

Computing Efficient and Envy-Free Allocations under Dichotomous Preferences using SAT

Ari Conati, Andreas Niskanen, Ronald de Haan, and Matti Järvisalo

Abstract

We study the problems of computing envy-free Pareto-efficient allocations in the context of fair allocation and hedonic games under dichotomous preferences. We establish Σ_2^P -completeness of deciding the existence of envy-free Pareto-efficient allocations, refining earlier related results. We also develop iterative SAT-based exact algorithms for computing envy-free Pareto-efficient allocations, and extend the approach to computing *minimum-envy* Pareto-efficient allocations under different combinations of aggregation functions. We provide open-source implementations of the algorithms and show empirically that the approach scales to computing envy-free Pareto-efficient allocations up to hundreds of agents.

This article is an abridged version of work published at the 24th International Conference on Autonomous Agents and Multiagent Systems [33].

1 Introduction

Allocation problems concerning multiple agents occur naturally in various real-world settings, from dividing computational resources in clusters [41] through assigning university courses to students [23] to food distribution [1], to name a few examples. Various allocation problems have been studied within computational social choice [26, 57, 3, 59]. We focus on two well-studied problem settings concerning allocations: *fair allocation* of indivisible goods [19, 44, 2, 51], where the task is to divide discrete items between agents based on the preferences of individual agents, and *hedonic games* [4, 39], where the task is to form a partitioning of a set of agents into coalitions based on the individual agents' preferences regarding the coalition they are assigned to. In particular, we study these tasks under dichotomous preferences [14, 15, 38, 7, 21], i.e., in a setting where agents express their preferences via propositional formulas [43, 34, 17, 6, 52].

Envy-freeness is a central and desirable property of fair allocations [40, 58]; an allocation is envy-free if no agent would prefer to be allocated a set of items allocated to another agent over the agent's own allocation. However, envy-freeness in itself admits non-satisfactory allocations, with not allocating any items to any agents as one extreme example. Such non-satisfactory allocations can be ruled out by a choice of a notion of *efficiency*; indeed, in the study of fair allocation it is typical to study combinations of envy-freeness and efficiency notions [17, 35, 18, 5, 9, 25]. Requiring that allocations are complete, i.e., that each item is allocated to some agent, already makes deciding the existence of an envy-free allocation NP-complete [46]. On the other hand, this also means that this decision problem can be polynomially encoded in propositional logic and decided via a single call to a Boolean satisfiability (SAT) solver [10]. Propositional encodings of similar flavor have also been presented [6] for envy-free partitioning into coalitions (or an envy-free allocation for short) in hedonic games. A more refined notion of efficiency in fair allocation is Pareto-efficiency, targeting envy-free allocations in which, intuitively, it is not possible to reallocate items in such a way that some agent would be better off without making some other agent worse off. However, deciding the existence of an envy-free Pareto-efficient allocation is computationally even harder, specifically Σ_2^P -complete [17, 35], which makes the development of efficient algorithms for computing envy-free Pareto-efficient allocations even more challenging. This is the setting we focus on.

In more detail, we focus on the problems of computing envy-free Pareto-efficient allocations in the contexts of fair allocation and hedonic games under dichotomous preferences. Our contributions

are four-fold: (1) we present new Σ_2^P -completeness results for fair allocation and hedonic games; (2) we develop exact SAT-based iterative algorithms for deciding the existence of envy-free Pareto-efficient allocations; (3) we further extend our algorithms to computing minimum-envy Pareto-efficient allocations under different notions of total envy; and (4) we provide open-source implementations of our algorithms and empirically evaluate their scalability.

In terms of complexity results, strengthening earlier results for fair allocation [17], we establish Σ_2^P -completeness for deciding the existence of envy-free Pareto-optimal allocations even when restricting each agent’s preferences to a 3DNF formula with at most four terms and each item to occur in the preferences of at most three agents. This also improves on other earlier results from the literature that only indicated (exponential) bounds on the number of parallel NP oracle queries needed to solve the problem [12]. For hedonic games, building on second-level completeness results in a different setting [53], we establish Σ_2^P -completeness in the specific context of envy-free Pareto-optimal allocations under dichotomous preferences even with similar restrictions as in the case of fair allocation. In terms of algorithms, we build on the earlier-proposed direct SAT encoding of envy-free fair allocation [17] by using the encoding as a base abstraction of our iterative SAT-based counterexample-guided abstraction refinement (CEGAR) [29, 30] approach to the Σ_2^P -complete problem of deciding the existence of a Pareto-efficient envy-free allocation under dichotomous preferences. The approach is motivated by earlier-proposed CEGAR-style approaches to other problem settings in computational social choice [16, 22, 32, 31]. Conceptually, our SAT-based CEGAR iterates between computing envy-free allocations (using a SAT solver) and checking for a counterexample for the claim that the latest found envy-free allocation is Pareto-efficient. Furthermore, due to the fact that our SAT-based CEGAR approach allows for including further constraints on the solutions of interest, our approach also directly captures hedonic games under dichotomous preferences by simply enforcing transitivity over pairwise allocations of agents into the same coalition through a propositional encoding [6]. Going beyond deciding the existence of envy-free Pareto-efficient allocations, it should be noted that a simple “no” answer can be considered insufficient as a “solution” in cases where there are no envy-free Pareto-efficient allocations. To this end, we show how our SAT-based approach can be extended to computing *minimum-envy* [46, 27, 24, 20, 50, 49, 28, 55] Pareto-efficient allocations, considering three combinations of aggregation functions which yield meaningful objective functions for minimizing envy under dichotomous preferences. Again, to the best of our understanding our approach to minimizing envy is the first to enable envy minimization under Pareto-efficiency. In particular, earlier works proposing ways of minimizing envy in fair allocation via, e.g., direct integer programming encodings [55] cannot be directly applied in this setting by complexity-theoretic assumptions. We provide an open-source implementation of our algorithms and show that our approach scales to hundreds of agents.

2 Preliminaries

We begin with an overview on fair allocation and hedonic games under dichotomous preferences [17, 6, 52].

Let $I = \{1, \dots, n\}$ be a set of *agents*, $O = \{o_1, \dots, o_p\}$ a set of *items*, and $R = \{\succeq_1, \dots, \succeq_n\}$ a preference *profile* where for each $i \in I$, $\succeq_i$ is a reflexive, transitive, and complete relation on 2^O . For subsets of items $A, B \subseteq O$, agent i prefers items A to items B if $A \succeq_i B$; if $A \succeq_i B$ and $B \not\succeq_i A$, agent i strictly prefers A to B , denoted by $A \succ_i B$. Preferences are *monotonic* if for any preference relation $\succeq_i \in R$ and $A \subseteq B \subseteq O$ we have $B \succeq_i A$. A preference relation $\succeq_i$ is *dichotomous* if there exists a collection of bundles of items $G_i \subseteq 2^O$ such that for all $A, B \subseteq O$, $A \succeq_i B$ if and only if $A \in G_i$ or $B \notin G_i$. We consider dichotomous preference profiles, i.e., preference profiles consisting only of dichotomous preference relations. An agent i is “happy” with a bundle A if $A \in G_i$, and “unhappy” if $A \notin G_i$.

An *allocation* π is a mapping $\pi: O \rightarrow I$. As standard in literature, we also use π to refer to the inverse

function, i.e., $\pi(i) = \{o \in O \mid \pi(o) = i\}$. An allocation is *envy-free* if, for each agent $i \in I$, $\pi(i) \succeq_i \pi(j)$ for all agents $j \neq i$. An allocation π' *dominates* an allocation π if (i) for all agents $i \in I$, $\pi'(i) \succeq_i \pi(i)$ and (ii) there is an agent $j \in I$ such that $\pi'(j) \succ_j \pi(j)$. An allocation π is *Pareto-efficient* if there is no π' which dominates π .

Hedonic games model scenarios where agents form coalitions amongst themselves, with agents only interested in the members of their own respective coalitions. The preference profile therefore contains preference relations over other agents instead of a set of items. Formally, $\succeq_i$ is a complete and transitive relation over $\{S \subseteq I \mid i \notin S\}$. Analogous to an allocation is a *partition* π of the agents into coalitions. For convenience, we interchangeably use the term “allocation” to refer to such partitions in the context of hedonic games scenarios. Let $\pi(i)$ denote the members of the coalition to which agent i belongs (not including i). A partition is envy-free if $\pi(i) \succeq_i \pi(j)$ for all pairs of agents i, j where i and j are not in the same coalition, i.e., $(\pi(i) \cup \{i\}) \neq (\pi(j) \cup \{j\})$.

Propositional Satisfiability The algorithms we develop are based on iteratively employing Boolean satisfiability (SAT) solvers [10]; we briefly recall necessary background on SAT. For a Boolean variable x there are two literals, x and $\neg x$. A clause C is a disjunction ($\vee$) of literals. A conjunctive normal form (CNF) formula F is a conjunction ($\wedge$) of clauses, while a disjunctive normal form (DNF) formula is a disjunction of conjunctions of literals. For convenience we view clauses as sets of literals and CNF formulas as sets of clauses. We denote by $V(F)$ and $L(F)$ the set of variables and literals of F , respectively. A truth assignment $\tau: V(F) \rightarrow \{0, 1\}$ maps each variable to 0 (false) or 1 (true), and is extended to literals via $\tau(\neg x) = 1 - \tau(x)$, to clauses via $\tau(C) = \max\{\tau(l) \mid l \in C\}$, and to formulas via $\tau(F) = \min\{\tau(C) \mid C \in F\}$. We interchangeably represent truth assignments τ as sets of non-contradictory literals: $\{l \in L(F) \mid \tau(l) = 1\}$. The *Boolean satisfiability* problem (SAT) asks if a given CNF formula F has an assignment τ with $\tau(F) = 1$; if so, F is satisfiable and otherwise unsatisfiable.

Dichotomous Preference Profiles Any dichotomous preference relation $\succeq_i$ can be represented as a propositional formula ϕ_i . Specifically, let $K = O$ if $\succeq_i$ is a preference relation over bundles of items (fair allocation), and $K = I \setminus \{i\}$ if it is a preference relation over sets of agents (hedonic games). Then, for a formula ϕ_i representing $\succeq_i$, we have $V(\phi_i) \supseteq \{p_k \mid k \in K\}$, and $\tau(\phi_i) = 1$ if and only if $\{k \in K \mid \tau(p_k) = 1\} \in G_i$. In other words, the formula is satisfied exactly by truth assignments corresponding to bundles (or coalitions) favored by agent i . In this work, we therefore assume without loss of generality that dichotomous preference profiles are represented as a set of (arbitrary) propositional formulas $R = \{\phi_1, \dots, \phi_n\}$.

Computational Problems We focus on the computational problems of deciding the existence of Pareto-efficient and envy-free allocations under dichotomous preferences in the settings of fair allocation and hedonic games. Formally, the problem input consists of a set I of agents, a set O of items (only in the setting of fair allocation), and a preference profile $R = \{\succeq_1, \dots, \succeq_n\}$ of dichotomous preferences that are specified as formulas $\phi_1, \dots, \phi_n$. The problem is to decide whether there exists an allocation π that is both Pareto-efficient and envy-free (and to output one if it exists).

3 Complexity Results

In this section, we present in brief complexity results establishing that deciding the existence of Pareto-efficient and envy-free allocations (under dichotomous preferences in logical form) is Σ_2^P -complete both for fair allocation and for hedonic games. We show that hardness holds even for the restricted setting where (i) each agent uses a constant-size DNF formula to express their preferences and (ii) each item

(resp. agent) is only mentioned in the formula of a constant number of agents. For fair allocation, this refines an earlier-established Σ_2^P -completeness result [17]. We present the proof of Theorem 1 in part, and we omit the proof of Theorem 2. For full proofs, see the conference version of this work [33].

Theorem 1. *The problem of determining whether there exists a Pareto-efficient and envy-free allocation for a given allocation problem with monotonic, dichotomous preferences under logical form is Σ_2^P -hard, even when: (i) each agent's preferences are expressed by a (positive) 3DNF formula with at most 4 terms; and (ii) each item occurs in at most 3 agents' expressed preferences.*

Proof. To show Σ_2^P -hardness, we give a reduction from $\exists\forall$ -QBF-SAT. Let $\chi = \exists X_1 \forall X_2 \psi$ be a quantified Boolean formula, where $\psi = t_1 \vee \dots \vee t_m$ is a (quantifier-free) formula in 3DNF. Without loss of generality we may assume that each variable in $X_1 \cup X_2$ appears at most twice positively and at most twice negatively in ψ . Let $X_1 = \{x_1, \dots, x_{n_1}\}$ and $X_2 = \{x_{n_1+1}, \dots, x_{n_2}\}$.

We will construct an allocation problem as follows, that admits a Pareto-efficient and envy-free allocation if and only if χ is true. As the set of agents, we take $I = \{(a, i), (b, i) \mid i = 1..n_1\} \cup \{(c, j), (f, j) \mid j = 1..m\} \cup \{(d, i), (e, i) \mid i = 1..n_2\} \cup \{(g, 1), (g, 2)\}$. The set O of items will contain the following items. For each agent of the form $(k, i) \in I$ such that $k \notin \{d, e, g\}$, we introduce an item w_i^k . Moreover, we have an item w^g . The idea behind introducing these items—which will be worked out in more detail below—is that these items are only desirable to the corresponding agent, so any Pareto-efficient allocation will assign these items to them, avoiding envy between these agents.

Next, for each $i = 1..n_1$, we have six items x_i^0, x_i^1, x_i^2 and z_i^0, z_i^1, z_i^2 . The items x_i^1, x_i^2 and z_i^1, z_i^2 correspond to the (at most) two occurrences of x_i and $\neg x_i$, respectively. Additionally, for each $j = 1..m$, we have an item u_j representing the j -th term t_j in ψ . We also have an item y_i for each $i = 1..n_2$, whose use can be intuitively described as forcing a choice between allocating the items corresponding to x_i or the items corresponding to $\neg x_i$ to the agent (d, i) .

The preferences of each agent $(k, i) \in I$ will be described by a formula $\varphi_{(k,i)}$. Before we describe these formulas, we briefly introduce some notation that we will use. For each term t_j in ψ , the formula $\llbracket \neg t_j \rrbracket$ is obtained from t_j as follows. We start with $\neg t_j$, and write it into negation normal form, i.e., as a clause. We then replace each positive literal x_i , corresponding to the ℓ -th occurrence of x_i in ψ , by x_i^ℓ . We replace each negative literal $\neg x_i$, corresponding to the ℓ -th occurrence of x_i in ψ , by z_i^ℓ .

$$\begin{aligned} \varphi_{a,i} &= x_i^0 \wedge w_i^a, & \varphi_{b,i} &= z_i^0 \wedge w_i^b & \text{for } i = 1..n_1 \\ \varphi_{c,1} &= (\llbracket \neg t_j \rrbracket \vee u_j) \wedge w_j^c, & \varphi_{f,j} &= (u_{j-1} \vee u_j) \wedge w_j^f & \text{for } j = 1..m \\ \varphi_{d,i} &= (x_i^0 \wedge x_i^1 \wedge x_i^2) \vee (z_i^0 \wedge z_i^1 \wedge z_i^2) \vee y_i, & \varphi_{e,i} &= y_i & \text{for } i = 1..n_2 \\ \varphi_{g,1} &= \varphi_{g,2} = u_0 \wedge w^g \end{aligned}$$

Next, we will show that this allocation problem admits a Pareto-efficient and envy-free allocation if and only if χ is true, using eight claims (of which we only provide a proof sketch for two). An allocation π is called *regular* if π assigns no agent k an item that is not mentioned in this agent's formula φ_k .

Claim 1. Let π be an allocation, and let π' be the regular allocation that assigns each agent k exactly those items in $\pi(k)$ that are mentioned in φ_k . Then π is Pareto-efficient if and only if π' is. Moreover, if one agent k_1 does not envy another agent k_2 under π then k_1 also does not envy k_2 under π' .

Claim 2. Under any regular allocation π , the following are the only pairs of agents that could possibly envy each other: agents (d, i) and (e, i) , for each $i = 1..n_2$; and agents $(g, 1)$ and $(g, 2)$.

Claim 3. Let π be a regular allocation that is Pareto-efficient and envy-free. Then for each $i = 1..n$, π satisfies both agents (d, i) and (e, i) .

Claim 4. Let π be a regular allocation that is Pareto-efficient and that for each $i = 1..n_2$ satisfies both agents (d, i) and (e, i) . Then for each $i = 1..n_1$, π satisfies at most one of (a, i) and (b, i) .

Claim 5. Let π be a regular allocation that is Pareto-efficient and envy-free. Then for each $i = 1..n_1$, π satisfies at most one of (a, i) and (b, i).

Claim 6. Let π be an allocation that is envy-free. Then neither of the agents (g, 1) and (g, 2) is satisfied.

Claim 7. If there exists an allocation π that is Pareto-efficient and envy-free. Then the QBF χ is true.

PROOF (SKETCH) OF CLAIM 7. Take an allocation π that is Pareto-efficient and envy-free. By Claim 1, we may assume without loss of generality that π is regular. We construct a (partial) truth assignment α to the variables in X_1 as follows. For each $i = 1..n_1$, we let $\alpha(x_i) = 1$ if π satisfies agent (b, i), we let $\alpha(x_i) = 0$ if π satisfies agent (a, i), and we let $\alpha(x_i)$ be undefined otherwise. By Claim 5 (which follows from Claims 3 and 4), we know that π satisfies at most one of (a, i) and (b, i), which means that α is well defined. One can show—using Claim 6—that $\psi[\alpha]$ is a valid 3DNF formula, i.e., that all truth assignments extending α make ψ true. This witnesses that the QBF χ is true. $\dashv$

Claim 8. If the QBF χ is true. Then there exists an allocation π that is Pareto-efficient and envy-free.

PROOF (SKETCH) OF CLAIM 8. Suppose that χ is true. This means that there is a truth assignment α to the variables in X_1 such that $\psi[\alpha]$ is a valid 3DNF formula. We will construct an allocation π that is Pareto-efficient and envy-free, as follows. To the agents (a, i) and (b, i), it assigns items as follows. Take some $i \in \{1, \dots, n_1\}$. If $\alpha(x_i) = 0$, it assigns to agent (a, i) item w_i^a only and it assigns to agent (b, i) the items w_i^b and z_i^1 . Conversely, if $\alpha(x_i) = 1$, it assigns to agent (b, i) item w_i^b only and it assigns to agent (a, i) the items w_i^a and x_i^1 . To each agent (d, i), it assigns the items z_i^0, z_i^1, z_i^2 if $\beta(x_i) = 1$, and it assigns the items x_i^0, x_i^1, x_i^2 if $\beta(x_i) = 0$. To each agent (e, i), it assigns the item y_i . To each agent (f, i) with $i \leq j_0$, it assigns the items u_{i-1}, w_i^f . To each agent (f, i) with $i > j_0$, it assigns the items u_i, w_i^f . To agent (g, 1), it assigns the item w^g , and to agent (g, 2), it assigns no items. Finally, we consider the assignment to the agents (c, j). Take an arbitrary $j \in \{1, \dots, m\}$. The allocation π' assigns to (c, j) all items mentioned in $\varphi_{c,j}$ that have not yet been assigned to other agents. In addition, if $j = j_0$, then it assigns to agent (c, j) in addition the item u_{j_0} , where j_0 is an arbitrary fixed index such that β makes the term t_{j_0} true. One can show that π is envy-free and Pareto-efficient using Claims 1, 2 and 4. $\square$

Theorem 2. *The problem of determining whether there exists a Pareto-efficient and envy-free partition for a given Boolean hedonic game is Σ_2^P -hard, even when: (i) each agent's preferences are expressed by a 3DNF formula with at most 3 terms; and (ii) each agent occurs in at most 3 other agents' expressed preferences.*

4 Encoding Allocations as SAT

It is well-known that under dichotomous preferences, both in the context of fair allocation [17] and hedonic games [6], envy-free allocations can be captured as satisfying truth assignments of a specific propositional formula, i.e., a SAT encoding. This motivates the use of SAT as a declarative paradigm for computing envy-free allocations in practice. We assume as input a set of agents $I = \{1, \dots, n\}$, a set of items O (for fair allocation), and a dichotomous preference profile $R = \{\phi_1, \dots, \phi_n\}$.

We begin by recalling the SAT encoding for envy-freeness in fair allocation [17]. We introduce variables $p_{i,o}$ for all $i \in I$ and $o \in O$ with the interpretation that for a truth assignment τ , it holds that $\tau(p_{i,o}) = 1$ if and only if item o is assigned to the bundle of agent i . To encode a *complete* allocation, we ensure that every item is allocated to *exactly one* agent via

$$\text{COMPLETENESS}(I, O) = \bigwedge_{o \in O} \left(\sum_{i \in I} p_{i,o} = 1 \right),$$

where $\sum_{i \in I} p_{i,o} = 1$ is an exactly-one constraint. Such cardinality constraints are converted to clauses by making use of readily-available CNF encodings [54].

Envy-freeness is then captured as follows [17]. We define for each $i \in I$ the formula $\phi_i^* = \phi_i[p_o \mapsto p_{i,o} \mid o \in O]$ which evaluates to true if agent i is happy with the bundle of items allocated to them. For

each pair $i, j \in I, i \neq j$, we then let $\phi_i^*[j] = \phi_i[p_o \mapsto p_{j,o} \mid o \in O]$, i.e., $\phi_i^*[j]$ evaluates to true iff agent i would be happy with the bundle assigned to agent j . Then envy-freeness is encoded via

$$\text{EF}(R) = \bigwedge_{\substack{i,j \in I \\ i \neq j}} \neg(\neg\phi_i^* \wedge \phi_i^*[j]) = \bigwedge_{\substack{i,j \in I \\ i \neq j}} (\phi_i^*[j] \rightarrow \phi_i^*).$$

In words, $\text{EF}(R)$ declares for every pair of agents $i, j \in I$ with $i \neq j$ that it is not the case that i envies j , i.e., if agent i would be happy with the bundle received by agent j (encoded by $\phi_i^*[j]$), they are happy with their respective bundle (encoded by ϕ_i^*).

In summary, for a fair allocation instance (I, O, R) , any truth assignment τ satisfying $\text{COMPLETENESS}(I, O) \wedge \text{EF}(R)$ corresponds to an envy-free allocation via $\pi(i) = \{o \in O \mid \tau(p_{i,o}) = 1\}$.

Adjustment to Hedonic Games. The above encoding is adjusted to hedonic games [6] as follows. We instead introduce variables $p_{i,j}$ for all $i, j \in I, i \neq j$, with the interpretation that $\tau(p_{i,j}) = 1$ if and only if agent j is a member of the coalition of agent i . Symmetry can be encoded by treating $p_{i,j}$ and $p_{j,i}$ as the same variable (following [6]). Transitivity is encoded using the constraint

$$\text{TRANSITIVITY}(I) = \bigwedge_{\substack{i,j,k \in I \\ i \neq j \neq k}} ((p_{i,j} \wedge p_{j,k}) \rightarrow p_{i,k}),$$

ensuring that satisfying assignments correspond to partitions of the set of agents. For encoding envy-freeness, we define for each $i \in I$ the formula $\phi_i^* = \phi_i[p_k \mapsto p_{i,k} \mid k \in I \setminus \{i\}]$, and for each pair $i, j \in I, i \neq j$,

$$\phi_i^*[j] = \phi_i[p_k \mapsto p_{j,k} \mid k \in I \setminus \{i, j\}][p_j \mapsto p_{i,j}]$$

which evaluates to true iff agent i would be happy by swapping partitions with agent j . Now for an instance of hedonic games (I, R) , truth assignments satisfying $\text{TRANSITIVITY}(I) \wedge \text{EF}(R)$ correspond to envy-free partitions via $\pi(i) = \{j \in I \setminus \{i\} \mid \tau(p_{i,j}) = 1\}$.

5 Iterative SAT for Efficient Allocations

We develop iterative procedures for identifying Pareto-efficient (non-dominated) allocations, including a SAT-based CEGAR algorithm for the Σ_2^P -complete task of finding allocations which are both Pareto-efficient and envy-free. Each of the algorithms takes as input either a fair allocation instance with agents I , items O , and preference profile R (consisting of agents' preferences over the items), or a hedonic game instance defined by agents I and preference profile R (consisting of agents' preferences over the remaining agents).

5.1 Maximizing Efficiency

Pareto-efficient allocations correspond to so-called maximally satisfiable subsets (MSSes) [45, 47] of the formulas $\{\phi_i^*, \dots, \phi_n^*\}$ under $\text{COMPLETENESS}(I, O)$ for fair allocation [17] or $\text{TRANSITIVITY}(I)$ for hedonic games [6]. In practice, computing such an MSS can be done with a series of calls to a SAT solver, making use of SAT calls under assumptions (i.e., partial assignments) in order to find an MSS incrementally, i.e., without starting the SAT solver from scratch after computing a satisfiable subset.

First, we assign a fresh variable name to each of the input formulas via $q_i \leftrightarrow \phi_i^*$ for each $i \in I$, and, for fair allocation, combine this with COMPLETENESS to form the SAT instance F defined as

$$\text{ABSTRACTION}_A(I, O, R) = \text{COMPLETENESS}(I, O) \wedge \bigwedge_{i \in I} (q_i \leftrightarrow \phi_i^*).$$

Algorithm 1 SAT-based CEGAR for finding EEF allocations. **Input:** Problem variant $S \in \{A, H\}$, agents I , items O (if $S = A$), and preference profile R .

```

1: if  $S = A$  then  $F \leftarrow \text{ABSTRACTION}_A(I, O, R)$ 
2: else if  $S = H$  then  $F \leftarrow \text{ABSTRACTION}_H(I, R)$ 
3: while true do
4:    $(result, \tau_{\text{abs}}) \leftarrow \text{SAT}(F \wedge \text{EF}(R))$ 
5:   if  $result = \text{unsat}$  then return false
6:    $\pi_{\text{abs}} \leftarrow \text{ALLOC}(\tau_{\text{abs}})$ 
7:   while  $result = \text{sat}$  do
8:      $F_{\text{dom}} \leftarrow F \wedge \bigwedge_{i \in \text{HAPPY}(\tau_{\text{abs}})} q_i \wedge \bigvee_{i \in \text{UNHAPPY}(\tau_{\text{abs}})} q_i$ 
9:      $(result, \tau) \leftarrow \text{SAT}(F_{\text{dom}} \wedge \text{EF}(R))$ 
10:    if  $result = \text{sat}$  then  $\tau_{\text{abs}} \leftarrow \tau, \pi_{\text{abs}} \leftarrow \text{ALLOC}(\tau)$ 
11:     $F_{\text{dom}} \leftarrow F \wedge \bigwedge_{i \in \text{HAPPY}(\tau_{\text{abs}})} q_i \wedge \bigvee_{i \in \text{UNHAPPY}(\tau_{\text{abs}})} q_i$ 
12:     $(result, \tau_{\text{cex}}) \leftarrow \text{SAT}(F_{\text{dom}})$ 
13:    if  $result = \text{unsat}$  then return  $\pi_{\text{abs}}$ 
14:     $F \leftarrow F \wedge \bigvee_{i \in \text{UNHAPPY}(\tau_{\text{cex}})} q_i$ 

```

Analogously, for hedonic games we employ TRANSITIVITY (instead of completeness) to form the SAT instance F defined as

$$\text{ABSTRACTION}_H(I, R) = \text{TRANSITIVITY}(I) \wedge \bigwedge_{i \in I} (q_i \leftrightarrow \phi_i^*).$$

In both cases, for a satisfying assignment τ to F and for all $i \in I$, it holds that $\tau(q_i) = 1$ iff $\tau(\phi_i^*) = 1$. We distinguish the happy and unhappy agents, respectively, via $\text{HAPPY}(\tau) = \{i \in I \mid \tau(q_i) = 1\}$ and $\text{UNHAPPY}(\tau) = \{i \in I \mid \tau(q_i) = 0\}$. To find a Pareto-efficient allocation (i.e., an MSS of $\{\phi_i^*, \dots, \phi_n^*\}$) we iteratively search for a satisfying assignment τ' to $F \wedge \bigwedge_{i \in \text{HAPPY}(\tau)} q_i \wedge \bigvee_{i \in \text{UNHAPPY}(\tau)} q_i$, set $\tau = \tau'$, and continue until this formula is unsatisfiable. In this case τ corresponds to a Pareto-efficient allocation.

5.2 Combining Envy-freeness and Efficiency

For determining existence of allocations or partitions which are both envy-free and efficient, we propose a SAT-based CEGAR algorithm presented as Algorithm 1. It returns a Pareto-efficient envy-free allocation if one exists, and *false* otherwise.

Firstly, we initialize an *abstraction* F , that is, a SAT instance whose solutions bijectively correspond to complete allocations of the items amongst the agents in the case of fair allocation (line 1), or partitions of the agents into disjoint coalitions in the case of hedonic games (line 2). We then iteratively solve this instance, further enforcing the constraint that the solution corresponds to an allocation which is envy-free with respect to the preferences of the agents (line 4). If there is no solution, we return *false*, as there is therefore no candidate allocation which is envy-free (line 5). Otherwise, we extract a candidate allocation π_{abs} from the obtained solution τ_{abs} (line 6). It remains to be checked, however, whether π_{abs} is efficient, that is, whether π_{abs} is dominated by any other allocation(s).

We first conduct a series of SAT solver calls to obtain a candidate allocation which is not dominated by any other *envy-free* allocation. This is accomplished by a similar procedure as outlined in Section 5.1 for computing a Pareto-efficient allocation. We construct a formula F_{dom} the solutions of which correspond to allocations which dominate π_{abs} (line 8), enforcing that all agents who are happy with the bundle allocated by τ_{abs} remain happy, while at least one of the unhappy agents becomes happy. Then, we query a SAT solver for a solution to F_{dom} (line 9). If a solution exists, we replace π_{abs} by the obtained candidate, which dominates π_{abs} (line 10). We repeat this process until F_{dom} is unsatisfiable, i.e., until π_{abs} is not dominated by any envy-free allocation.

Algorithm 2 CEGAR for finding a minimum-envy Pareto-efficient allocation. **Input:** Problem variant $S \in \{A, H\}$, agents I , items O (if $S = A$), preference profile R , global envy aggregator $\star$, and local envy aggregator $\square$.

```

1:  $\pi^* \leftarrow \emptyset, e^* \leftarrow \infty$ 
2: if  $S = A$  then  $F \leftarrow \text{ABSTRACTION}_A(I, O, R) \wedge \mathcal{E}(I)$ 
3: else if  $S = H$  then  $F \leftarrow \text{ABSTRACTION}_H(I, R) \wedge \mathcal{E}(I)$ 
4: while true do
5:    $(result, \tau_{\text{abs}}) \leftarrow \text{SAT}(F \wedge \text{BOUND}(\star, \square, e^*))$ 
6:   if  $result = \text{unsat}$  then return  $\pi^*$ 
7:    $\pi_{\text{abs}} \leftarrow \text{ALLOC}(\tau_{\text{abs}})$ 
8:   while  $result = \text{sat}$  do
9:      $F_{\text{dom}} \leftarrow F \wedge \bigwedge_{i \in \text{HAPPY}(\tau_{\text{abs}})} q_i \wedge \bigvee_{i \in \text{UNHAPPY}(\tau_{\text{abs}})} q_i$ 
10:     $(result, \tau) \leftarrow \text{SAT}(F_{\text{dom}} \wedge \text{BOUND}(\star, \square, e^*))$ 
11:    if  $result = \text{sat}$  then  $\tau_{\text{abs}} \leftarrow \tau, \pi_{\text{abs}} \leftarrow \text{ALLOC}(\tau)$ 
12:     $F_{\text{dom}} \leftarrow F \wedge \bigwedge_{i \in \text{HAPPY}(\tau_{\text{abs}})} q_i \wedge \bigvee_{i \in \text{UNHAPPY}(\tau_{\text{abs}})} q_i$ 
13:     $(result, \tau_{\text{ceex}}) \leftarrow \text{SAT}(F_{\text{dom}})$ 
14:    if  $result = \text{unsat}$  then  $\pi^* \leftarrow \pi_{\text{abs}}, e^* \leftarrow \text{ENVY}(\pi_{\text{abs}})$ 
15:    else  $F \leftarrow F \wedge \bigvee_{i \in \text{UNHAPPY}(\tau_{\text{ceex}})} q_i$ 

```

Finally, we check for a *counterexample* to the candidate allocation π_{abs} . This is achieved by dropping the envy-freeness constraint and querying the SAT solver for an allocation that dominates π_{abs} (lines 11–12). If no counterexample exists, π_{abs} is non-dominated, so we return π_{abs} as a Pareto-efficient envy-free allocation (line 13). Otherwise, the obtained solution τ_{ceex} corresponds to an allocation π_{ceex} which dominates π_{abs} (but is not envy-free). In this case we continue by *refining* the abstraction by conjoining to F a clause enforcing that at least one agent who is unhappy with the counterexample allocation π_{ceex} must be happy with allocations extracted from solutions to subsequent calls (line 14), excluding τ_{ceex} as a solution.

5.3 Minimizing Envy

Since an efficient envy-free allocation is not guaranteed to exist, we additionally consider the task of minimizing envy [50, 28], outlining an algorithm which finds a Pareto-efficient allocation with minimum envy, that is, envy at least as low as any other efficient allocation. The amount of envy in a given allocation is quantified via a pair of aggregation functions [28]. A *local* aggregation function defines the local degree of envy for a *single* agent i , i.e., $\text{ENVY}(i) = \square_{j \neq i} e_{i,j}$, where $\square$ is an aggregation function, and $e_{i,j}$ is a Boolean variable assigned to 1 iff agent i envies agent j . Then, a *global* aggregation function $\star$ combines the envy of all of the agents together, that is, $\text{ENVY}(\pi) = \star_{i \in I} \text{ENVY}(i)$. We consider three combinations of aggregators which correspond to reasonable optimization objectives in the context of dichotomous preferences:

- local aggregator $\square = \bigvee$ and global aggregator $\star = \sum$ correspond to $\text{ENVY}(\pi) = \sum_{i \in I} \left(\bigvee_{j \neq i} e_{i,j} \right)$, that is, the number of envious agents;
- using $\square = \star = \sum$ as both the local and global aggregator defines $\text{ENVY}(\pi) = \sum_{i \in I} \left(\sum_{j \neq i} e_{i,j} \right)$, i.e., the total number of (i, j) pairs where agent i envies agent j (absolute envy);
- local aggregator $\square = \sum$ and global aggregator $\star = \max$ correspond to $\text{ENVY}(\pi) = \max_{i \in I} \left(\sum_{j \neq i} e_{i,j} \right)$, that is, the maximum number of agents envied by any single agent (maximum envy).

Towards a CEGAR algorithm, we define constraints

$$\mathcal{E}(I) = \bigwedge_{\substack{i,j \in I \\ i \neq j}} (e_{i,j} \leftrightarrow (\neg \phi_i^* \wedge \phi_i^*[j])).$$

In words, variable $e_{i,j}$ is true if and only if agent i envies agent j . Further, we denote by $\text{BOUND}(\star, \square, k)$ a cardinality constraint [54] which bounds the envy of a candidate allocation to less than k . For the three combinations we consider,

$$\text{BOUND}(\sum, \vee, k) = \sum_{i \in I} \left(\bigvee_{j \neq i} e_{i,j} \right) \leq k - 1$$

enforces such a bound on the number of envious agents,

$$\text{BOUND}(\sum, \sum, k) = \sum_{i \in I} \sum_{j \neq i} e_{i,j} \leq k - 1$$

similarly constrains absolute envy, and finally

$$\text{BOUND}(\max, \sum, k) = \bigwedge_{i \in I} \left(\sum_{j \neq i} e_{i,j} \leq k - 1 \right)$$

enforces a bound on maximum envy.

Our CEGAR algorithm for finding a minimum-envy Pareto-efficient allocation is detailed as Algorithm 2, as an extension of Algorithm 1. We start by initializing the best known solution so far to $\emptyset$ and its total envy to ∞ (line 1). We initialize an abstraction which encodes complete or transitive allocations, and adds constraints $\mathcal{E}(I)$ for pairs of envious agents (lines 2–3). We iteratively solve this instance under the constraint that the total envy of the allocation is strictly less than the total envy of the best known solution (line 5). If no solution is found, we return the current best allocation, as there is no candidate allocation with less total envy (line 6). Otherwise the obtained solution corresponds to a candidate allocation π_{abs} (line 7) which is not known to be efficient. We obtain through a sequence of SAT solver calls a candidate allocation which is not dominated by any other allocation *with less envy than the best known solution* (lines 8–11). A counterexample to the candidate allocation π_{abs} is another allocation which dominates it. We drop the constraint which bounds the total envy and search for such an allocation (lines 12–13). If there is no counterexample, π_{abs} is efficient, so we set it as the best known solution and update the total envy (line 14). Otherwise, the obtained solution corresponds to an allocation dominating π_{abs} , but which has more total envy. As in Algorithm 1, we refine the abstraction via a clause which states that at least one agent who is unhappy with the counterexample allocation must be happy in subsequent iterations (line 14).

6 Empirical Evaluation

We implemented the CEGAR approach and its extension to minimizing envy on top of PySAT [42], using CaDiCaL [11] (version 1.9.5) incrementally as the underlying SAT solver. We use the sequential counter encoding [56] for exactly-1 constraints and incremental totalizers [48] for at-most-k constraints, offered by PySAT. The implementation, benchmark generators and experiment data are openly available at <https://bitbucket.org/coreo-group/satfair>. The experiments were run on 2.40-GHz Intel Xeon Gold 6148 CPUs and 381-GB memory using a per-instance 30-minute time and 16-GB memory limit (the memory limit was exceeded only when minimizing absolute envy for 300 agents).

We generated benchmarks as follows. Agents express their instances in negation-free DNF. Each conjunction in a DNF represents a bundle of items preferred by the agent, and the bundles may overlap. We assigned each agent 5–10 preferred bundles with 3–4 items per bundle, selected for each agent uniformly at random. These parameter values reflect the complexity results in Section 3. We generated

Table 1: Envy-free Pareto-efficient fair allocation.

#a	#i	#solved	avg. time (s)	#yes	#no
300	5	100	38.22	0	100
300	10	100	90.61	0	100
300	15	100	60.55	100	0
300	20	100	36.15	100	0
300	25	100	39.72	100	0
300	30	100	41.76	100	0
400	5	100	84.40	0	100
400	10	100	181.82	0	100
400	15	100	167.22	100	0
400	20	100	64.72	100	0
400	25	100	72.98	100	0
400	30	100	76.80	100	0
500	5	100	159.73	0	100
500	10	100	306.96	0	100
500	15	100	547.38	74	26
500	20	100	106.98	100	0
500	25	100	117.38	100	0
500	30	100	123.48	100	0
600	5	100	225.05	0	100
600	10	100	520.79	0	100
600	15	100	1047.41	10	90
600	20	100	168.11	100	0
600	25	100	188.34	100	0
600	30	100	203.11	100	0
700	5	100	313.09	0	100
700	10	100	860.89	0	100
700	15	99	1389.19	0	99
700	20	100	240.12	100	0
700	25	100	260.57	100	0
700	30	100	303.29	100	0

instances for $n = 300, 400, \dots, 700$ agents, and report for each n on a range of values for the total number of items to ensure that we obtained challenging-enough, both “yes” instances (where an EEF allocation exists) and “no” instances (where an EEF allocation does not exist). Intuitively, and as observed in our experiments, if items are sufficiently abundant, one can expect there to be various EEF allocations. In contrast, by considering relatively low numbers of items wrt agents, there can be expected to be few EEF allocations or none at all.

An overview of the results for deciding EEF fair allocation are shown in Table 1 and Figure 1. Overall, our CEGAR approach scales up to 700 agents (#a) for each number i of items (#i) and n of agents considered. By varying i and n , we observe that there exists a sharp transition from all instances being “no” to all instances being “yes” at specific thresholds i/n . The hardest-to-solve instances for our CEGAR algorithm appear to be those “critical” instances which were generated near the threshold for each n . In addition to average runtimes (y axis) individually for $n = 500, 600, 700$ agents for different numbers of items (x axis), Figure 1 also shows curves fit to the data shown in Table 1 through polynomial interpolation for the potential no-to-yes transition points for each n . Interestingly, a phase transition phenomenon has been previously reported on in the context of envy-free fair allocation under additive preferences [36]. The problem we consider here is naturally related but different (and also harder in




Figure 1: Fair allocation: average runtimes over solved instances, % of ‘yes’ instances.

terms of computational complexity). Nevertheless, the phase-transition-like behaviour we observe here could warrant further investigation of independent interest.

To obtain benchmarks for minimum-envy Pareto-efficient fair allocation, we generated smaller instances (reflecting the increased difficulty of the problem), using $n = 100, 200, 300$ and keeping other parameters the same, and solved the instances with the CEGAR algorithm to the decision problem to find 100 “no” instances for each n . An overview of the runtime performance of our CEGAR approach extended to minimizing envy on the resulting benchmark instances is shown in Table 2. Here mea, mae, and mme refer to the three objectives of minimizing the number of agents envious of some other agent, minimizing absolute envy in terms of the number of times the agents are envious of other agents, and minimizing the maximum number of envied agents over the individual agents, respectively. Overall, as expected we observe that minimizing envy is indeed empirically harder than deciding the existence of an envy-free Pareto-efficient allocation. Regardless we can solve instances at least up to 200 agents. Minimizing maximum envy appears to be the easiest for our approach based on these benchmarks, while minimizing absolute envy appears the hardest.

Finally, we consider hedonic games. We generated benchmarks for deciding the existence of EEF coalitions as follows. As in the fair allocation instances, each agent’s preferences are expressed as a DNF formula with 5–10 DNF terms and 3–4 agents per term, selected uniformly at random. As the grand coalition is trivially a solution if the DNF contains only positive variables, we flip each literal in the DNFs with probabilities 0.25 and 0.5, expressing that an agent prefers a coalition which another specific agent is not in. We generated instances for $n = 20, 30, \dots, 100$ agents. An overview of the results is shown in Table 3. Interestingly, the probability used for negating literals has a significant impact on the runtime of our approach. With the lower probability 0.25, our approach performs somewhat modestly, with an increasing number of timeouts beyond 60 agents. Using the higher probability 0.5 results in noticeably

Table 2: Minimum-envy Pareto-efficient fair allocation.

		#solved (avg. time (s))					
#a	#i	mea		mae		mme	
100	5	100	(52.31)	100	(90.29)	100	(16.54)
100	10	100	(78.34)	100	(110.04)	100	(27.59)
200	5	100	(600.10)	100	(1158.43)	100	(163.98)
200	10	25	(1474.77)	4	(1515.76)	88	(411.60)
300	5	13	(1515.53)	0	(−)	100	(753.81)
300	10	0	(−)	0	(−)	16	(977.27)

Table 3: Results on hedonic games.

#a	neg. = 0.5		neg. = 0.25	
	#solved	avg. time (s)	#solved	avg. time (s)
20	100	0.26	100	0.34
30	100	0.47	100	1.46
40	100	0.99	100	6.42
50	100	1.65	100	28.02
60	100	3.02	98	238.61
70	100	4.76	53	632.73
80	100	7.52	23	212.61
90	100	11.16	15	291.49
100	100	16.97	7	294.78

better scalability, with all instances solved up to 100 agents (and likely beyond). It should be noted that while we were able to observe a yes-no transition for the fair allocation benchmark generation model, for hedonic games the benchmark generation parameters we used here appear to yield mainly “yes” instances, with very few “no” instances. These observations suggest a more involved study into more fine-grained benchmark generation models for hedonic games as well as further investigating how choices within the large parameter space for generated benchmark instances affect algorithmic performance.

7 Conclusions

We presented new complexity results and SAT-based algorithms for the Σ_2^P -complete problems of computing envy-free Pareto-efficient allocations in the context of fair allocation and hedonic games under dichotomous preferences. Refining earlier related results, we establish Σ_2^P -completeness of deciding the existence of envy-free Pareto-efficient allocations even when limiting the number of preferences and preferences to restricted-size DNFs. The SAT-based CEGAR algorithms we developed for deciding envy-free Pareto-efficient allocations constitute to the best of our knowledge the first practical approaches proposed for these problems. Further, we extended the algorithmic approach to computing *minimum-envy* Pareto-efficient allocations under reasonable measures of global envy, and, providing an open-source implementation of the algorithms showed empirically that our approach scales reasonably. An interesting direction for further work would be to extend the algorithmic approach to handling additive preferences, and, in the case of hedonic games, to notions of stability such as core and strict-core stable coalitions [8, 13, 37]. Furthermore, investigating the potentially underlying phase-transition-like phenomenon in this problem setting further would be interesting; and additionally, evaluating the algorithmic approach on real-world allocation instances would also be of significant interest, especially as randomly generated benchmarks tend to typically be significantly harder than more structured instances for modern SAT solvers.

Acknowledgments

Work financially supported by Research Council of Finland (347588, 356046), and the Finnish Doctoral Program Network in Artificial Intelligence AI-DOC (VN/3137/2024-OKM-6). The authors thank the Finnish Computing Competence Infrastructure (FCCI) for computational and data storage resources.

References

- [1] Martin Aleksandrov, Haris Aziz, Serge Gaspers, and Toby Walsh. Online fair division: Analysing a food bank problem. In *IJCAI 2015*, pages 2540–2546. AAAI Press, 2015.
- [2] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322:103965, 2023.
- [3] Haris Aziz. Developments in multi-agent fair allocation. In *AAAI 2020*, pages 13563–13568. AAAI Press, 2020.
- [4] Haris Aziz and Rahul Savani. Hedonic games. In *Handbook of Computational Social Choice*, pages 356–376. Cambridge University Press, 2016.
- [5] Haris Aziz, Serge Gaspers, Simon Mackenzie, and Toby Walsh. Fair assignment of indivisible objects under ordinal preferences. *Artif. Intell.*, 227:71–92, 2015.
- [6] Haris Aziz, Paul Harrenstein, Jérôme Lang, and Michael J. Wooldridge. Boolean hedonic games. In *KR 2016*, pages 166–175. AAAI Press, 2016.
- [7] Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. Fair mixing: The case of dichotomous preferences. *ACM Trans. Economics and Comput.*, 8(4):18:1–18:27, 2020.
- [8] Suryapratim Banerjee, Hideo Konishi, and Tayfun Sönmez. Core in a simple coalition formation game. *Soc. Choice Welf.*, 18(1):135–153, 2001.
- [9] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish. Finding fair and efficient allocations. In *EC 2018*, pages 557–574. ACM, 2018.
- [10] Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh, editors. *Handbook of Satisfiability - Second Edition*, volume 336 of *FAIA*. IOS Press, 2021. ISBN 978-1-64368-160-3.
- [11] Armin Biere, Tobias Faller, Katalin Fazekas, Mathias Fleury, Nils Froleyks, and Florian Pollitt. Cadical 2.0. In *CAV 2024*, volume 14681 of *LNCS*, pages 133–152. Springer, 2024.
- [12] Bernhard Bliem, Robert Brederbeck, and Rolf Niedermeier. Complexity of efficient and envy-free resource allocation: Few agents, resources, or utility levels. In *IJCAI 2016*, pages 102–108. IJCAI/AAAI Press, 2016.
- [13] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games Econ. Behav.*, 38(2):201–230, 2002.
- [14] Anna Bogomolnaia and Herve Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004. ISSN 00129682, 14680262.
- [15] Anna Bogomolnaia, Hervé Moulin, and Richard Stong. Collective choice under dichotomous preferences. *J. Econ. Theory*, 122(2):165–184, 2005.
- [16] Arthur Boixel and Ulle Endriss. Automated justification of collective decisions via constraint solving. In *AAMAS 2020*, pages 168–176. IFAAMAS, 2020.
- [17] Sylvain Bouveret and Jérôme Lang. Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. *J. Artif. Intell. Res.*, 32:525–564, 2008.
- [18] Sylvain Bouveret, Ulle Endriss, and Jérôme Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *ECAI 2010*, volume 215 of *FAIA*, pages 387–392. IOS Press, 2010.
- [19] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.
- [20] Steven J. Brams, Michael A. Jones, and Christian Klamler. Divide-and-conquer: A proportional, minimal-envy cake-cutting algorithm. *SIAM Rev.*, 53(2):291–307, 2011.
- [21] Florian Brandl and Dominik Peters. Approval voting under dichotomous preferences: A catalogue of characterizations. *J. Econ. Theory*, 205:105532, 2022.
- [22] Robert Brederbeck, Aleksander Figiel, Andrzej Kaczmarczyk, Dusan Knop, and Rolf Niedermeier. High-multiplicity fair allocation made more practical. In *AAMAS 2021*, pages 260–268. ACM, 2021.
- [23] Eric Budish. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. In *BQGT 2010*, page 74:1. ACM, 2010.

- [24] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On low-envy truthful allocations. In *ADT 2009*, volume 5783 of *LNCS*, pages 111–119. Springer, 2009.
- [25] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.
- [26] Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaître, Nicolas Maudet, Julian A. Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1):3–31, 2006.
- [27] Yann Chevaleyre, Ulle Endriss, Sylvia Estivie, and Nicolas Maudet. Reaching envy-free states in distributed negotiation settings. In *IJCAI 2007*, pages 1239–1244. AAAI Press, 2007.
- [28] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Distributed fair allocation of indivisible goods. *Artif. Intell.*, 242:1–22, 2017.
- [29] Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM*, 50(5):752–794, 2003.
- [30] Edmund M. Clarke, Anubhav Gupta, and Ofer Strichman. SAT-based counterexample-guided abstraction refinement. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, 23(7):1113–1123, 2004.
- [31] Ari Conati, Andreas Niskanen, Ronald de Haan, and Matti Järvisalo. Complexity results and algorithms for manipulation and bribery in judgment aggregation. In *ECAI 2024*, volume 392 of *FAIA*, pages 3597–3604. IOS Press, 2024.
- [32] Ari Conati, Andreas Niskanen, and Matti Järvisalo. Declarative approaches to outcome determination in judgment aggregation. *J. Artif. Intell. Res.*, 81:793–836, 2024.
- [33] Ari Conati, Andreas Niskanen, Ronald de Haan, and Matti Järvisalo. Computing efficient and envy-free allocations under dichotomous preferences using SAT. In *AAMAS 2025 (to appear)*. IFAAMAS, 2025.
- [34] Sylvie Coste-Marquis, Jérôme Lang, Paolo Liberatore, and Pierre Marquis. Expressive power and succinctness of propositional languages for preference representation. In *KR 2004*, pages 203–212. AAAI Press, 2004.
- [35] Bart de Keijzer, Sylvain Bouveret, Tomas Klos, and Yingqian Zhang. On the complexity of efficiency and envy-freeness in fair division of indivisible goods with additive preferences. In *ADT 2009*, volume 5783 of *LNCS*, pages 98–110. Springer, 2009.
- [36] John P. Dickerson, Jonathan R. Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. The computational rise and fall of fairness. In *AAAI 2014*, pages 1405–1411. AAAI Press, 2014.
- [37] Dinko Dimitrov, Peter Borm, Ruud Hendrickx, and Shao Chin Sung. Simple priorities and core stability in hedonic games. *Soc. Choice Welf.*, 26(2):421–433, 2006.
- [38] Conal Duddy. Fair sharing under dichotomous preferences. *Math. Soc. Sci.*, 73:1–5, 2015.
- [39] Edith Elkind and Jörg Rothe. Cooperative game theory. In *Economics and Computation, An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, Springer texts in business and economics, pages 135–193. Springer, 2016.
- [40] Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7:45–98, 1967.
- [41] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *NSDI 2011*, pages 323–336. USENIX Association, 2011.
- [42] Alexey Ignatiev, António Morgado, and João Marques-Silva. Pysat: A python toolkit for prototyping with SAT oracles. In *SAT 2018*, volume 10929 of *LNCS*, pages 428–437. Springer, 2018.
- [43] Jérôme Lang. Logical preference representation and combinatorial vote. *Ann. Math. Artif. Intell.*, 42(1-3):37–71, 2004.
- [44] Jérôme Lang and Jörg Rothe. Fair division of indivisible goods. In *Economics and Computation, An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, Springer texts in business and economics, pages 493–550. Springer, 2016.
- [45] Mark H. Liffiton and Karem A. Sakallah. Algorithms for computing minimal unsatisfiable subsets of constraints. *J. Autom. Reason.*, 40(1):1–33, 2008.

- [46] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *ÉC 2004*, pages 125–131. ACM, 2004.
- [47] João Marques-Silva, Federico Heras, Mikolás Janota, Alessandro Previti, and Anton Belov. On computing minimal correction subsets. In *IJCAI 2013*, pages 615–622. IJCAI/AAAI, 2013.
- [48] Ruben Martins, Saurabh Joshi, Vasco Manquinho, and Inês Lynce. Incremental cardinality constraints for maxsat. In *CP 2014*, volume 8656 of *LNCS*, pages 531–548. Springer, 2014.
- [49] Arnon Netzer, Amnon Meisels, and Roie Zivan. Distributed envy minimization for resource allocation. *Auton. Agents Multi Agent Syst.*, 30(2):364–402, 2016.
- [50] Trung Thanh Nguyen and Jörg Rothe. Minimizing envy and maximizing average nash social welfare in the allocation of indivisible goods. *Discret. Appl. Math.*, 179:54–68, 2014.
- [51] Trung Thanh Nguyen and Jörg Rothe. Complexity results and exact algorithms for fair division of indivisible items: A survey. In *IJCAI 2023*, pages 6732–6740. ijcai.org, 2023.
- [52] Dominik Peters. Complexity of hedonic games with dichotomous preferences. In *AAAI 2016*, pages 579–585. AAAI Press, 2016.
- [53] Dominik Peters. Precise complexity of the core in dichotomous and additive hedonic games. In *ADT 2017*, volume 10576 of *LNCS*, pages 214–227. Springer, 2017.
- [54] Olivier Roussel and Vasco Manquinho. Pseudo-boolean and cardinality constraints. In *Handbook of Satisfiability - Second Edition*, volume 336 of *FAIA*, pages 1087–1129. IOS Press, 2021.
- [55] Parham Shams, Aurélie Beynier, Sylvain Bouveret, and Nicolas Maudet. Minimizing and balancing envy among agents using ordered weighted average. In *ADT 2021*, volume 13023 of *LNCS*, pages 289–303. Springer, 2021.
- [56] Carsten Sinz. Towards an optimal CNF encoding of boolean cardinality constraints. In *CP 2005*, volume 3709 of *LNCS*, pages 827–831. Springer, 2005.
- [57] William Thomson. Introduction to the theory of fair allocation. In *Handbook of Computational Social Choice*, pages 261–283. Cambridge University Press, 2016.
- [58] Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9(1):63–91, 1974.
- [59] Toby Walsh. Fair division: The computer scientist’s perspective. In *IJCAI 2020*, pages 4966–4972. ijcai.org, 2020.

Ari Conati
 University of Helsinki
 Helsinki, Finland
 Email: ari.conati@helsinki.fi

Andreas Niskanen
 University of Helsinki
 Helsinki, Finland
 Email: andreas.niskanen@helsinki.fi

Ronald de Haan
 Institute for Logic, Language, and Computation (ILLC), University of Amsterdam
 Amsterdam, The Netherlands
 Email: me@ronalddehaan.eu

Matti Järvisalo
University of Helsinki
Helsinki, Finland
Email: matti.jarvisalo@helsinki.fi