Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

Stable Invitations

Hooyeon Lee
Computer Science Department
Stanford University
haden.lee@stanford.edu

Yoav Shoahm
Computer Science Department
Stanford University
shoham@cs.stanford.edu

Abstract

We consider the situation in which an organizer is trying to convene an event, and needs to choose whom out of a given set of agents to invite. Agents have preferences over how many attendees should be at the event and possibly also who the attendees should be. This induces a stability requirement: All invited agents should prefer attending to not attending, and all the other agents should not regret being not invited. The organizer's objective is to find an invitation of maximum size, subject to the stability requirement. We investigate the computational complexity of finding such an invitation when agents are truthful, as well as the mechanism design problem when agents act strategically.

1 Introduction

Imagine an event organizer trying to convene an event – for example, a fundraiser. We assume that the time and venue for the event are fixed, and that the only remaining decision for the organizer to make is whom to invite among a set of agents. An *invitation* is simply defined to be a subset of agents. The goal of the organizer is to maximize attendance (for example, in order to maximize donations), but the potential invitees have their own preferences over how many attendees there should be at the event and possibly also who the potential attendees should be. For example, a given donor may not want to attend if too few attendees show up, but she may not want the event to be overly crowded. Another donor may want to attend the event only if her friends attend and her business competitor does not.

We first consider agents with anonymous preferences over invitations – agents only care about how many attendees are at the event (but not the identities of attendees). An invitation is *stable* if all invitees prefer attending to not attending and if no person who is not invited wishes she had been invited. Stability is desirable, but a stable invitation may not exist in general. This naturally raises the question of how hard it is to determine whether it exists for a given setting, and if it does, what the maximum stable invitation is. These questions take an extra meaning in the strategic case, in which agents may misreport their preferences. Can the organizer incentivize the agents to disclose their true preferences? We call this problem the *Anonymous Stable Invitation Problem* (ASIP). If we assume truthful agents, we have an algorithm design problem, and we obtain positive results in this case. If we assume strategic agents, we have a mechanism design problem, and we obtain an impossibility result in general as well as positive results for a special case of the problem.

We then remove the assumption of anonymous preferences and define the *General Stable Invitation Problem* (GSIP) where each agent can specify her acceptance set of agents and rejection set of agents (in addition to her preference over sizes). An agent is willing to attend only if everyone in her acceptance set attends, no one in her rejection set attends, and the number of attendees is acceptable to her. We generalize the definition of stable invitations, accordingly. We then ask the same set of questions as in ASIP. In the non-strategic case we show that computational complexity depends on the size of the largest acceptance and rejection sets. In the strategic case, an impossibility for ASIP directly implies the same impossibility for GSIP. We provide another impossibility result for a different sub-class of GSIP.

In Section 2 we discuss related work in the literature. In Section 3 we define the Anonymous Stable Invitation Problem (ASIP), and investigate both the non-strategic and strategic cases. In Section 4 we define the General Stable Invitation Problem (GSIP), and investigate both the non-strategic and strategic cases. In Section 5 we discuss our contributions and directions for future work.

2 Related Work

Darmann et al. (2012) consider the *Group Activity Selection Problem* (GASP), in which the objective is to assign agents to activities where agents have anonymous preferences over activities as well as the number of participants. The *Anonymous Stable Invitation Problem* (ASIP in Section 3) can be viewed as a sub-class of GASP with a single activity. Yet there are several differences between our work and the work by Darmann et al. First, our main results are anchored in the *General Stable Invitation Problem* (GSIP in Section 4) where agents no longer have anonymous preferences; therefore their hardness results for GASP do not imply similar results for GSIP. Second, we consider strategic agents in both ASIP and GSIP, but Darmann et al. only consider truthful agents. Finally, our impossibility results imply the same negative results for GASP.

The Stable Invitations Problems (both ASIP and GSIP) are closely related to hedonic games. In a hedonic game each agent has preferences over coalitions that she can be a member of (usually given by a weak order over all coalitions). If each agent is indifferent among all coalitions of the same size that includes her, then a given hedonic game is said to be anonymous. Much work has been devoted to analyzing solution concepts in hedonic coalition games such as stability and Pareto-optimality (Bogomolnaia and Jackson. 2002; Dreze and Greenberg. 1980; Aziz and Brandl. 2012). In this work we take a solution concept of (Nash) stability for granted, and focus on analyzing complexity of finding stable solutions. In particular, we are interested in finding Nash-stable invitations and individually rational invitations (we provide formal definitions in Section 3).

Ballester (2004) provides a number of computational complexity results (in fact, hardness results) for finding a core-stable, Nash-stable, or individually rational outcome in hedonic games and anonymous hedonic games. These hardness results do not imply similar hardness results for GSIP for two reasons. First, the hardness results on anonymous hedonic games do not hold for GSIP because GSIP is not equivalent to anonymous hedonic games. Second, one can transform an instance of GSIP into a (non-anonymous) hedonic game in a naïve manner by listing all possible coalitions, but this increases the size of input exponentially. Therefore our results discussed in this work are original and should not be considered as a derivative of the work by Darmann et al. or Ballester.

3 Anonymous Stable Invitation Problem

3.1 Definitions and Notation

Definition 1. An instance of the *Anonymous Stable Invitation Problem* (ASIP) is a pair (N, P) where $N = \{a_1, a_2, \ldots, a_n\}$ is a set of n agents and P is an n-tuple of preferences of agents where $P = (P_1, P_2, \ldots, P_n)$. For each agent a_i, we define P_i to be a total preorder ($\succeq_i$) on the set of outcomes, $X = \{0, 1, 2, \ldots, n\}$. An outcome $x \in (X \setminus \{0\})$ denotes the number of attendees and $x = 0$ is a special outcome of not attending. For any $x_1, x_2 \in (X \setminus \{0\})$, $x_1 \succeq_i x_2$ is interpreted as agent a_i weakly preferring attending the event if x_1 attendees are present (including herself) to attending if x_2 attendees are present (including herself). We use $\succ_i$ and $\sim_i$ to denote the induced strict preferences and indifference relations, respectively. We drop the subscript (i) if it is clear from the context.

We assume that for each a_i and each $x \in (X \setminus \{0\})$, either $x \succ_i 0$ or $0 \succ_i x$. That is, no agent is indifferent between not attending and any other outcome. This assumption is made for convenience and does not change our technical results.

We now formally define invitations and notion of stability.

Definition 2. Given an instance (N, P), an *invitation* S is a subset of N, and is interpreted as the organizer inviting the agents in S. An invitation S satisfies *individual rationality (IR)* if for every agent $a_i \in S$ it holds that $|S| \succ_i 0$. An invitation S exhibits *no exclusion-regret with addition (a-ER)* if for every agent $a_j \notin S$ it holds that $|S \cup \{a_j\}| = (|S|+1) \prec_j 0$. An invitation S exhibits *no exclusion-regret with replacement (r-ER)* if for every agent $a_j \notin S$ it holds that $|S| \prec_j 0$. An invitation is *a-stable* (*r-stable*, respectively) if it satisfies IR and exhibits no a-ER (IR and no r-ER, respectively).

In words, a-ER states that an agent who is not invited should not wish that she had been invited in addition to those who are invited (hence the name 'addition'), while r-ER states that an agent who is not invited should not wish that she had been invited in place of someone else who is invited (hence the name 'replacement'). These two definitions are not equivalent. In most parts of this paper we only use a-ER and a-stable and refer to them simply as ER and stable. However our technical results are unchanged when we instead use r-ER and r-stable (see Section 4.5 for details).

For each agent a_i, we can naturally induce from P_i her preference over the set of all invitations. Given S and S', if both contain a_i, then the preference relation between S and S' is induced from P_i on $|S|$ and $|S'|$ (their cardinality). If neither contains a_i, then both invitations are equivalent to 0 (the outside option). If S contains a_i and S' does not, then the preference relation is induced from P_i on $|S|$ and 0. We overload our notation $(\sim_i, \succeq_i, \succ_i)$ for the induced preferences over invitations. The preferences of agents over invitations are now well-defined.

Let us define two special classes of preferences of agents, increasing (INC) and decreasing (DEC). Informally, agents with INC-preferences (DEC, respectively) prefer the event with more attendees (fewer, respectively). In particular INC-preference implies that n is the most favorable outcome and DEC-preference that 1 is the most favorable outcome.

Definition 3. Agent a_i has an INC-preference if there is some threshold l_i such that $n \succeq_i (n-1) \succeq_i \cdots \succeq_i (l_i) \succ_i 0 \sim k$ for all $1 \leq k < l_i$. Agent a_i has a DEC-preference if there is some threshold h_i such that $1 \succeq_i 2 \succeq_i \cdots \succeq_i (h_i) \succ_i 0 \sim k$ for all $h_i < k \leq n$.

We assume that agents can have arbitrary preferences, but these two special types preferences play an important role in the strategic case in Section 3.4.

3.2 Examples of ASIP

Example 1 (Stable invitations are not unique). Let us consider two agents with identical preferences:

$$P_1 : 1 \succ 0 \succ 2, \quad P_2 : 1 \succ 0 \succ 2$$

Note that both agents have DEC-preferences (with $h_1 = h_2 = 1$). The two stable invitations are $S_1 = \{a_1\}$ and $S_2 = \{a_2\}$. The empty invitation ($\emptyset$) exhibits ER while the full invitation (N) is not IR.

Example 2 (A stable invitation may not exist). Let us consider three agents with the following preferences:

$$P_1 : 3 \succ 2 \succ 0 \succ 1,\ P_2 : 1 \succ 0 \succ 2 \sim 3,\ P_3 : 0 \succ 1 \sim 2 \sim 3$$

Note that a_1 has an INC-preference (with $l_1 = 2$) while a_2 has a DEC-preference (with $h_2 = 1$). a_3 is simply unwilling to attend the event (no outcome is acceptable to her). Due to a_3, any invitation including a_3 is not IR. Between a_1 and a_2 only, the empty invitation exhibits ER due to a_2, $S_1 = \{a_1\}$ is not IR due to a_1, $S_2 = \{a_2\}$ exhibits ER due to a_1, and $\{a_1, a_2\}$ is not IR due to a_2.

3.3 The Non-strategic Case

We first consider the non-strategic case of ASIP in which agents are truthful. As Theorem 1 states, one can find a maximum stable invitation (if it exists) in polynomial time.[1] Due to space we only state our theorem here (formal proof can be found in a long version).

Theorem 1. *There exists a polynomial time algorithm that, given an instance of* ASIP*, determines whether a stable invitations exists, and finds a maximum one if it exists.*

3.4 The Strategic Case

In the strategic case of ASIP, we assume that agents may act strategically in reporting their preferences to the event organizer. Example 3 shows how agents may have an incentive to act strategically.

Example 3 (An agent may act strategically.). Let us revisit Example 1 with two agents with identical preferences:

$$P_1 : 1 \succ 0 \succ 2, \quad P_2 : 1 \succ 0 \succ 2$$

Recall from Example 1 that $S_1 = \{a_1\}$ and $S_2 = \{a_2\}$ are the only two stable invitations provided that both agents are truthful. If S_1 were to be chosen by the organizer, a_2 would have an incentive to act strategically, by reporting $\hat{P}_2(1 \succ 2 \succ 0)$ instead of P_2. Given $(P_1, \hat{P}_2)$, the only stable invitation is S_2 (S_1 now exhibits ER due to a_2). Notice that a_2 strictly prefers S_2 over S_1 (because $1 \succ_2 0$) and has an incentive to misreport in this example. By symmetry, a_1 may act strategically if S_2 were to be chosen.

We first provide a formal definition of a mechanism in the context of ASIP with strategic agents. We then state several impossibility results for general cases of ASIP, and also provide a strategy-proof mechanism for special cases of ASIP. Although we only consider deterministic mechanisms here, we discuss how one can generalize to randomized mechanisms at the end of this section.

Definition 4. Given an instance (N, P) of ASIP, we define V_i (the set of available actions to a_i) to be the set of all preferences over X where $X = \{0, 1, 2, \ldots, n\}$ is the set of outcomes. A (deterministic) *mechanism* is a pair (V, Z) where $V = (V_1 \times \cdots \times V_n)$ is the set of action profiles of all agents (i.e., V_i is a subset of total preorder on X) and $Z : V \mapsto U$ is a mapping from each action profile to an invitation in U where $U = 2^N$. Let $V_{-i} = (V_1 \times \cdots \times V_{i-1} \times V_{i+1} \times \cdots \times V_n)$ be the set of action profiles available to all agents but agent a_i. A mechanism (V, Z) is said to be *strategy-proof* if for all $a_i \in N$ it holds that $Z(P_i, v_{-i}) \succeq_i Z(v_i, v_{-i})$ for all $v_i \in V_i$ and $v_{-i} \in V_{-i}$.

We now formally state our first impossibility result.

Theorem 2. *No strategy-proof mechanism can find a stable invitation, even if it exists, for arbitrary instances of* ASIP.

Proof. Example 3 can serve as a proof; no strategy-proof mechanism can find a stable invitation for this instance. ☐

[1]We acknowledge that Theorem 1 is implied by the work by Darmann et al. (2012)

Intuitively this result is due to the conflicting interests of the organizer and agents – the organizer is trying to maximize attendance while the agents (with DEC-preferences) to minimize. Since no strategy-proof mechanisms can find a stable invitation, one can instead seek to design a strategy-proof mechanism that can find a non-empty IR invitation that may exhibit exclusion-regret. We show that this is also impossible as Theorem 3 states.

Theorem 3. *No strategy-proof mechanism can find a non-empty individually rational (IR) invitation, even if it exists, for arbitrary instances of* ASIP.

Proof. Consider three agents with preferences as follows:

$$P_1 : 3 \succ 0 \succ 2 \sim 1, \ P_2 : 2 \succ 3 \succ 2 \succ 1, \ P_3 : 3 \succ 2 \succ 0 \succ 1$$

There are two non-empty IR invitations: $S_1 = \{a_1, a_2, a_3\}$ and $S_2 = \{a_2, a_3\}$. If a mechanism chooses S_1 given (P_1, P_2, P_3), a_2 can report $(2 \succ 0 \succ 3 \sim 1)$ and make S_2 the only non-empty IR invitation. Similarly, if a mechanism chooses S_2 given (P_1, P_2, P_3), a_3 can report $(3 \succ 0 \succ 2 \sim 1)$ so as to make S_3 the only non-empty IR invitation. The rest of the proof is similar to that of Theorem 2. ☐

Earlier we emphasized that the conflict between agent(s) and the organizer is the main factor that leads to an impossibility result. Indeed, in the example we used in the proof of Theorem 2, both agents have DEC-preferences while the organizer's goal is to maximize attendance.

We now consider a special case of ASIP in which all agents have INC-preferences. We obtain a positive result in this case as Theorem 4 states.

Theorem 4. *There is a strategy-proof mechanism for* INC-*instances of* ASIP*, which can also find a maximum stable invitation in linear time (after sorting).*

Proof sketch. For simplicity let us assume that each agent reports her threshold value (i.e., l_i as defined in Definition 3) as L_i (L_i may differ from l_i). Our mechanism then chooses the largest k such that $L_k \leq k$ holds and chooses the set of k agents with largest threshold values (if no such k exists, mechanism chooses the empty invitation). ☐

Although our mechanism is simple, proof of Theorem 4 is not trivial. First, the full invitation is not necessarily stable (if at least one agent is unwilling to attend at all). Second, an agent may have an incentive to under-report such that the organizer would invite more agents than when the agent is truthful. It is indeed possible to under-report ($L_i < l_i$) and lead to a larger invitation, but we show that this larger invitation would contain less than l_i agents.

3.5 Extensions

Although we have so far only discussed deterministic mechanisms, our impossibility results can be extended to randomized mechanisms. First we define Z to be a mapping from V to $\Pi(U)$ where $\Pi(U)$ denotes the set of all probability distributions over U. The definition of a strategy-proof mechanism must change accordingly – we do this by adopting the axioms in the von Neumann-Morgenstern utility theorem (Von Neumann and Morgenstern 1947). We introduce lotteries over invitations and define preferences of agents over lotteries. Given a probability distribution over invitations, one can compute the expected cardinal utility of lotteries. We then define a strategy-proof mechanism analogously to Definition 4: for each a_i, it must hold that the expected utility of $Z(P_i, v_{-i})$ is no less than the expected utility of $Z(v_i, v_{-i})$ for all $v_i \in V_i$ and for all $v_{-i} \in V_{-i}$. The impossibility result given by Theorem 2 still holds: If (V, Z) is a strategy-proof mechanism, then $Z(P_1, P_2)$ must assign zero probability to both $\{a_1\}$ and $\{a_2\}$, yet these are the only two stable invitations. All other impossibility results we mention in this work can be extended in this manner.

We can extend our results in a different direction by considering multiple time-alternatives for the event. In such settings, agents may have preferences over time-alternatives for the event, in addition to size of invitations. In the non-strategic case, our easiness result is still applicable: One can run the algorithm used in Theorem 1 iteratively for each time-alternative, and choose the maximum stable invitation among all. In the strategic case, our impossibility results immediately imply the same negative results. For INC-instances of ASIP, we obtain a similar impossibility result even if there is only two time-alternatives. The intuition is that over-reporting ($L_i > l_i$) can give the veto power to an agent, which prevents us from designing a strategy-proof mechanism even for INC-instances of ASIP. Note that all of our impossibility results can be naturally extended to the Group Activity Selection Problem by Darmann et al. (2012) since ASIP is a sub-class of GASP with a single activity.

4 General Stable Invitation Problem

We now allow agents to specify which agents they like or do not like, in addition to preferences over sizes of invitations.

4.1 Definitions and Notation

Let us formally define the General Stable Invitation Problem (GSIP) and solution concepts.

Definition 5. An instance of the *General Stable Invitation Problem* (GSIP) is a tuple (N, P, F, R) where N and P are defined the same as before (see Definition 1), $F = (F_1, F_2, \ldots, F_n)$ is a collection of *acceptance sets*, and $R = (R_1, R_2, \ldots, R_n)$ is a collection of *rejection sets* where $F_i \subseteq (N \setminus \{a_i\})$ and $R_i \subseteq (N \setminus \{a_i\})$ for all i.[2] We interpret (P_i, F_i, R_i) such that a_i is willing to attend the event only if all agents in F_i attend, no agent in R_i attends, and the number of attendees is acceptable according to P_i.

Given an instance (N, P, F, R) of GSIP, we say that it is an (α, β)-instance, where $\alpha = \max_{a_i \in N} |F_i|$ and $\beta = \max_{a_i \in N} |R_i|$. It holds by definition that $0 \leq \alpha, \beta \leq n-1$; in particular empty acceptance sets and rejection sets are allowed under our definition. We will later see that our easiness and hardness results rely on (α, β) values. Notice that any ASIP instance is a $(0, 0)$-instance of GSIP, and therefore ASIP is a special case of GSIP. We now define *stability* of invitations analogously to Definition 2.

[2] One may assume $F_i \cap R_i = \emptyset$ for all i, but this is not necessary in this work.

Definition 6. An *invitation* is a subset S of N as before. An invitation S satisfies *individual rationality (IR)* if for every agent $a_i \in S$ it holds that $F_i \subseteq S$, $R_i \cap S = \emptyset$, and $|S| \succ_i 0$. An invitation S exhibits *no exclusion-regret with addition (a-ER)* if for every agent $a_j \notin S$ it holds that $0 \succ_j (S \cup \{a_j\})$. An invitation S exhibits *no exclusion-regret with replacement (r-ER)* if for every agent $a_j \notin S$ there is no $a_i \in S$ such that $(S \setminus \{a_i\} \cup \{a_j\}) \succ_j 0$. An invitation is *a-stable* (*r-stable*, respectively) if it satisfies IR and exhibits no a-ER (IR and no r-ER, respectively).

Note that Definition 6 coincides with Definition 2 when all $\{F_i, R_i\}$'s are empty. From here one we use the definitions of a-ER and a-stable and refer to them simply as ER and stable. However our technical results are unchanged when we instead use r-ER and r-stable (see Section 4.5).

For each agent a_i, we can naturally induce from P_i the preference of a_i over the set of all invitations, 2^N, in the same manner as we did in Section 3.1. Therefore, the preferences of agents over invitations are well-defined. Note that individual rationality and exclusion-regret are properties of an invitation, not preferences of agents, and therefore we define the induced preferences over invitations to be independent of the properties of a solution concept. Next we define a special class of preferences, called *simple* preferences.

Definition 7. A preference P_i is *simple*, if agent a_i strictly prefers any outcome $x \in X$ with $x \neq 0$ to 0. That is, for all $x \in X$ with $x \neq 0$, $x \succ_i 0$.

Note that when an agent has a simple preference, she may still have an arbitrary preference ordering over outcomes, but she strictly prefers attending to not attending.

We emphasize that in GSIP agents still have preferences (P) over sizes of invitations. Most of our hardness and impossibility results for GSIP are stated while assuming that all agents have simple preferences – this assumption strengthens our negative results because they directly imply the same negative results for general cases. On the other hand, we provide our easiness result for GSIP assuming arbitrary preferences, which of course implies the same positive result for GSIP with simple preferences.

4.2 An Example of GSIP

We present two examples of GSIP.

Example 4. Consider four agents with simple preferences and the following acceptance-rejection sets:

$$F_1 = \{a_2\},\ F_2 = \{a_1\},\ F_3 = \{a_4\},\ F_4 = \{a_3\},$$
$$R_1 = \{a_3\},\ R_2 = \{a_4\},\ R_3 = \{a_1\},\ R_4 = \{a_2\}.$$

Agents a_1 and a_2 have each other in their acceptance sets, while they reject agents a_3 and a_4, respectively. Similarly, agents a_3 and a_4 have each other in their acceptance sets, while they reject agents a_1 and a_2, respectively. Among $2^{|N|} = 16$ possible invitations, there are three stable invitations: $\emptyset$, $\{a_1, a_2\}$, and $\{a_3, a_4\}$ (the latter two being maximum). One can easily verify that all other invitations are not stable; for example, $S = \{a_1, a_3, a_4\}$ is not IR due to agents a_1 and a_3. Although we did not describe preferences of agents on sizes, we can still find stable invitations because they all have simple preferences. Note that this example is a $(1,1)$-instance of GSIP with simple preferences.

4.3 The Non-strategic Case

The decision problem of GSIP is whether a stable invitation of size k exists given (N,P,F,R). Computational complexity of this problem depends on (α,β) values – recall that an instance of GSIP is an (α,β)-instance where $\alpha = \max_{a_i \in N} |F_i|$ and $\beta = \max_{a_i \in N} |R_i|$. We state NP-hardness results while assuming agents with simple preferences (this strengthens our negative results) and easiness results while assuming agents with arbitrary preferences (this strengthens our positive results) whenever possible.

Note that NP-hardness for GSIP implies NP-completeness because GSIP is clearly in NP (i.e., one can efficiently check whether a given invitation is stable).

Theorem 5 states that the decision problem of GSIP is NP-hard even if the size of all acceptance sets and rejection sets are at most one. Theorem 6 delivers a similar negative result even if acceptance sets are empty and rejection sets contain at most two agents. Most of our NP-hardness results are obtained by reducing form the MAX-2-SAT problem or the 3-SAT problem. Due to space we omit proofs.

Theorem 5. *It is NP-hard to decide whether a $(1,1)$-instance of* GSIP *admits a stable invitation of size k, even if all agents have simple preferences.*

Theorem 6. *It is NP-hard to decide whether a $(0,2)$-instance of* GSIP *admits a stable invitation of size k, even if all agents have simple preferences.*

We now consider the remaining cases of GSIP whose computational complexity is not implied by Theorems 5 and 6. We know that $(0,0)$-instances of GSIP with arbitrary preferences are solvable in polynomial time (due to Theorem 1) because those are instances of ASIP. In addition if we are given $(1,0)$-instances or $(0,1)$-instances of GSIP with arbitrary preferences, we can find a maximum stable invitation in polynomial time.

Theorem 7. *One can find a maximum stable invitation in polynomial time, given any $(1,0)$-instance of* GSIP *(with arbitrary preferences, P).*

Proof sketch. One con construct a directed graph where each node corresponds to an agent and each edge corresponds to membership of an agent in a rejection set. Since each node contributes at most one edge, we know that each component of G either is a tree or contains a cycle such that each node on the cycle is the root of a tree. By utilizing this structure, we can determine for fixed k whether a stable invitation of size k exists in polynomial time via a dynamic programming algorithm. □

Theorem 8. *One can find a maximum stable invitation in polynomial time, given any $(0,1)$-instance of* GSIP *(with arbitrary preferences, P).*

Proof sketch. We can use a similar technique that was described in the proof sketch of Theorem 7. □

We emphasize that our algorithms for Theorems 7 and 8 rely on the restriction that each agent's acceptance set (rejection set, respectively) is limited to singleton or empty sets.

Finally, we consider $(2,0)$-instances of GSIP. In this subclass of GSIP, the decision problem is in P if we assume simple preferences, whereas it is NP-hard if we assume arbitrary preferences.

Theorem 9. *It is NP-hard to decide whether a $(2,0)$-instance of* GSIP *admits a stable invitation of size k, when agents may have arbitrary preferences.*

Lemma 1. *Given any $(\alpha,0)$-instance of* GSIP *with simple preferences, the full invitation is the unique maximum stable invitation.*

Proof. Consider $S = N$, the full invitation. S exhibits no ER by definition because there is no $a_j \notin S$. For all $a_i \in S$, we know that $F_i \subseteq S = N$ and $R_i = \emptyset$ as $\beta = 0$. Therefore S is individually rational and stable. Since $S = N$, it is indeed the unique maximum stable invitation. □

While Lemma 1 is easy to prove, it is worth noting that the interests of the agents and organizer align in this case, which enables us to efficiently find a maximum stable invitation.

We summarize in Table 1 (on left) our analysis of computational complexity for the non-strategic case of GSIP. P denotes the existence of polynomial time algorithms and NP-C denotes NP-completeness. The entries in boldface remark the results shown in this work, and the other entries are implied by those results. For comparison, we also summarize hardness and easiness results for finding maximum individually rational (IR) invitations on the right. Notice that the only difference between the two problems is the entry for $(\alpha,0)$-instances with $\alpha \geq 2$ with arbitrary preferences. Also note that the full invitation is trivially the maximum if we look for no-ER invitations.

In what follows we state our hardness and easiness results for finding maximum IR invitations in GSIP.

Theorem 10. *It is NP-hard to decide whether a $(1,1)$-instance of* GSIP *admits an individually rational invitation of size k, even if all agents have simple preferences.*

Theorem 11. *It is NP-hard to decide whether a $(0,2)$-instance of* GSIP *admits an individually rational invitation of size k, even if all agents have simple preferences.*

Proof sketch. We reduce from the Maximum-Independent-Set problem in (undirected) cubic graphs, which was shown to be NP-complete (Michael R. Garey and Stockmeyer. 1976). Given a cubic graph with n nodes, we first create n agents each of which corresponds to a node. We then assign either of the two directions to each edge such that each directed edge corresponds to an agent belonging to another agent's rejection set. If we can find an assignment of directions such that each node has at most two outgoing edges, then we can construct a $(0,2)$-instance of GSIP by putting at most two agents in each agent's rejection set. Using the Lovász local lemma we can show that such assignment always exists (we use Theorem 2 of the work by Shearer (1985)). □

	Finding Max. Stable Invitations						Finding Max. IR Invitations					
	Simple Preferences			Arbitrary Preferences			Simple Preferences			Arbitrary Preferences		
	$\beta=0$	$\beta=1$	$\beta\geq 2$	$\beta=0$	$\beta=1$	$\beta\geq 2$	$\beta=0$	$\beta=1$	$\beta\geq 2$	$\beta=0$	$\beta=1$	$\beta\geq 2$
$\alpha=0$	**P**	P	**NP-C**	**P**	**P**	NP-C	P	P	**NP-C**	**P**	**P**	NP-C
$\alpha=1$	**P**	**NP-C**	NP-C	**P**	NP-C	NP-C	P	**NP-C**	NP-C	**P**	NP-C	NP-C
$\alpha\geq 2$	**P**	NP-C	NP-C	**NP-C**	NP-C	NP-C	P	NP-C	NP-C	**P**	NP-C	NP-C

Table 1: Computational complexity of finding maximum stable invitations (on left) and maximum IR invitations (on right).

Theorem 12. *One can find a maximum individually rational invitation in polynomial time, given any $(\alpha, 0)$-instance of* GSIP *(with arbitrary preferences, P).*

Theorem 13. *One can find a maximum individually rational invitation in polynomial time, given any $(0, 1)$-instance of* GSIP *(with arbitrary preferences, P).*

4.4 The Strategic Case

Recall that the impossibility results for the strategic case of ASIP (namely, Theorems 2 and 3) immediately imply the same negative results for GSIP since ASIP is a sub-class of GSIP. These results are mainly due to arbitrary preferences over size of invitations, and we now consider another sub-class of GSIP where all agents have simple preferences. Recall that a_i has a simple preference if $x \succ_i 0$ for all $x \in (X \setminus \{0\})$.

When $\beta = 0$ (i.e., all rejection sets are empty), we immediately obtain an optimal, strategy-proof mechanism due to Lemma 1. We know that the full invitation is the unique maximum stable invitation given $(\alpha, 0)$-instances of GSIP with simple preferences. Therefore a trivial mechanism that invites everyone is strategy-proof and optimal. When $\beta > 0$, we show an impossibility result as stated in Theorem 14. Although we do not formally define a mechanism and strategy-proofness in the context of GSIP, the reader should assume definitions analogous to the one provided in Section 3.4.

Theorem 14. *No strategy-proof mechanism can find a stable invitation, even if it exists, for an arbitrary (α, β) instance of* GSIP *when $\beta > 0$. This remains true even if all agents have simple preferences.*

Proof. Consider a $(0, 1)$-instance of GSIP with two agents $N = \{a_1, a_2\}$. Assume both agents have simple preferences and their rejects sets are given by $R_1 = \{a_2\}$ and $R_2 = \{a_1\}$. Given (R_1, R_2), the only two stable invitations are $S_1 = \{a_1\}$ and $S_2 = \{a_2\}$. Suppose that a mechanism chooses S_2 given (R_1, R_2). Then a_1 has an incentive to misreport as if her rejection set were empty (i.e., $\hat{R}_1 = \emptyset$). Given $(\hat{R}_1, R_2)$, the only stable invitation is now S_1 (and S_2 exhibits ER due to a_1). Since a_1 prefers S_1 to S_2 (because $S_1 \sim 1 \succ 0 \sim S_2$ for a_1), a_1 has an incentive to report $\hat{R}_1$. By symmetry a_2 would have an incentive to misreport if S_1 to be chosen. Thus there is no strategy-proof mechanism that can find a stable invitation for this instance. $\square$

4.5 From a-stability to r-stability

Recall that we defined a-stable and r-stable invitations in Definition 6, but we have so far only discussed the former. However, our technical results (both easiness and hardness results) remain unchanged when we consider r-stable invitations instead. Specifically in the non-strategic case of GSIP, our results in Table 1 do not change (this includes ASIP). In the strategic-case, Theorem 3 implies an impossibility result for both ASIP and GSIP regardless of which definition of exclusion-regret (ER) and stability is chosen.

5 Contributions and Future Work

The main contribution of this work is a thorough analysis of the Stable Invitation Problem from both computational complexity perspective and game-theoretic perspective. For the former, we provide a number of easiness and hardness results on ASIP and GSIP with truthful agents. We show that finding a stable invitation or an individually rational invitation is computationally hard in general. For the latter, we show several impossibility results for mechanism design, and also provide a strategy-proof, optimal mechanism for special cases of ASIP when all agents have INC-preferences. It is worth emphasizing that we obtained positive results when the interests of agents and organizer align (via INC-preferences or empty rejection sets). We also provided a few interesting extensions of ASIP and GSIP.

While we answered many interesting questions, there are many interesting directions for future work. One interesting direction is to study a case where the organizer is interested in hosting a certain number of (identical) events such that each agent is invited to one of those events and that each event admits a stable invitation. This generalization is well-motivated (for instance, consider a series corporate info-sessions for recruiting). Other interesting directions include probabilistic algorithms and mechanisms – for instance, we can relax the individual rationality constraint and allow algorithms or mechanisms to 'fail' with a small probability. This relaxation may let us circumvent the impossibility results stated in Theorems 2, 3, and 14.

Missing Proofs

Due to space we only provided proof sketches and omitted formal proofs. They can be found in the extended version.

Acknowledgements

This work was funded in part by the National Science Foundation (under grants IIS-0963478 and IIS-1347214), the U.S. Army (under grant W911NF1010250), and the Kwanjeong Educational Foundation. The authors also thank the reviewers, Eric Huang, Euiwoong Lee, and Qiang Zhang for providing fruitful suggestions.

References

Aziz, H., and Brandl., F. 2012. Existence of stability in hedonic coalition formation games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, 763–770.

Ballester, C. 2004. Np-completeness in hedonic games. *Games and Economic Behavior* 49(1):1–30.

Bogomolnaia, A., and Jackson., M. O. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2):201–230.

Darmann, A.; Elkind, E.; Kurz, S.; Lang, J.; Schauer, J.; and Woeginger, G. 2012. Group activity selection problem. In *Proceedings of the Internet and Network Economics*, 156–169.

Dreze, J. H., and Greenberg., J. 1980. Hedonic coalitions: Optimality and stability. *Journal of the Econometric Society* 987–1003.

Michael R. Garey, D. S. J., and Stockmeyer., L. 1976. Some simplified np-complete graph problems. *Theoretical Computer Science* 1(3):237–267.

Shearer, J. B. 1985. On a problem of spencer. *Combinatorica* 5(3):241–245.

Von Neumann, J., and Morgenstern, O. 1947. *The Theory of Games and Economic Behavior*. Princeton University Press.