

On Simplified Group Activity Selection

Andreas Darmann¹(✉), Janosch Döcker², Britta Dorn², Jérôme Lang³,
and Sebastian Schneckenburger²

¹ University of Graz, Graz, Austria

`andreas.darmann@uni-graz.at`

² University of Tübingen, Tübingen, Germany

³ Université Paris-Dauphine, Paris, France

Abstract. Several real-world situations can be represented in terms of agents that have preferences over activities in which they may participate. Often, the agents can take part in at most one activity (for instance, since these take place simultaneously), and there are additional constraints on the number of agents that can participate in an activity. In such a setting we consider the task of assigning agents to activities in a reasonable way. We introduce the simplified group activity selection problem providing a general yet simple model for a broad variety of settings, and start investigating the case where upper and lower bounds of the groups have to be taken into account. We apply different solution concepts such as envy-freeness and core stability to our setting and provide a computational complexity study for the problem of finding such solutions.

1 Introduction

Several real-world situations can be represented in terms of agents that have preferences over activities in which they may participate, subject to some feasibility constraints on the way they are assigned to the different activities. Here ‘activity’ should be taken in a wide sense; here are a few examples, each with its specificities which we will discuss further:

1. a group of co-workers may have to decide in which project to work, given that each project needs a fixed number of participants;
2. the participants to a big workshop, who are too numerous to fit all in a single restaurant, want to select a small number of restaurants (say, between two and four) out of a wider selection, with different capacities, and that serve different types of food, and to assign each participant to one of them;
3. a group of pensioners have to select two movies out of a wide selection, to be played simultaneously in two different rooms, and each of them will be able to see at most one of them;
4. a group of students have to choose one course each to follow out of a selection, given that each course opens only if it has a minimum number of registrants and has also an upper bound;

5. a set of voters want to select a committee of k representatives, given that each voter will be represented by one of the committee members.

While these examples seem to vary in several aspects, they share the same general structure: there is a set of *agents*, a set of available *activities*; each agent has preferences over the possible activities; there are constraints bearing on the selection of activities and the way agents are assigned to them; the goal is to assign each agent to one activity, respecting the constraints, and respecting as much as possible the agents' preferences.

Sometimes the set of selected activities is fixed (as Example 1), sometimes it will be determined by the agents' preferences. The nature of the constraints can vary: sometimes there are constraints that are *local* to each activity (typically, bounds on the number of participants, although we might imagine more complex constraints), as Examples 1, 2, 4, 5, and also 3 if the rooms have a capacity smaller than the number of pensioners; sometimes there are *global* constraints, that bear on the whole assignment (typically, bounds on the number of activities that can be selected; once again, we may consider more complex constraints), as in Examples 2, 3. Sometimes each agent *must* be assigned to an activity (as in Examples 1 and 5), sometimes she has the option of not being assigned to any activity.

This class of problems can be seen as a simplified version of the *group activity selection problem* (GASP), which asks how to assign agents to activities in a "good" way. In the original form introduced by Darmann et al. [5], agents express their preferences both on the activities and on the number of participants for the latter; in general, these preferences are expressed by means of weak orders over pairs "(activity, group size)". Darmann [4] considers the variant of GASP in which the agents' preferences are strict orders over such pairs and analyzes the computational complexity of finding assignments that are stable or maximize the number of agents assigned to activities.

Our model considers a simplified version of the group activity selection problem, called *s-GASP*. Here, agents only express their preferences over the set of activities. However, the activities come with certain constraints, such as restrictions on the number of participants, concepts like balancedness, or more global restrictions. The goal is again to find a "good" assignment of agents to activities, respecting both the agents' preferences as well as the constraints.

But what is a good assignment? Clearly, this essentially depends on the application on hand, but there are several concepts in the social choice and game theory literature that propose for an evaluative solution. We consider two classes of criteria for assessing the quality of an assignment:

- *solution concepts* that mainly come from game theory and that aim at telling whether an assignment is stable enough (that is, immune to some types of deviations) to be implemented. First, *individual rationality* requires that each agent is assigned to an activity she likes better than not being assigned to any activity at all. Then, a solution concept considered both in hedonic games, where coalition building is studied, and in matching theory, is the notion

of *stability*. It asks whether the assignment is stable in the sense that no agent would want to or be able to deviate from her coalition, her match, or in our case, her assigned activity. Besides considering different variants of *core stability*, it also makes sense in our setting to investigate variations of *virtual stability*, meaning that it is not possible that an agent deviates from her assigned activity due to the given constraints.

- *criteria* that mainly come from social choice and that measure, qualitatively or quantitatively, the welfare of agents. A common quality measure in terms of efficiency of an assignment is the notion of *Pareto optimality*: there should be no feasible assignment in which there is an agent that is strictly better off, while the remaining agents do not change for the worse. More generally, one may wish to *optimize social welfare*, for some notion of utility derived from the agents' preferences: for instance, one may simply be willing to maximize the number of agents assigned to an activity. If fairness is important in the design, the notion of *envy-freeness* makes sense: an assignment respecting the constraints is envy-free if no agent strictly prefers the group another agent is assigned to.

Related Work. Apart from GASP, our model is related to various streams of work:

Course allocation, e.g. [2,6,10,14]. Students bear preferences over courses they would like to be enrolled in (these preferences are typically strict orders), and there are typically constraints given on the size of the courses. Courses will only be offered if a minimum number of participants is found, and there are upper bounds due to space or capacity limitations. In particular, Cechlárová and Fleiner [2] consider a course-allocation framework, so for them it makes sense that one agent can be matched to more than one activity (course), while [10,14] consider the case in which an agent can be assigned to at most one activity (project). The latter works are very close to our setting with constraints over group sizes. In contrast to above works however, our setting contains a dedicated outside option (the *void activity*), and agents' preferences are represented by weak orders over activities instead of strict rankings.

Hedonic games (see the recent survey by Aziz and Savani [1]) are coalition formation games where each agent has preferences over coalitions containing her. The stability notions we will focus on are derived from those for hedonic games. However, in our model, agents do not care about who else is assigned to the same activity as them, but only on the activity to which they are assigned to.¹

¹ Still, it is possible to express simplified group activity selection within the setting of hedonic games, by adding special agents corresponding to activities, who are indifferent between all locally feasible coalitions. See the work by Darmann et al. [5] for such a translation for the more general group activity selection problem. But it is a rather artificial, and overly complex, representation of our model, which moreover does not help characterizing and computing solution concepts.

In *multiwinner elections*, there is a set of candidates, voters have preferences over single candidates, and a subset of k candidates has to be elected. In some approaches to multiwinner elections, each voter is assigned to one of the members of the elected committee, who is supposed to represent her. Sometimes there are no constraints on the number of voters assigned to a given committee member (as is the case for the *Chamberlin-Courant* rule [3]), in which case each voter is assigned to her most preferred committee member; on the other hand, for the *Monroe* rule [13], the assignment has to be balanced. A more general setting, with more general constraints, has been defined by Skowron et al. [16]. Note also that multiwinner elections can also be interpreted as *resource allocation* with items that come in several units (see again [16]) and as *group recommendation* [12]. While assignment-based multiwinner elections problems are similar to simplified group activity, an important difference is that for the former, stability notions play no role, as the voters are not assumed to be able to deviate from their assigned representatives.

Contents and Outline. In this work, we will take into account various solution concepts and ask two questions: First, do “good” assignments exist? Can we decide this efficiently? And if they exist, can we find them efficiently? Our second concern is optimization: we are looking for desirable assignments that maximize the number of agents which can be assigned to an activity. Again, we may ask whether an assignment that is optimal in this sense exists, and we can try to find it.

We will focus on one family of constraints concerning the size of the groups—we assume that each activity comes with a lower and an upper bound on the number of participants—and give a detailed analysis of the described problems for this class.

Our results for this class are twofold. First, we show that it is often possible to find assignments with desirable properties in an efficient way: we propose several polynomial time algorithms to find good assignments or to optimize them. We complement these findings with NP- and coNP-completeness results for certain solution concepts. Whenever we encounter computational hardness, we identify tractable special cases: we will see that all our problems can be solved in polynomial time if there is no restriction on the minimum number of participants for the activities to take place. An overview of our computational complexity results is given in Table 1 in Sect. 3; due to space constraints, we do not elaborate all proofs. Second, we show that also in this class of problems considered, there is a certain tension between the concepts of envy-freeness and Pareto-optimality, even for small instances.

The remainder of this work is organized as follows. In Sect. 2, we formally introduce the simplified model as well as possible constraints and several solution concepts. Section 3 is the main part of the paper and provides an analysis of the computational complexity of the questions described above. Section 4 deals with the tension between envy-freeness and Pareto optimality. In Sect. 5, we conclude and discuss future directions of research connected to s-GASP.

2 Model, Constraints, and Solution Concepts

We start with defining our model and with introducing the solution concepts we want to consider.

Simplified Group Activity Selection, Constraints. An instance (N, A, P, R) of the *simplified group activity selection problem* (s-GASP) is given as follows. The set $N = \{1, \dots, n\}$ denotes a set of agents and $A = A^* \cup \{a_\emptyset\}$ a set of activities with $A^* = \{a_1, \dots, a_m\}$, where a_\emptyset stands for the *void activity*. An agent who is assigned to a_\emptyset can be thought of as not participating in any activity. The preference profile $P = \langle \succsim_1, \dots, \succsim_n \rangle$ consists of n votes (one for each agent), where \succsim_i is a weak order over A for each $i \in N$. The set R is a set of side constraints that restricts the set of assignments.

A mapping $\pi : N \rightarrow A$ is called an *assignment*. Given assignment π , $\#(\pi) = |\{i \in N : \pi(i) \neq a_\emptyset\}|$ denotes the number of agents π assigns to a non-void activity; for activity $a \in A$, $\pi^a := \{i \in N : \pi(i) = a\}$ is the set of agents π assigns to a .

The goal will be to find “good” assignments that satisfy the constraints in R . The structure of the set R depends on the application. Some typical kinds of constraints are (combinations of) the following cases:

1. each activity comes with a lower and/or upper bound on the number of participants;
2. no more than k activities can have some agent assigned to them;
3. the number of voters per activity should be balanced in some way;

Intuitively, if there are no constraints or the constraints are flexible enough, then agents go where they want and the problem becomes trivial. If the constraints are tight enough (e.g., perfect balancedness, provided $|A|$ and $|V|$ allow it), then some agents are generally not happy, but they are unable to deviate because most deviations violate the constraints. The interesting cases can therefore be in between these two extreme cases.

In this work, we will start investigations for s-GASP for the first class of constraints: We assume that each activity $a \in A^*$ comes with a lower bound $\ell(a)$ and an upper bound $u(a)$, and all constraints in R are of the following type: for each $a \in A^*$, $|\pi^a| \in \{0\} \cup [\ell(a), u(a)]$.

Feasible Assignments, Solution Concepts. Let an instance (N, A, P, R) of s-GASP be given. A *feasible assignment* is an assignment meeting the constraints in R . We will consider the following properties. A feasible assignment π is

- *envy-free* if there is no pair of agents $(i, j) \in N \times N$ with $\pi(j) \in A^*$ such that $\pi(j) \succ_i \pi(i)$ holds;
- *individually rational* if for each $i \in N$ we have $\pi(i) \succsim_i a_\emptyset$;

- *individually stable* if there is no agent i and no activity $a \in A$ such that (i) $a \succ_i \pi(i)$ and (ii) the mapping π' defined by $\pi'(i) = a$ and $\pi'(k) = \pi(k)$ for $k \in N \setminus \{i\}$ is a feasible assignment;
- *core stable* if there is no set $E \subseteq N$ and no activity $a \in A$ such that (i) $a \succ_i \pi(i)$ for all $i \in E$, (ii) $\pi^a \subset E$ holds if $a \in A^*$, and (iii) the mapping π' defined by $\pi'(i) = a$ for $i \in E$ and $\pi'(k) = \pi(k)$ for $k \in N \setminus E$ is a feasible assignment; (Note that the respective activity a to which the set E of agents wishes to deviate must be either a_\emptyset or currently unused.)
- *strictly core stable* if there is no set $E \subseteq N$ and no activity $a \in A$ such that (i) $a \succsim_i \pi(i)$ for all $i \in E$ where $a \succ_i \pi(i)$ for at least one $i \in E$, (ii) $\pi^a \subset E$ holds if $a \in A^*$, and (iii) the mapping π' defined by $\pi'(i) = a$ for all $i \in E$ and $\pi'(k) = \pi(k)$ for $k \in N \setminus E$ is a feasible assignment;
- *Pareto optimal* if there is no feasible assignment $\pi' \neq \pi$ such that $\pi'(i) \succsim_i \pi(i)$ for all $i \in N$ and $\pi'(i) \succ_i \pi(i)$ for at least one $i \in N$;

Finally, an individually rational assignment π is *maximum individually rational* if for all individually rational assignments π' we have $\#(\pi) \geq \#(\pi')$. Analogously, maximum feasible/envy-free/.../Pareto optimal assignments are defined.

For the class of constraints we consider, the notion of *virtual stability* is interesting. It requires that any deviation from the assigned towards a more preferred activity $a \in A^*$ violates the capacity constraints of a . Formally, we define the following stability concepts.

A feasible assignment π is

- *virtually individually stable* if there is no agent i and no activity $a \in A$ with $\ell(a) \leq |\pi^a| + 1 \leq u(a)$ such that $a \succ_i \pi(i)$ holds;
- *virtually core stable* if there is no set $E \subseteq N$ and no activity $a \in A$ with $\ell(a) \leq |E| \leq u(a)$ such that $a \succ_i \pi(i)$ for all $i \in E$, and (ii) $\pi^a \subset E$ holds if $a \in A^*$;
- *virtually strictly core stable* if there is no set $E \subseteq N$ and no activity $a \in A$ with $\ell(a) \leq |E| \leq u(a)$ such that (i) $a \succsim_i \pi(i)$ for all $i \in E$ where $a \succ_i \pi(i)$ for at least one $i \in E$, and (ii) $\pi^a \subset E$ holds if $a \in A^*$.

Note that as in the definition of core stability, also in virtual core stability the respective activity a to which the set E of agents wishes to deviate must be either a_\emptyset or currently unused.

The relationships between the solution concepts is shown in Fig. 1 (for an overview of the relationships between solution concepts in hedonic games we refer to [1]).

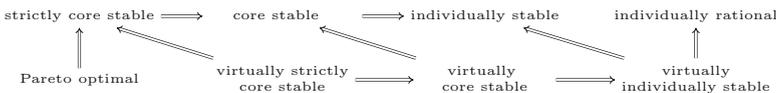


Fig. 1. Relations between the solution concepts we consider.

3 Computational Complexity for s-GASP with Group Size Constraints

We will now consider the computational complexity of s-GASP for various solution concepts. An overview of our results is given in Table 1.

Table 1. Overview of results for constraints $|\pi^a| \in \{0\} \cup [\ell(a), u(a)]$, $a \in A^*$.

Find assignment that is	general	$u(a) = n$	$\ell(a) = 1$
Feasible	in P (Proposition 1)	in P (Proposition 1)	in P (Proposition 1)
Individually rational	in P (Theorem 2)	in P (Theorem 2)	in P (Corollary 2)
Envy-free	in P (trivial)	in P (trivial)	in P (trivial)
Individually stable	in P (Theorem 1)	in P (Theorem 1)	in P (Corollary 2)
Core stable	in P (Theorem 1)	in P (Theorem 1)	in P (Corollary 2)
Strictly core stable	in P (Theorem 1)	in P (Theorem 1)	in P (Corollary 2)
Virtually individually stable	in P (Theorem 2)	in P (Theorem 2)	in P (Corollary 2)
Virtually core stable	NP-c (Corollary 1)	NP-c (Corollary 1)	in P (Corollary 2)
Virtually strictly core stable	NP-c (Theorem 3)	NP-c (Theorem 3)	in P (Corollary 2)
Pareto optimal	?	?	in P (Theorem 7)
Is there an assignment π with $\#(\pi) \geq k$ ($k \in \mathbb{N}$) that is	general	$u(a) = n$	$\ell(a) = 1$
Feasible	in P (Proposition 1)	in P (Proposition 1)	in P (Proposition 1)
Individually rational	NP-c (Theorem 4)	NP-c (Theorem 4)	in P (Theorem 5)
Envy-free	NP-c (Theorem 6)	in P (trivial)	?
Virtually individually stable	NP-complete	NP-complete	in P (Corollary 2)
Virtually core stable	NP-c (Corollary 1)	NP-c (Corollary 1)	in P (Corollary 2)
Virtually strictly core stable	NP-c (Theorem 3)	NP-c (Theorem 3)	in P (Corollary 2)
Pareto optimal	?	?	in P (Theorem 7)
Given assignment π , is π PO?	coNP-c (Theorem 8)	coNP-c (Theorem 8)	in P (Theorem 9)

3.1 Finding “Good” Assignments

The first interesting question is whether “good” assignments exist and how to find them. Obviously, assigning the void activity to every agent results in a feasible, individually rational and envy-free assignment. However, this is not a satisfying solution in terms of stability because agents will want to deviate. The good news is that for several stability concepts, a corresponding assignment always exists and can efficiently be found, as shown in the following theorem.

Theorem 1. *A strictly core stable assignment always exists and can be found in polynomial time.*

Proof. We sketch the basic algorithmic idea. Starting with a feasible assignment π , for each agent i and each activity b which i prefers to $\pi(i)$ we check whether there is a subset of agents including agent i that wants to deviate to b such that the resulting assignment is feasible. That is, we check whether there is a subset $E \supset \pi^b$ such that (i) for all $j \in E$ we have that $b \succ_j \pi(j)$ holds (recall that for agent i $b \succ_i \pi(i)$ holds) and (ii) π' with $\pi'(i) = b$ for $i \in E$ and $\pi'(j) = \pi(j)$ for $j \in N \setminus E$ is a feasible assignment. In order to do so, for each activity $c \in A \setminus \{b\}$, we compute the possible numbers of agents in the set π^c that agree with joining b and can be removed from π^c while still enabling a feasible assignment. Finally, given these numbers, we need to verify if—including i and the agents in π^b —these add up to an integer contained in $[\ell(b), u(b)]$ by taking exactly one number from each activity. The latter problem reduces to the MULTIPLE-CHOICE SUBSET-SUM problem (see Pisinger [15]), which, in our case, allows for an overall polynomial time algorithm for finding a strictly core stable assignment. ■

Recall that a strictly core stable assignment is also core stable and individually stable. Hence, as a consequence of the above theorem, also a core stable and an individually stable assignment always exist.

Theorem 2. *A virtually individually stable assignment always exists and can be found in polynomial time.*

Proof. In an instance (N, A, P, R) of s-GASP, we initially assign each agent to a_\emptyset , i.e., set $\pi(i) := a_\emptyset$ for $i \in N$. For $a \in A^*$ with $\ell(a) \geq 2$, if no agent is assigned to such a , then $\ell(a) \leq |\pi^a| + 1$ cannot hold. Hence, in what follows, we only consider activities $a \in A^*$ with $\ell(a) = 1$. For $1 \leq i \leq n$, assign agent i to the best ranked such activity $a \succ_i a_\emptyset$ with $|\pi^a| < u(a)$ and update π (i.e., set $\pi(i) := a$ while $\pi(j)$ remains unchanged for $j \in N \setminus \{i\}$). It is easy to see that the resulting assignment π is virtually individually stable. ■

In contrast, a virtually core stable (and thus a virtually strictly core stable) assignment does not always exist, as the following example shows.

Example 1. Let $N = \{1, 2, 3\}$ and $A^* = \{a, b, c\}$, with $a \succ_1 b \succ_1 c \succ a_\emptyset$, $b \succ_2 c \succ_2 a \succ a_\emptyset$, and $c \succ_3 a \succ_3 b \succ a_\emptyset$. The restrictions on the activities are given by $|\pi^x| \in \{0\} \cup [2, 3]$, for each $x \in A^*$. By the restrictions given, there is at most one non-void activity to which agents can be assigned. Clearly, for any activity $z \in A$ there is a $y \in A^*$ such that two agents prefer y to z . As a consequence, there can be no virtually core stable assignment.

In addition, the problem to decide whether or not a virtually strictly core stable assignment exists turns out to be computationally difficult.

Theorem 3. *It is NP-complete to decide if there is a virtually strictly core stable assignment, even when for each activity $a \in A^*$ we have $u(a) = n$.*

Proof. Membership in NP is not difficult to verify. The proof proceeds by a reduction from EXACT COVER BY 3-SETS (X3C). The input of an instance of X3C consists of a pair $\langle X, \mathcal{Z} \rangle$, where $X = \{1, \dots, 3q\}$ and $\mathcal{Z} = \{Z_1, \dots, Z_p\}$ is a collection of 3-element subsets of X ; the question is whether we can cover X with exactly q sets of \mathcal{Z} . X3C is known to be NP-complete even when each element of X is contained in exactly three sets of \mathcal{Z} (see [7, 8]); note that in such a case $p = 3q$ holds. For each $i \in X$, let the sets containing i be denoted by $Z_{i_1}, Z_{i_2}, Z_{i_3}$ with $i_1 < i_2 < i_3$.

Define instance $\mathcal{I} = (N, A, P, R)$ of s-GASP as follows. Let $N = \{V_{i,1}, V_{i,2}, V_{i,3} \mid 1 \leq i \leq p\}$ and $A^* = \{y_i, a_i, b_i, c_i \mid 1 \leq i \leq p\}$. For $1 \leq i \leq p$, let $\ell(a_i) = \ell(b_i) = \ell(c_i) = 2$, and $\ell(y_i) = 9$. For each $a \in A^*$, let $u(a) = |N|$. Since any virtually strictly core stable assignment is individually rational, in the profile P we omit the activities ranked below a_0 ; for each $i \in \{1, \dots, p\}$, let the ranking of the agents $V_{i,1}, V_{i,2}, V_{i,3}$ (each of which represents element $i \in X$) be given as follows:

$$\begin{aligned} V_{i,1} : & y_{i_1} \succ_{i,1} y_{i_2} \succ_{i,1} y_{i_3} \succ_{i,1} a_i \succ_{i,1} b_i \succ_{i,1} c_i \succ_{i,1} a_0 \\ V_{i,2} : & y_{i_2} \succ_{i,2} y_{i_3} \succ_{i,2} y_{i_1} \succ_{i,2} b_i \succ_{i,2} c_i \succ_{i,2} a_i \succ_{i,2} a_0 \\ V_{i,3} : & y_{i_3} \succ_{i,3} y_{i_1} \succ_{i,3} y_{i_2} \succ_{i,3} c_i \succ_{i,3} a_i \succ_{i,3} b_i \succ_{i,3} a_0 \end{aligned}$$

Note that each set Z contains three elements, and hence each y_i , $1 \leq i \leq p$, is preferred to a_0 by exactly 9 agents. We show that there is an exact cover in instance $\langle X, \mathcal{Z} \rangle$ if and only if there is a virtually strictly core stable assignment in instance \mathcal{I} .

Assume there is an exact cover C . Consider the assignment π defined by $\pi(V_{i,h}) = y_j$ if $i \in Z_j$ and $Z_j \in C$, for $i \in \{1, \dots, p\}$ and $h \in \{1, 2, 3\}$. Since C is an exact cover, assignment π is well-defined and feasible; note that each agent is assigned to an activity she ranks first, second or third. In addition, note that for $Z_j \in C$, each agent that prefers y_j to a_0 is assigned to y_j . Assume a set of agents E wishes to deviate to another activity d , such that at least one member $i \in E$ prefers d over $\pi(i)$ while there is no $j \in E$ with $\pi(j) \succ_j d$. By the definition of π , $d \in \{y_i \mid 1 \leq i \leq p\}$ holds. Observe that $\pi^d = \emptyset$ holds because C is an exact cover. Due to $\ell(d) = 9$, it hence follows that each agent of those who prefer d to a_0 must prefer d to the assigned activity, which is impossible since, by construction of the instance, for at least one of these agents j the assigned activity is top-ranked, i.e., $\pi(j) \succ_j d$ holds. Therewith, π is virtually strictly core stable.

Conversely, assume there is a virtually strictly core stable assignment π . Assume that there is an agent $V_{i,h}$ who is not assigned to one of the activities $y_{i_1}, y_{i_2}, y_{i_3}$. Then, by $\ell(y_i) = 9$ and the fact that exactly 9 agents prefer y_i to a_0 for each $i \in \{1, \dots, p\}$, it follows that no agent is assigned to one of $y_{i_1}, y_{i_2}, y_{i_3}$; in particular none of $V_{i,1}, V_{i,2}, V_{i,3}$ is assigned to one of these activities. Analogously to Example 1 it then follows that there is no virtually strictly core stable assignment, in contradiction with our assumption.

Thus, π assigns each agent $V_{i,h}$ to one of the activities $y_{i_1}, y_{i_2}, y_{i_3}$. For each $i \in \{1, \dots, p\}$, by $\ell(y_i) = 9$ and the fact that exactly 9 agents prefer y_i to a_0

it follows that to exactly one of $y_{i_1}, y_{i_2}, y_{i_3}$ exactly 9 agents are assigned, while no agent is assigned to the remaining two activities. As a consequence, the set $C = \{Z_i \mid |\pi^{y_i}| = 9, 1 \leq i \leq p\}$ is an exact cover in instance $\langle X, Z \rangle$. ■

In the instance considered in the above proof, an assignment is virtually strictly core stable if and only if it is virtually core stable. As a consequence, we get the following corollary.

Corollary 1. *It is NP-complete to decide if there is a virtually core stable assignment, even if for each activity $a \in A^*$ we have $u(a) = n$.*

However, for the case of $\ell(a) = 1$ for each $a \in A^*$, we get a positive complexity result (see Sect. 3.2). In particular, we can show that in this case a virtually strictly core stable assignment that maximizes the number of agents assigned to a non-void activity can be found in polynomial time.

Turning to Pareto optimality, in the special case of $\ell(a) = 1$ for each $a \in A^*$, there is a simple algorithm to compute a Pareto optimal assignment. In that case, it is easy to see that a Pareto optimal assignment is always individually rational. Thus, neglecting activities ranked below a_\emptyset , we start with the assignment $\pi(i) = a_\emptyset$ for each $i \in N$ and iteratively assign an agent to the best-ranked among the activities a with $|\pi^a| < u(a)$. However, in the case of $\ell(a) = 1$ for each $a \in A^*$ we can even find a Pareto optimal assignment that maximizes the number of agents assigned to a non-void activity in polynomial time (see Sect. 3.2).

3.2 Maximizing the Number of Agents Assigned to a Non-void Activity

We now turn to an optimization problem: Among all feasible assignments that feature a certain property, one is usually interested in finding one that maximizes the number of agents that are assigned to a non-void activity, thus keeping the number of agents who cannot be enrolled in any activity low.

Proposition 1. *In polynomial time we can find a feasible assignment that maximizes the number of agents assigned to a non-void activity.*

But already for individual rational assignments, it is hard to decide whether all agents can be assigned to a non-void activity, as the following theorem shows. We omit its proof which is again a reduction from the EXACT COVER BY 3-SETS problem.

Theorem 4. *It is NP-complete to decide if there is an individually rational assignment that assigns each agent to some $a \in A^*$, even if for each activity $a \in A^*$ we have $u(a) = n$.*

However, if we assume that each activity admits a group size of 1, then we can find an optimal individually rational assignment efficiently.

Theorem 5. *If for each activity $a \in A^*$ we have $\ell(a) = 1$, then in polynomial time we can find a maximum individually rational assignment.*

Proof. Reduction to max integer flow with upper bounds. Given an instance $\mathcal{I} = (N, A, P, R)$ of \mathfrak{s} -GASP with $\ell(a) = 1$ for all $a \in A^*$, we construct an instance \mathcal{M} of max integer flow with directed graph $G = (V, E)$. Set $V := \{s, t\} \cup N \cup A^*$, and let the edges and their capacities be given as follows: for each $i \in N$, introduce edge (s, i) with capacity 1; for each $a \in A^*$ and $i \in N$ introduce an edge (i, a) of capacity 1 if $a \succ_i a_\emptyset$ holds; for each $a \in A^*$, introduce edge (a, t) of capacity $u(a)$. It is easy to see that a max integer flow from s to t induces a maximum individually rational assignment in \mathcal{I} and vice versa. ■

For envy-freeness, optimizing the number of “active” agents turns again out to be a hard problem which can be shown by a reduction from EXACT COVER BY 3-SETS as well.

Theorem 6. *It is NP-complete to decide if there is an envy-free assignment that assigns each agent to some $a \in A^*$.*

We obtain tractability for envy-freeness if we loosen the constraints on the upper bounds of the group sizes: Clearly, if there is an activity with “unlimited” capacity (i.e., its upper bound equals n), we can assign all agents to it and obtain envy-freeness.

3.3 Pareto Optimality

In this subsection, we consider the computational complexity involved in Pareto optimal assignments.

In the framework of course allocation, if all agents have strict preferences it is known that a Pareto optimal matching—that assigns an agent to an activity (course) only if the activity is acceptable for the agent—can be found in polynomial time (see [2, 10]). Since in our setting (i) the agents’ preferences are represented by weak orders and (ii) Pareto optimality does not require individual rationality, these results do not immediately translate. For the latter reason, the computational intractability result of [2] for finding a Pareto optimal matching maximizing the number of agents assigned to a non-void activity if each agent can be assigned to at most one activity does not immediately translate to our setting either. In particular, in general we do not know the computational complexity status of finding a Pareto optimal assignment (or of finding one that maximizes the number of agents assigned to non-void activities) in \mathfrak{s} -GASP. As the following theorem shows, the latter issue is computationally tractable if we relax the constraint on the lower bound of the group sizes.

Theorem 7. *If for each activity $a \in A^*$ we have $\ell(a) = 1$, then in polynomial time we can find a Pareto optimal assignment that maximizes the number of agents assigned to a non-void activity.*

Proof. In that case, any Pareto optimal assignment is individually rational. Let k be the maximum number of agents assigned to non-void activities by an individually rational assignment. Hence, it is sufficient to find a Pareto optimal

assignment π with $\#(\pi) = k$. Given an instance $\mathcal{I} = (N, A, P, R)$ of s-GASP with $\ell(a) = 1$ for all $a \in A^*$, we construct an instance \mathcal{F} of the minimum cost flow problem. Instance \mathcal{F} corresponds to instance \mathcal{M} of the proof of Theorem 5 except that we add the following edge costs: for each $a \in A^*$ and $i \in N$ edge (i, a) has cost $-(1 + |\{b \in A^* | a \succ_i b, b \succ_i a_\emptyset\}|)$, all remaining edges have zero cost. Let f be a minimum integer cost flow of size k in instance \mathcal{F} . Then f induces the assignment π by setting $\pi(i) = a$ iff f sends a unit of flow through edge (i, a) . Clearly, π is Pareto optimal since otherwise a flow f' of lower total cost than the total cost of f could be induced. ■

Note that in the case $\ell(a) = 1$ for each $a \in A^*$, also any strictly core stable, core stable, or individually stable assignment is individually rational. In addition, in this case virtually (strict) core stability coincides with (strict) core stability, and virtually individually stability coincides with individual stability. Hence we can state the following corollary.

Corollary 2. *If for each activity $a \in A^*$ we have $\ell(a) = 1$, then in polynomial time we can find a maximum individually rational assignment that is Pareto optimal, (virtually) individually stable, (virtually) core stable and (virtually) strictly core stable.*

However, checking whether a given assignment is Pareto optimal turns out to be coNP-complete, as Theorem 8 shows. We omit the proof which makes use of the NP-completeness of X3C.

Theorem 8. *It is coNP-complete to decide if a given assignment is Pareto optimal, even if for each activity $a \in A^*$ we have $u(a) = n$.*

Again, if there are no restrictions on the minimum number of participants of each activity, the latter problem becomes tractable.

Theorem 9. *If for each activity $a \in A^*$ we have $\ell(a) = 1$, then in polynomial time we can decide if a given assignment is Pareto optimal.*

Proof. Given instance $\mathcal{I} = (N, A, P, R)$ of s-GASP with $\ell(a) = 1$ for all $a \in A^*$ and assignment π , we construct instance \mathcal{C} of the minimum cost flow problem as follows with lower and upper edge capacities. Note that π must be individually rational. In instance \mathcal{C} , the directed graph $G = (V, E)$, edge costs and capacities are given as follows. $G = (V, E)$ has vertex set $V := \{s, t\} \cup N \cup A^*$, the edge set E consists of the following edges:

- for $i \in N$, edge (s, i) of zero cost, and, for $a \in A^*$ with $a \succ_i \pi(i)$, edge (i, a) of cost -1 if $a \succ_i \pi(i)$ and of cost 0 if $a \sim_i \pi(i)$;
- for $a \in A^*$ edge (a, t) of upper capacity bound $u(a)$.

The lower and upper capacity bound of edge (s, i) is 1 iff $\pi(i) \succ a_\emptyset$ holds. Unless otherwise specified, the lower capacity bound of edge $e \in E$ is 0 and the upper capacity bound is 1, and its cost is 0.

Assume there is an integer flow f of negative total cost. Consider the assignment π' defined by $\pi'(i) = a$ iff f sends flow through edge (i, a) . Then, by construction we must have $\pi'(i) \sim_i \pi(i)$ or $\pi'(i) \succ_i \pi(i)$ for each $i \in N$, where the latter holds for at least one agent $i \in N$ by the negative total cost of f . Thus, π is not Pareto optimal.

If, on the other hand, π is not Pareto optimal, then there is an assignment π' with $\pi'(i) \sim_i \pi(i)$ or $\pi'(i) \succ_i \pi(i)$ for each $i \in N$, where the latter holds for at least one $i \in N$. The integer flow f' that sends flow along the edges $(s, i), (i, a), (a, t)$ iff $\pi'(i) = a$ holds, has negative total cost.

Therewith, for verifying if π is Pareto optimal it is sufficient to find an integer minimum cost flow in instance \mathcal{C} . ■

4 Envy-Freeness vs. Pareto Optimality

In many social choice settings, there is a tension between envy-freeness and Pareto optimality. This is also the case for our simplified group activity selection problem, as the following proposition and the subsequent corollary show.

Proposition 2. *For any $k \geq 2$, there is an instance (N, A, P, R) of s-GASP with $|N| = k$ and $\ell(a) = 1$ for each $a \in A^*$, for which there does not exist an assignment π which is both Pareto optimal and envy-free.*

Proof. We provide a proof for $k = 2$, which easily extends to $n = k$ for any $k > 2$. Consider the instance with $N = \{1, 2\}$, $A^* = \{a\}$, with the rankings $a \succ_1 a_\emptyset$ and $a \succ_2 a_\emptyset$, and the restrictions given by $\ell(a) = u(a) = 1$. Any Pareto optimal assignment assigns exactly one agent to a , which is clearly not envy-free. ■

Corollary 3. *There is no mechanism that determines an assignment that is both Pareto optimal and envy-free for each given instance (N, A, P, R) of s-GASP, even if $\ell(a) = 1$ holds for each $a \in A^*$.*

Interestingly, this tension also holds if the only relevant constraint is the lower bound of the activities (i.e., $u(a) = n$ for all a).

Proposition 3. *For any $k \geq 6$, there is an instance (N, A, P, R) of s-GASP with $|N| = k$ and $u(a) = k$ for each $a \in A^*$, for which there does not exist an assignment π which is both Pareto optimal and envy-free.*

Proof. We provide a proof for $k = 6$, which easily extends to $k = n$ for any $n > 6$. Consider the instance of s-GASP with $N = \{1, 2, 3, 4, 5, 6\}$, $A^* = \{a, b, c\}$ and for any $x \in A^*$ we have $\ell(x) = 3, u(x) = 6$. The rankings are

$$\begin{array}{ll}
 \succsim_1: a \succ_1 b \succ_1 c \succ_1 a_\emptyset & \succsim_4: a \succ_4 b \succ_4 c \succ_4 a_\emptyset \\
 \succsim_2: b \succ_2 c \succ_2 a \succ_2 a_\emptyset & \succsim_5: b \succ_5 c \succ_5 a \succ_5 a_\emptyset \\
 \succsim_3: c \succ_3 a \succ_3 b \succ_3 a_\emptyset & \succsim_6: c \succ_6 a \succ_6 b \succ_6 a_\emptyset
 \end{array}$$

Due to the feasibility constraints, there are only 4 types of feasible assignments:

- (i) 3–5 agents are assigned to the same activity $x \neq a_\emptyset$, and the rest to a_\emptyset .
- (ii) All agents are assigned to the void activity.
- (iii) All agents are assigned to the same activity $x \neq a_\emptyset$.
- (iv) 3 agents are assigned to the same activity $x \neq a_\emptyset$ and the other 3 agents are assigned to another activity $y \notin \{x, a_\emptyset\}$.

The assignments of type (i) and (ii) are Pareto dominated by some assignment of type (iii). An assignment π_1 of type (iii) is envy-free but not Pareto optimal. Due to the symmetrical construction of the preferences profiles, we can assume without loss of generality $\pi_1^a = N$. But then the assignment is Pareto dominated by the assignment π_2 with $\pi_2^a = \{1, 3, 4\}$ and $\pi_2^b = \{2, 5, 6\}$. An assignment of type (iv) cannot be envy-free. Without loss of generality we can assume $x = a$ and $y = b$. Assume, for the sake of contradiction, that there is an envy-free assignment. Agents 1 and 4 must be assigned to activity a and agents 2 and 5 to activity b . As the preference profiles of the remaining agents both rank a strictly better than b , the assignment cannot be an envy-free assignment. ■

Corollary 4. *There is no mechanism that determines an assignment that is both Pareto optimal and envy-free for each given instance (N, A, P, R) of s-GASP, even if $u(a) = n$ holds for each $a \in A^*$.*

5 Conclusion

We have formulated a simplified version of GASP where the assignment of agents to activities depends on the agents' preferences as well as on exogenous constraints. This model is powerful enough to capture many real world applications. We have made a first step by analyzing one family of constraints and have studied several solution concepts for this family.

An obvious next step is to drive a similar analysis for other interesting classes of constraints as described in Sect. 2. In particular, it would be interesting to characterize families of constraints guaranteeing or not guaranteeing existence of a stable solution for the different solution concepts we considered, or exploring forbidden structures that prevent stability. Also, it would be nice to provide a detailed analysis of the parameterized complexity of the hard cases, as done by Lee and Williams [11] for the stable invitation problem and by Igarashi et al. [9] for GASP on social networks. Another variant would be to consider typed agents as in the paper by Spradling and Goldsmith [17].

Acknowledgement. This work was partly supported by COST Action IC1205 on Computational Social Choice.

References

1. Aziz, H., Savani, R.: Hedonic games. In: Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.) Handbook of Computational Social Choice. Cambridge University Press (2015). Chap. 15

2. Cechlárová, K., Fleiner, T.: Pareto optimal matchings with lower quotas. *Math. Soc. Sci.* **88**, 3–10 (2017)
3. Chamberlin, J.R., Courant, P.N.: Representative deliberations and representative decisions: proportional representation and the Borda rule. *Am. Polit. Sci. Rev.* **77**(3), 718–733 (1983)
4. Darmann, A.: Group activity selection from ordinal preferences. In: Walsh, T. (ed.) ADT 2015. LNCS (LNAI), vol. 9346, pp. 35–51. Springer, Cham (2015). doi:[10.1007/978-3-319-23114-3_3](https://doi.org/10.1007/978-3-319-23114-3_3)
5. Darmann, A., Elkind, E., Kurz, S., Lang, J., Schauer, J., Woeginger, G.: Group activity selection problem. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 156–169. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-35311-6_12](https://doi.org/10.1007/978-3-642-35311-6_12)
6. Goto, M., Hashimoto, N., Iwasaki, A., Kawasaki, Y., Ueda, S., Yasuda, Y., Yokoo, M.: Strategy-proof matching with regional minimum quotas. In: Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, pp. 1225–1232. International Foundation for Autonomous Agents and Multiagent Systems (2014)
7. Hein, J., Jiang, T., Wang, L., Zhang, K.: On the complexity of comparing evolutionary trees. *Discrete Appl. Math.* **71**, 153–169 (1996)
8. Hickey, G., Dehne, F., Rau-Chaplin, A., Blouin, C.: SPR distance computation for unrooted trees. *Evol. Bioinform.* **4**, 17–27 (2008)
9. Igarashi, A., Bredereck, R., Elkind, E.: On parameterized complexity of group activity selection problems on social networks (2017). CoRR abs/1703.01121
10. Kamiyama, N.: A note on the serial dictatorship with project closures. *Oper. Res. Lett.* **41**(5), 559–561 (2013)
11. Lee, H., Williams, V.V.: Parameterized complexity of group activity selection. In: Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017), IFAAMAS, May 2017
12. Lu, T., Boutilier, C.: Budgeted social choice: From consensus to personalized decision making. In: IJCAI (2011)
13. Monroe, B.L.: Fully proportional representation. *Am. Polit. Sci. Rev.* **89**, 925–940 (1995)
14. Monte, D., Tumennasan, N.: Matching with quorums. *Econ. Lett.* **120**(1), 14–17 (2013)
15. Pisinger, D.: An $O(nr)$ Algorithm for the Subset Sum Problem. DIKU-Rapport, Datalogisk Institut, Københavns Universitet, Datalogisk Institut København (1995)
16. Skowron, P., Faliszewski, P., Slinko, A.M.: Achieving fully proportional representation: Approximability results. *Artif. Intell.* **222**, 67–103 (2015)
17. Spradling, M., Goldsmith, J.: Stability in role based hedonic games. In: Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2015, Hollywood, Florida, 18–20 May 2015, pp. 85–90 (2015)