because it would cover the same vertex twice by directed triangles. Secondly, if a vertex is left unmatched, then it is also left uncovered by the triangle cover, which is a contradiction.

To show the opposite direction, we first assume that a perfect matching M exists in the original instance. We transform it into directed triangles in the created directed graph D as above: the edges of M will be the gadget with (a, b, c) (the black triangles), the rest of the edges will be the gadgets without (a, b, c) (the gray triangles). We need to show that these triangles form a cover. The gadget vertices are trivially covered exactly once. The rest of the vertices are matched in M along some hyperedge, and the gadget belonging to this hyperedge covers them with a triangle inside the gadget. \square

Remark 1. It is an obvious observation that both TRIANGLE COVER and DIRECTED TRIANGLE COVER are hard only in graphs for which $|V|$ is a multiple of 3. We will thus assume that G in the input of these hard problems has this property. Besides this, we can also assume that G is simple, and that each vertex is of degree at least 2, because graphs with an isolated or a degree 1 vertex are trivial no-instances. For the directed version, one can even assume that each vertex has at least one outgoing and at least one incoming edge.




Fig. 2. The gadget replacing each hyperedge $(u, v, w) \in E(H)$. The unfilled vertices are gadget vertices, and these are not connected to any vertex outside of this gadget. Triangles in this gadget are either black or gray. The gray triangles cover all of these, and leave the remaining three vertices uncovered, while the black triangles cover all vertices in the gadget.

We close this section with the final element in our toolbox of hard problems. The input of UNARY BIN PACKING is a set of item sizes, and a bin size, all encoded in unary. The goal is to group all items into bins so that the total item size in each bin is exactly the bin size. This problem has been shown to be NP-complete by Garey and Johnson [12].

Problem 4 UNARY BIN PACKING

Input: A set of item sizes i_1, i_2, \dots, i_n , and a bin size b , all encoded in unary.

Question: Does there exist a partitioning of i_1, i_2, \dots, i_n so that the sum of item sizes in each set of the partition adds up to b ?

Remark 2. For our proofs, we will assume that the smallest item size is at least 2. The hardness of this variant is easy to see. If we take an input of UNARY BIN PACKING and multiply all item and bin sizes by 2, then we get an equally hard problem that can be encoded in twice as many bits as the original one.

3 Verification

In this section, we show the hardness of verification for all cases. We present two proofs: in Theorem 1, the worst roommate defines the base of comparison for two coalitions, while it is the best roommate who counts in Theorem 2. Other than this, we restrict our reduction to the least general case of the problem, having strict and complete lists, and 3-person rooms.

Theorem 1 (Verification, strict and complete lists, worst roommate counts, 3-person rooms). *Supposing that the preferences of a player depend on the worst roommate, verifying whether a given assignment is Pareto optimal is a coNP-complete task even if all preferences are strict and complete, and every room is of size 3.*

Proof. To each instance G of TRIANGLE COVER we will construct an instance of our verification problem, consisting of players with strict and complete preferences, and an assignment on which Pareto optimality is to be verified. We will show that a triangle cover exists in the first instance if and only if the assignment has a Pareto improvement.




Fig. 3. The assignment instance constructed to G in the proof of Theorem 1. The numbers on the edges mark the preferences of the players. Players in V_1 rank other players in V_1 higher than players in V_2 and V_3 , as the numbers on their solid black and dashed gray edges indicate. Players in V_2 and V_3 prefer their dotted blue edges to their dashed gray edges. The dotted blue edges form a triangle cover of $V_2 \cup V_3$ —they connect the leftmost two vertices in V_3 with the rightmost vertex in V_2 .

First we draw graph G and also make two further copies of its vertex set V_1 . We denote these copies by V_2 and V_3 . Each vertex in $V_1 \cup V_2 \cup V_3$ represents a player in our POA instance. To show hardness in the most general case, we construct our POA instance with complete preference lists, which translates into a complete graph on the vertex set $V_1 \cup V_2 \cup V_3$. The edges of this graph can be partitioned into four classes.

- The original edges of G are solid and black in Figure 3. These edges are the best choices of both of their end vertices, and the order among them can be chosen arbitrarily. Recall that each vertex in V_1 has at least two of these black edges, as argued in Remark 1 in Section 2.
- The copied vertices V_2 and V_3 are connected by dotted blue triangles, as shown in Figure 3. Each such triangle connects three vertices that originate from three different vertices of G . Notice that these triangles can cover the entire set $V_2 \cup V_3$, because both $|V_2|$ and $|V_3|$ are multiples of 3, as noted in Remark 1. The rank of these edges is the highest possible, and their order among themselves does not matter.
- The three copies of the same vertex in G are connected by a dashed gray triangle in Figure 3. These edges are ranked lower than the solid black and the dotted blue edges. The order among them does not matter.
- All edges not visible in Figure 3 are ranked lower than the listed edges, in an arbitrary order.

The verification will happen with respect to the assignment R built by all dashed gray edges.

If G has a triangle cover, then these solid black triangles cover the entire set V_1 . In our assignment problem, these triangles translate into coalitions of size 3. The dotted blue triangles on V_2 and V_3 complete the alternative assignment. It is easy to see that every agent is better off by switching to the solid black and the dotted blue edges, since they are always ranked higher than the dashed gray edges. Thus, R is not a POA.

Now we show the opposite direction. If there is a Pareto improvement to the assignment marked by the dashed gray edges, then it may use none of the invisible edges, since they are all worse than the dashed gray ones, and no player is allowed to receive a worse partner in the improved assignment. So we need to find an alternative assignment using only the first 3 types of edges. It is clear that by breaking any dashed gray coalition, at least one agent in each vertex group should be reassigned to a different coalition. Now, players in $V_2 \cup V_3$ can only choose a dotted blue edge instead of a dashed gray one. The sparsity of edges gives us that if there exists one dotted blue edge in a triangle, then not only the whole triangle is dotted blue, but we also must have all of the dotted blue triangles in the new assignment—otherwise players remain unassigned. Thus, in the new assignment, we must have all of the dotted blue triangles, and V_1 's vertices are free to be grouped up among themselves. This they can only do using the solid black edges, which correspond to the original edges in G . Thus if we have a Pareto improvement of the dashed gray assignment, then we can cover V_1 with disjoint, solid black triangles. It means that G has a triangle cover. \square

Theorem 2 (Verification, strict and complete lists, best roommate counts, 3-person rooms). *Supposing that the preferences of a player depend on the best roommate, verifying whether a given assignment is Pareto optimal is a coNP-complete task even if all preferences are strict and complete, and every room is of size 3.*

Proof. This proof follows the lines of the previous one, but it reduces our problem to DIRECTED TRIANGLE COVER. To each instance D of DIRECTED TRIANGLE COVER we will construct an instance of our verification problem, consisting of players with preferences, and an assignment on which Pareto optimality is to be verified. We will show that a directed triangle cover exists in the first instance if and only if the assignment has a Pareto improvement.

First we draw the directed graph D and make two further copies of its vertex set V_1 . We denote these copies by V_2 and V_3 . Just as in our previous proof, the edges of this complete graph can be partitioned into four classes. Notice that the preferences differ from the ones in our previous proof.

- The original directed edges of D are solid and black in Figure 4. These edges are the best choices of their starting vertex and they are ranked below the dashed gray edges at their end vertex.




Fig. 4. The assignment instance constructed to the directed graph D in the proof of Theorem 2. The numbers on the edges mark the preferences of the players. Players in V_1 rank their outgoing edges first, then their edges to players in V_2 and V_3 , and then their incoming edges, as the numbers on their solid black and dashed gray edges indicate. Preferences in the dashed gray and dotted blue triangles are cyclic. Players in V_2 and V_3 rank their two dashed gray edges sandwiched between their two dotted blue edges.

The order among all outgoing and among all incoming directed edges of the same vertex can be chosen arbitrarily. Recall that each vertex in V_1 has at least one outgoing and at least one incoming edge, as stated in Remark 1.

- The copied vertices V_2 and V_3 are connected by dotted blue triangles, as shown in Figure 4. Each such triangle connects three vertices that originate from three different vertices of D . The rank of these edges is either 1 or 4, so that the triangle forms a preference cycle, i.e. each edge is ranked first by one end vertex and fourth by the other one. Edges in the middle, ranked second and third, will be the dashed gray edges.
- The three copies of the same vertex in G are connected by a dashed gray triangle in Figure 4. These edges are ranked between outgoing and incoming edges in D at vertices in V_1 , and they are ranked second and third by vertices in $V_2 \cup V_3$. The order among them matters: the triangles themselves must form a preference cycle.
- All edges not visible in Figure 4 are ranked lower than the listed edges, in an arbitrary order.

Again, the verification will happen with respect to the assignment R built by all dashed gray edges.

Suppose that there exists a DIRECTED TRIANGLE COVER in D . This means that each vertex in V_1 has an outgoing edge in the triangle cover, which is a first-choice roommate in the assignment problem. So switching to the coalitions marked by the triangle partition would be a Pareto im-

provement for the vertices in V_1 . We need to take care of the vertices in $V_2 \cup V_3$ too. The dotted blue triangles complete the alternative assignment, and due to the cyclic nature of the preferences on them, they too assign each player a first choice roommate, which was not present in the original assignment.

Now suppose that there exists a Pareto improvement to the assignment R marked by the dashed gray edges. In R , the best roommate of each player is ranked right below its outgoing black edges for vertices in V_1 , and second for vertices in V_2 and V_3 . In order to make at least one player, say, player i better off, the alternative assignment R' must allocate i to a more preferred neighbor j . This j must be the end vertex of the directed solid black edge (i, j) , if $i \in V_1$, and i 's first choice roommate (dotted blue edge) if $i \in V_2 \cup V_3$. This is only possible if we find a (non-directed) triangle partition in the constructed graph so that it contains at least one solid black or dotted blue edge, which is (i, j) . We now search for a third player to complete the coalition in R' . Since j just received a bad roommate in the person of i , she must have her first or second choice in the room as well, in order to keep her satisfied. Her second choice is one of her copies j' , but j' only ranks j third, and i beyond all listed players, so her situation would worsen if we put her up in a room with i and j . On the other hand, j 's first choice edge (j, ℓ) is of the same color as (i, j) . The position of ℓ is not worse than in the original assignment if and only if (i, j, ℓ) forms a blue or a directed black triangle. This shows that only red, blue, or directed black triangles can appear in a Pareto improved assignment. The existence of such a monochromatic triangle partition implies that there is a directed triangle cover in D . \square

4 Existence

In all investigated cases, if the instance admits a feasible assignment, then it also admits a Pareto optimal one. Due to monotonicity in the rank of the worst or best roommate, a chain of Pareto improvements starting at any feasible assignment must end in a POA, which is thus guaranteed to exist if a feasible assignment exists. This is the case if lists are complete, since all assignments filling up all rooms are feasible then. Not so for incomplete lists—by declaring some other players unacceptable, players can easily reach a situation where not even feasible assignments exist. We now show that in all cases with incomplete lists, deciding whether a feasible assignment, and thus, a POA exists, is NP-complete. Notice that since feasibility is already hard, it bears no importance whether a player judges the coalition based on the best or the worst roommate.

Theorem 3 (Existence, strict and incomplete lists, best or worst roommate counts, 3-person rooms). *If lists are incomplete, then deciding whether a feasible assignment exists is NP-complete even if all preferences are strict, and every room is of size 3.*

Proof. We show hardness via a reduction from TRIANGLE COVER. Given a graph G as the input of this problem, we associate players with vertices and acceptable roommate pairs with edges. The preferences within the set of acceptable agents can be chosen arbitrarily, because they play no role in feasibility anyway. Since feasible assignments must form coalitions of size exactly 3, each assignment in our problem corresponds to a triangle cover and vice versa. \square

This theorem shows that it is computationally infeasible to find a POA if lists are incomplete, because even deciding whether any exists is NP-complete. Moreover, even if we are given a feasible assignment—which guarantees the existence of a POA—it is also computationally infeasible to find a Pareto improvement, since deciding whether the given assignment itself is a POA is coNP-complete,

as shown in Theorems 1 and 2. From this point on, we can thus restrict our attention to instances with complete lists.

5 Finding a Pareto optimal assignment

As already mentioned in the previous section, a POA is guaranteed to exist if lists are complete. Here we distinguish all 2^3 cases based on three basic features of the problem, listed as points (ii)–(iv) in the Introduction.

5.1 Easy cases

We start with describing and analyzing the tailored variants of serial dictatorship in the cases where it delivers a POA. As usual, the core of serial dictatorship is that each dictator specifies a set of solutions that guarantee her one of her most desirable outcomes, and all later dictators must choose similarly, but within the already specified set. We show here algorithms that implement this principle under the problem settings we study.

Our algorithms are simple, and suited for the very specific problems in question, whereas the Preference Refinement Algorithm (PRA) of Aziz et al. [2] for computing a Pareto optimal and individually rational assignment in hedonic games is generic, but also more complicated. Besides this, it invokes an oracle solving a problem they call 'Perfect Partition'. Perfect Partition asks for an assignment that gives each player one of her best outcomes. (Notice that in this paper, we use the term perfectness in its traditional, graph-theoretical sense, where it means that each player is matched.) PRA starts off coarsening, and then refining the preferences of players according to certain rules. Then, it calls the oracle to compute a Perfect Partition with the refined preferences, which is shown to be Pareto optimal in the original instance. Here we give a much more direct interpretation of serial dictatorship, focusing only on the specific problem variant we discuss in Theorems 4 and 5. Our interpretations are then illustrated in Example 1.

Theorem 4 (Finding a POA, strict and complete lists, best roommate counts, 3-person rooms). *If lists are strict and complete, all rooms are of capacity 3, and the best roommate counts, then serial dictatorship delivers a POA.*

Proof. The exact implementation of serial dictatorship works in rounds, as follows. The first dictator points at her first choice. We fix this pair, and immediately assign them their final room R_1 , which will be completed by the third player later. The same happens in each round: the current dictator points at her first choice *available* roommate, we fix this pair, and assign them a room. If one player in the new pair is already in a room, the other one joins her, otherwise a new room is opened for them.

For dictator i , player j is available, if both of the following hold.

1. The number of roommates already assigned to a room together with either i or j is at most one in total, or i and j are already assigned to the same room.
2. If there is no further room to open, then j must be a player already assigned to a room.

We will now show that this procedure indeed delivers a POA. Assume indirectly that an assignment $R' = R'_1, R'_2, \dots, R'_k$ Pareto dominates the outcome $R = R_1, R_2, \dots, R_k$ of our algorithm. Among the players who are better off in R' , we choose the one who comes earliest in the order of

dictators. Let this player be i . We know that j , the best roommate of i in R' , is strictly better than i 's best roommate in R . In i 's turn in our algorithm, j thus was not available to i . The reason for this must be one of the above two points.

If the first point was the reason, then some of the earlier dictators already reserved i and j for themselves. Since we assumed that i is the earliest dictator who is better off in R' than in R , the choices of all earlier dictators must stay intact in R' , otherwise some of them receives a worse best roommate, which contradicts the fact that R' is a Pareto improvement. The second case is when all k rooms had already been open as i was considering to point to j , and neither of them had been assigned to a room yet. At least one of the fixed pairs must be split up in R' , if i and j are assigned to the same room in it. The earlier dictator in this pair must precede i in the order of dictators, and thus any change in her best allocated roommate must worsen her situation, which contradicts the Pareto improvement property. \square

Theorem 5 (Finding a POA, strict and complete lists, worst roommate counts, r_i -person rooms, including 3-person rooms). *If lists are strict and complete, the rooms are of arbitrary capacity, and the worst roommate counts, then serial dictatorship delivers a POA.*

Proof. If the worst roommate counts, serial dictatorship can be interpreted as follows. In each round, the dictator moves into one of the smallest available rooms of size r_i with her best $r_i - 1$ choice roommates. The coalition is fixed and the room is removed from the set of available rooms.

To see correctness, we apply induction. Clearly, the price for improving the position of a dictator is to harm some previous dictator, because serial dictatorship gave her the fewest possible top choices still available on her list. Thus the output of the mechanism is a POA. \square

Example 1. Figure 5 illustrates an instance. We run serial dictatorship in the three settings in which it delivers a POA. We assume the order of dictators to be 1, 2, \dots , 9.

1: 5 4 7 3 9 6 8 2	4: 3 6 7 2 9 6 8 1	7: 1 2 9 3 4 6 8 5
2: 1 4 5 9 8 6 3 7	5: 3 6 2 7 8 4 1 9	8: 6 3 7 1 9 5 4 2
3: 2 5 4 9 1 6 7 8	6: 7 2 8 5 4 9 1 3	9: 2 4 1 6 7 3 8 5

Fig. 5. An instance with 9 players and strictly ordered complete lists.

- **strict and complete lists, best roommate counts, 3-person rooms** (Theorem 4)
 The first dictator, player 1 chooses her first choice partner player 5 and they become a fixed pair. Then, the second dictator, player 2 chooses player 1, so these three form a fixed triplet. The third dictator cannot choose her first or second choice players 2 and 5, thus she becomes a fixed pair with player 4. Now it is exactly player 4 who is next to choose, but since she is already coupled up with her first choice roommate player 3, she does not change the current assignment. Player 5 is already in a fixed room. Player 6 opens the third room together with player 7, who then adds player 9 to this room, because her first two choices are already taken. Player 8 then joins the room of player 3. The outcome is the following partition: (1 2 5), (3 4 8), (6 7 9).
- **strict and complete lists, worst roommate counts, 3-person rooms** (Theorem 5)
 The first dictator, player 1 chooses her first and second choice partners player 5 and 4, and they occupy a room. The next dictator is player 2, whose first 3 choices are taken, thus she moves into a room with her two best available choices, players 9 and 8. The remaining 3 players are assigned to the last room. The outcome is the following partition: (1 4 5), (2 8 9), (3 6 7).

- **strict and complete lists, worst roommate counts, r_i -person rooms** (Theorem 5)

Let the rooms have capacity 2, 3, and 4, respectively. The first dictator, player 1 chooses her first choice partner player 5, and they occupy the smallest room. The next dictator is player 2, who cannot choose her first choice player 1, since her assignment is already fixed, thus she moves into the 3-person room with her best available choices, players 4 and 9. The remaining 4 players are assigned to the largest room. The outcome is the following partition: (1 5), (2 4 9), (3 6 7 8).

5.2 Hard cases

In all remaining cases, finding a POA is computationally infeasible, even though it is guaranteed to exist. We show this in two steps, similarly to the technique used by Aziz et al. [2]. First we observe that either all POAs of an instance or none of them satisfy the property that all players receive one of their best outcomes. Then, we show that an NP-complete problem can be reduced to the decision problem of answering whether a POA exists with this property. From this follows that by computing *any* POA in polynomial time, one could answer the NP-complete decision problem in polynomial time. Below we give three proofs for three different settings.

Theorem 6 (Finding a POA, strict and complete lists, best roommate counts, r_i -person rooms). *If lists are strict and complete, the rooms are of arbitrary capacity, and the best roommate counts, then computing a POA is at least as hard as the NP-complete UNARY BIN PACKING problem.*

Proof. We construct an instance of the POA problem to each input of UNARY BIN PACKING with item size at least 2. We will show that all POAs order each player to a room with a first ranked roommate if and only if a bin packing exists.

Each item of size $i \geq 2$ (recall Remark 2) in UNARY BIN PACKING corresponds to i players in the POA problem. The players of one item have their unique first choice player among themselves, in a circular manner. The rest of the preference lists can be chosen arbitrarily. To guarantee that all players have a first choice roommate, we need to keep every preference cycle together. This is equivalent to keeping the items of the bin packing problem unsplit, and thus it is possible if and only if there is a perfect bin packing.

This reduction from UNARY BIN PACKING with bin size b immediately implies that the proof is valid even if all rooms are of equal size b . \square

Theorem 7 (Finding a POA, ties, complete lists, best roommate counts, 3-person rooms). *If lists are weakly ordered and complete, the rooms are of capacity 3, and the best roommate counts, then computing a POA is at least as hard as the NP-complete DIRECTED TRIANGLE COVER problem.*

Proof. We construct an instance of the POA problem to each input of DIRECTED TRIANGLE COVER. Let us consider a digraph D , where all vertices have at least one outgoing and at least one incoming edge, as stated in Remark 1. Vertices in D correspond to players in our assignment problem. If a player i had an outgoing edge towards player j in D , then i ranks j first. All other players are ranked second by i . A directed triangle cover exists in D if and only if there is an assignment where each player has at least one first choice roommate. This latter happens if and only if all POAs have this property. \square

Theorem 8 (Finding a POA, ties, complete lists, worst roommate counts, 3-person rooms). *If lists are weakly ordered and complete, the rooms are of capacity 3, and the worst roommate counts, then computing a POA is at least as hard as the NP-complete TRIANGLE COVER problem.*

Proof. To each input G of the TRIANGLE COVER problem, we now construct an instance of the POA problem. Remember that according to Remark 1, G has no isolated or degree 1 vertex. Starting with G , let us assign rank 1 to all neighbors of each player. Now we complete G by adding all missing edges and assign rank 2 on both end vertices of such an edge.

We claim that there is a POA that gives every player two of her first ranked roommates if and only if a triangle cover exists in G . If a triangle cover exists in G , then it delivers an assignment consisting of original edges only, thus it is possible to assign each player into a room with first-choice roommates only. Since each player reaches her best outcome in this assignment, it also must be Pareto optimal. To see the other direction, we assume that there is a POA that orders each player to a room with only first ranked roommates. This assignment must then consist of the edges of G exclusively, and thus it is a triangle cover in G . \square

6 Conclusion

We have studied complexity issues in the Pareto optimal coalition formation problem in which players have preferences over each other, and the coalitions must be of a fixed size. We have investigated a number of variants of this problem and determined the complexity of verifying Pareto optimality, deciding the existence of a POA, and finding a POA.

One natural direction of future research is to forgo the requirement on the perfectness of the assignment. In this case, due to the feasibility of the empty assignment and monotonicity, a POA trivially exists, so our question raised in Section 4 about the existence of an optimal solution does not apply. However, allowing the assignment to be imperfect leads to unnatural strategies in serial dictatorship. If the worst roommate matters, then the dictator is better off choosing her single best roommate and not letting anyone else into the room, however large it is. Besides this, one needs to clarify how to deal with the option of staying alone in a large room, which was prohibited in our setting. Depending on the total capacity of rooms, and the way of calculating the value of an imperfect coalition, one can define several natural variants of the problem. We remark that our results in Sections 3 and 5 remain valid for imperfect assignments as well, if the total capacity of rooms equals the number of players, but some agents might be left unassigned—one needs to marginally adjust serial dictatorship, while the hardness proofs carry over.

Besides this, one might consider to investigate analogous problems with cardinal preferences instead of ordinal ones, an outside option for players, or strategic behaviour.

References

1. D. J. Abraham and D. F. Manlove. Pareto optimality in the Roommates problem. Technical Report TR-2004-182, University of Glasgow, Department of Computing Science, December 2004.
2. H. Aziz, F. Brandt, and P. Harrenstein. Pareto optimality in coalition formation. *Games and Economic Behavior*, 82:562–581, 2013.
3. H. Aziz, F. Brandt, and H. G. Seedig. Stable partitions in additively separable hedonic games. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 183–190. International Foundation for Autonomous Agents and Multiagent Systems, 2011.
4. C. Ballester. NP-completeness in hedonic games. *Games and Economic Behavior*, 49(1):1 – 30, 2004.
5. S. Banerjee, H. Konishi, and T. Sönmez. Core in a simple coalition formation game. *Social Choice and Welfare*, 18(1):135–153, 2001.
6. A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.

7. K. Cechlárová and J. Hajduková. Stability of partitions under \mathcal{WB} -preferences and \mathcal{BW} -preferences. *International Journal of Information Technology and Decision Making*, 3(4):605–618, 2004.
8. K. Cechlárová and J. Hajduková. Stable partitions with \mathcal{W} -preferences. *Discrete Applied Mathematics*, 138:333–347, 2004.
9. A. Darmann. Stable and Pareto optimal group activity selection from ordinal preferences. *International Journal of Game Theory*, 47(4):1183–1209, Nov 2018.
10. J. H. Dreze and J. Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica: Journal of the Econometric Society*, pages 987–1003, 1980.
11. W. A. Gamson. A theory of coalition formation. *American Sociological Review*, pages 373–382, 1961.
12. M. R. Garey and D. S. Johnson. *Computers and Intractability*. Freeman, San Francisco, CA., 1979.
13. C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*, volume 32. 01 1982.
14. J. P. Kahan and A. Rapoport. *Theories of coalition formation*. Psychology Press, 2014.
15. R. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
16. A. S. Kelso Jr and V. P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.
17. D. Knuth. *Mariages Stables*. Les Presses de L’Université de Montréal, 1976. English translation in *Stable Marriage and its Relation to Other Combinatorial Problems*, volume 10 of CRM Proceedings and Lecture Notes, American Mathematical Society, 1997.
18. T. Morrill. The roommates problem revisited. *Journal of Economic Theory*, 145(5):1739, 2010.
19. C. Ng and D. Hirschberg. Three-dimensional stable matching problems. *SIAM Journal on Discrete Mathematics*, 4:245–252, 1991.
20. D. Peters and E. Elkind. Simple causes of complexity in hedonic games. In *IJCAI*, volume 15, pages 617–623, 2015.
21. O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1):165–200, 1998.