

Graphical House Allocation

Hadi Hosseini
Penn State University
University Park, United States
hadi@psu.edu

Justin Payan
University of Massachusetts Amherst
Amherst, United States
jpayan@umass.edu

Rik Sengupta
University of Massachusetts Amherst
Amherst, United States
rsengupta@umass.edu

Rohit Vaish
IIT Delhi
Delhi, India
rvaish@iitd.ac.in

Vignesh Viswanathan
University of Massachusetts Amherst
Amherst, United States
vviswanathan@umass.edu

ABSTRACT

The classical house allocation problem involves assigning n houses (or items) to n agents according to their preferences. A key criteria in such problems is satisfying some fairness constraints such as envy-freeness. We consider a generalization of this problem wherein the agents are placed along the vertices of a graph (corresponding to a social network), and each agent can only experience envy towards its neighbors. Our goal is to minimize the *aggregate envy* among the agents as a natural fairness objective, i.e., the sum of the envy value over all edges in a social graph.

When agents have identical and *evenly-spaced* valuations, our problem reduces to the well-studied problem of *linear arrangements*. For identical valuations with possibly uneven spacing, we show a number of deep and surprising ways in which our setting is a departure from this classical problem. More broadly, we contribute several structural and computational results for various classes of graphs, including NP-hardness results for disjoint unions of paths, cycles, stars, or cliques; we also obtain fixed-parameter tractable (and, in some cases, polynomial-time) algorithms for paths, cycles, stars, cliques, and their disjoint unions.

Additionally, a conceptual contribution of our work is the formulation of a structural property for disconnected graphs that we call *separability* which results in efficient parameterized algorithms for finding optimal allocations.

KEYWORDS

Fair Allocation; House Allocation; Envy Minimization

ACM Reference Format:

Hadi Hosseini, Justin Payan, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. 2023. Graphical House Allocation. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 9 pages.

1 INTRODUCTION

The house allocation problem has attracted interest from the computer science and multiagent systems communities for a long time. The classical problem deals with assigning a set of n houses to n agents with (possibly) different valuations over the houses. It is often desirable to find assignments that satisfy some economic property of interest. In this work, we focus on the well-motivated

economic notion of *fairness*, and in particular, study the objective of minimizing the *aggregate envy* among the agents.

Despite the historical interest in this problem, to the best of our knowledge, the house allocation problem has not been studied thoroughly over *graphs*, a setting in which the agents are placed on the vertices of an undirected graph G and each agent’s potential envy is only towards its neighbors in G [4, 25].

Incorporating the social structure over a graph enables us to capture the underlying restrictions of dealing with partial information, which is representative of constraints in many real-world applications. Thus, the classical house allocation problem is the special case of our problem when the underlying graph is complete.

Our work is in line with recent literature on examining various problems in computational social choice on social networks, including voting [10, 16, 32], fair division [2, 7], and hedonic games [18, 27]. By focusing on graphs, we aim to gain insights into how the structure of the social network impacts fairness in house allocation. We focus on *identical* valuation functions and show that even under this seemingly strong restriction, the problem is computationally hard, yet structurally rich. We provide a series of observations and insights about graph structures that help identify, and in some cases overcome, these computational bottlenecks.

1.1 Overview and Our Contributions

We assume that the agents are placed at the vertices of a graph representing a social network, and that they have identical valuation functions over the houses. Our objective is to find an allocation of the houses among the agents to minimize the total envy in the graph. We call this the *graphical house allocation problem*.

This is a beautiful combinatorial problem in its own right, as it can be restated as the problem where, given an undirected graph and a multiset of nonnegative numbers, the numbers need to be placed on the vertices in a way that minimizes the sum of the edgewise absolute differences.

In Section 2, we present the formal model and set up some preliminaries, including the connection between the graphical house allocation problem and the *minimum linear arrangement* problem, which has several notable similarities and differences.

In Section 3, we present computational lower bounds and inapproximability results for the problem, even for very simple graphs. In particular, we show NP-hardness even when the graph is a disjoint union of paths, cycles, stars, or cliques, which all have known polynomial-time algorithms for linear arrangements.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

In Section 4, we focus on connected graphs and completely characterize optimal allocations when the graph is a path, cycle, star, or a complete bipartite graph. We also prove a technically involved structural result for *rooted binary trees*, and discuss general trees.

In Section 5, we focus on disconnected graphs, starting with a fundamental difference between graphical house allocation and linear arrangements, motivating our definitions of *separable*, *strongly separable*, and *inseparable* disconnected graphs. We employ these characterizations to prove algorithmic results for a variety of graphs. In particular, we show that disjoint unions of paths, cycles, stars, and equal-sized cliques are strongly separable and develop natural fixed parameter tractable algorithms for these graphs. Moreover, we show that disjoint unions of arbitrary cliques satisfy separability (but not strong separability) and admit XP algorithms.

In the interest of space, we mainly provide proof sketches and defer the details to the full version of the paper [17].

1.2 Related Work

House allocation has been traditionally studied in the economics literature under the *housing market* model, where agents enter the market with a house (or an endowment) each and are allowed to engage in cyclic exchanges [30]. This model has found important practical applications, most notably in kidney exchange [1, 28].

While the initial work on house allocation focused on the economic notions of core and strategyproofness [31], subsequent work has explored *fairness* issues. Gan et al. [13] study the house allocation problem under ordinal preferences (specifically, weak rankings) and provide a polynomial-time algorithm for determining the existence of an envy-free allocation. By contrast, the problem becomes NP-hard when agents' preferences are specified as a set of pairwise comparisons [21]. Kamiyama et al. [22] study house allocation under cardinal preferences (similar to our work) and examine the complexity of finding a “fair” assignment for various notions of fairness such as proportionality, equitability, and maximizing the number of envy-free agents (they do not consider aggregate envy). They show that the latter problem is hard to approximate under general valuations, and remains NP-hard even for the restricted case of binary valuations. For binary valuations, the problem of finding the largest envy-free partial matching has also been studied [3].

Recent studies have considered *graphical* aspects of house allocation (similar to our work), though with different objectives. Massand and Simon [25] consider house allocation under externalities and study various kinds of stability-based objectives. Beynier et al. [4] study *local envy-freeness* in house allocation, which entails checking the existence of an allocation with no envy along any edge of the graph. Their work is close to ours, but their model involves agents with *possibly distinct* ordinal preferences and only the zero-envy condition.

There is also a growing literature on fair allocation of indivisible objects among agents who are part of a social network. Bredereck et al. [7] present fixed-parameter tractability results, mainly parametrized by the number of agents, though they leave results using graph structure to future work. Eiben et al. [11] extend these results, showing a number of parametrized complexity results relating the treewidth, cliquewidth, number of agent types, and number

of item types to the complexity of determining if an envy-free allocation exists on a graph. This line of work again focuses on deciding if envy-free allocations exist, rather than minimizing envy. Other works seek to obtain envy-free allocations, maximum welfare allocations, or other objectives by swapping objects along a graphical structure [5, 15, 19, 24].

2 PRELIMINARIES

We use $[t]$ to denote the set $\{1, 2, \dots, t\}$. There is a set of n agents $N = [n]$ and n houses $H = \{h_1, h_2, \dots, h_n\}$ (often called *items*). Each agent i has a *valuation function* $v_i : H \rightarrow \mathbb{R}_{\geq 0}$; $v_i(h)$ indicates agent i 's value for house $h \in H$.

An *allocation* π is a bijective mapping from agents to houses. For each $i \in N$, $\pi(i)$ is the house allocated to agent i under the allocation π , and $v_i(\pi(i))$ is its utility.

Given a problem instance consisting of agents and houses, our goal is to generate an allocation π that is “fair” to all the agents, for some reasonable definition of fairness. A natural way to define fairness is using *envy*. An agent i is said to envy agent j under allocation π if $v_i(\pi(i)) < v_i(\pi(j))$. While we would ideally like to find *envy-free* allocations, this may not always be possible — consider a simple example with two agents and two houses but (exactly) one of the houses is valued at 0 by both agents. Therefore, we instead focus on the magnitude of envy that agent i has towards agent j , for a fixed allocation π . This is defined as $envy_i(\pi, j) := \max\{v_i(\pi(j)) - v_i(\pi(i)), 0\}$.

We define an undirected graph $G = (N, E)$ over the set of agents, which represents the underlying social network. Our goal is to compute an allocation that minimizes the *total envy* along the edges of the graph, defined as $Envy(\pi, G) := \sum_{(i,j) \in E} (envy_i(\pi, j) + envy_j(\pi, i))$; note that edges are undirected. An allocation π^* that minimizes the total envy is referred to as a *minimum envy allocation*. When there are multiple minimum envy allocations, we abuse notation slightly and use π^* to denote any of them.

When the graph G is a complete graph K_n , a minimum envy allocation can be computed in polynomial time by means of a reduction to a bipartite minimum-weight matching problem (proved in the full version [17]). However, it is known that for several other simple graphs like paths and matchings, computing a minimum envy allocation is NP-complete [4]. Given this computational intractability, we therefore explore a natural restriction of the problem, when all agents have identical valuations, to gain insights into the computational and structural aspects of fairness in social networks. More formally, we assume that there is a fixed valuation function v such that $v_i = v$ for all $i \in N$. Identical valuations capture a natural aspect of real-world housing markets, where the house values are independent of agents.

When all agents have the same valuation function v , the total envy of an allocation π along the edges of a graph $G = (N, E)$ can be written as $Envy(\pi, G) = \sum_{(i,j) \in E} |v(\pi(i)) - v(\pi(j))|$. This formulation also yields a new expression for envy along an edge $e = (i, j) \in E$ as $envy_e(\pi) = |v(\pi(i)) - v(\pi(j))|$. This value equals $envy_i(\pi, j) + envy_j(\pi, i)$, as one of those terms is zero under identical valuations.

When G is K_n , under identical valuations, an optimal allocation is trivially computable, as all allocations are equivalent.

For the rest of this paper, we will assume without loss of generality that the house values are all distinct (refer to the full version [17] for a formal justification). In particular, every agent’s valuation function (denoted by v) gives each house a unique nonnegative value, with $v(h_1) < v(h_2) < \dots < v(h_n)$. We will say $h_1 < h_2$ to mean $v(h_1) < v(h_2)$.

For an allocation π and a subset $N' \subseteq N$, we will refer to the set of houses received by N' as $\pi(N')$. If G' is a subgraph of G , we will use $\pi(G')$ in the same way.

We will use $G_1 + G_2$ to mean the disjoint union of G_1 and G_2 .

Definition 2.1. For an instance of the graphical house allocation problem, the valuation interval is defined as the closed interval $[v(h_1), v(h_n)] \subset \mathbb{R}_{\geq 0}$ with each $v(h_k)$ marked.

The motivation for Definition 2.1 is as follows. For an arbitrary allocation π , for each edge $e = (i, j) \in E$, we can draw a line segment from $v(\pi(i))$ to $v(\pi(j))$. This line segment has length $|v(\pi(i)) - v(\pi(j))| = \text{envy}_e(\pi)$. It follows that $\text{Envy}(\pi, G)$ is the sum of the lengths of all such line segments. An optimal allocation π^* is any allocation that attains this minimum sum. See Figure 1 for an example of a valuation interval, together with a graph G , and a particular allocation on G depicted under the valuation interval.




Figure 1: (Left) A graph G on five agents along with a particular allocation π . The valuations are identical and are given by $\vec{v} = (1, 2, 4, 5, 6)$. (Right) The valuation interval is shown via the thick horizontal line in black. The five line segments in red denote the envy along the five edges of the graph G . The total length of these line segments is $\text{Envy}(\pi, G) = 15$.

A subset of houses $H' = \{h_{i_1}, \dots, h_{i_k}\} \subseteq H$ with $h_{i_1} < \dots < h_{i_k}$ is called *contiguous* if there is no house $h' \in H \setminus H'$ with $h_{i_1} < h' < h_{i_k}$. Geometrically, the values in H' form an uninterrupted sub-interval of the valuation interval, with no value outside of H' appearing inside that sub-interval. In Figure 1 above, the subsets $\{h_1, h_2, h_3\}$ and $\{h_5\}$ are contiguous, whereas the subsets $\{h_1, h_2, h_5\}$ and $\{h_3, h_5\}$ are not.

For simplicity, we will interchangeably talk about allocating h_i and allocating $v(h_i)$ to an agent, and we will also sometimes refer to houses as being marked points on the valuation interval.

2.1 Connection to the Linear Arrangement Problem

The *minimum linear arrangement problem* is the problem where, given an undirected n -vertex graph $G = (V, E)$, we want to find a bijective function $\pi : V \rightarrow [n]$ that minimizes $\sum_{(i,j) \in E} |\pi(i) - \pi(j)|$. The minimum linear arrangement problem is a special case of the graphical house allocation problem where the valuation interval has evenly spaced values ($H = [n]$).

We may assume without loss of generality that the underlying graph G in any instance of the minimum linear arrangement problem is connected. This is a consequence of the following folklore observation, whose proof we omit (see, for instance, [29]).

Proposition 2.2 (Seidvasser [29]). *If G is any instance of the linear arrangement problem, then at least one optimal solution assigns contiguous subsets of $[n]$ to the connected components of G .*

We will see in Section 5 that in the graphical house allocation problem, it no longer suffices to only consider connected graphs.

Finding a minimum linear arrangement is NP-hard for general graphs [14], with a best known run-time of $O(2^n m)$, where $|V| = n$, $|E| = m$. The problem remains NP-hard even for bipartite graphs [12]. However, the minimum linear arrangement problem can be solved in polynomial time on paths, cycles, stars, wheels, and trees. Minimum linear arrangement on trees takes time $O(n^{\log_2 3})$ [9].

One of the fundamental difficulties with the minimum linear arrangement problem is that there are no natural approaches to prove lower bounds on the objective: the best known general lower bounds are trivial ones based on degrees or edge counts. All other known lower bounds are graph-specific and not easy to generalize.

3 HARDNESS AND LOWER BOUNDS

In this section, we prove lower bounds on the general graphical house allocation problem.

Our problem is NP-complete (for arbitrary graphs and valuations), because the special case of minimum linear arrangements is already NP-complete, as stated in Section 2.

A different NP-completeness proof, reducing from the MINIMUM BISECTION PROBLEM, results in inapproximability results as well.

Definition 3.1. The MINIMUM BISECTION PROBLEM asks, for an n -vertex graph G and a natural number k , if there is a partition of $V(G)$ into two parts of size $n/2$, with at most k edges crossing the cut.

The MINIMUM BISECTION PROBLEM is a known NP-complete problem [14]. Furthermore, it is also known to be hard to approximate efficiently, a fact that is useful in light of the following observation.

THEOREM 3.2. *There is a linear time reduction from the MINIMUM BISECTION PROBLEM to the graphical house allocation problem with identical valuations.*

PROOF SKETCH. Given an instance $\langle G, k \rangle$ of MINIMUM BISECTION, we construct an instance of the house allocation problem by creating a cluster of $n/2$ values concentrated around 0, and a cluster of $n/2$ values concentrated around 1 on the valuation interval. Then, G has an allocation with total envy at most $k + \epsilon$ (for sufficiently small ϵ) if and only if $\langle G, k \rangle \in \text{MINIMUM BISECTION}$. This reduction is clearly linear time. \square

It follows immediately that the inapproximability results for minimum bisection carry over to the graphical house allocation problem. In particular, for any fixed constant $\epsilon > 0$, unless $P = NP$, there is no polynomial-time algorithm that can approximate the optimal total envy under the house allocation problem within an additive term of $n^{2-\epsilon}(v(h_n) - v(h_1))$ [8]. Additionally, the house allocation problem has no PTAS unless NP has randomized algorithms in subexponential time [23]. Thus our problem is hard to approximate *even* with identical valuations.

Finally, we show that graphical house allocation is NP-complete even on simple instances of graphs which are solvable in linear time in the case of linear arrangements, such as disjoint unions of paths, cycles, cliques, or stars (and any combinations of them).

THEOREM 3.3 (HARDNESS OF DISJOINT UNIONS). *Let \mathcal{A} be any collection of connected graphs, such that there is a polynomial time one-to-one mapping from each nonnegative integer t (given in unary) to a graph in \mathcal{A} of size t . Let \mathcal{G} be the class of graphs whose members are the finite sub-multisets of \mathcal{A} (as connected components). Then, finding a minimum envy house allocation is NP-hard on the class \mathcal{G} .*

PROOF SKETCH. We reduce from UNARY BIN PACKING, which asks, given a finite set I of items with sizes $s(i)$ for $i \in I$, a bin capacity B , and an integer k , all given in unary, whether there exists a packing of all the items into at most k bins. This problem is known to be NP-complete (see, for instance, [20]).

Given an arbitrary instance of UNARY BIN PACKING, we create an instance of graphical house allocation by having k equispaced clusters of width ϵ separated by intervals of size C (for sufficiently large C and small ϵ), with each cluster containing B values. Our graph G is a disjoint union of the graphs in \mathcal{A} that form the image of the sizes $s(i)$ over all $i \in I$. Note that $G \in \mathcal{G}$. Then, the given instance is in UNARY BIN PACKING if and only if the graphical house allocation instance has an allocation with envy less than C . \square

Corollary 3.4. *The house allocation problem under identical valuations is NP-complete on: (a) disjoint unions of arbitrary paths, (b) disjoint unions of arbitrary cycles, (c) disjoint unions of arbitrary stars, and (d) disjoint unions of arbitrary cliques.*

In Section 5 we show that despite the hardness suggested by Corollary 3.4, it is possible to exploit a structural property to develop FPT algorithms for the first three problems.

4 CONNECTED GRAPHS

In this section, we characterize optimal house allocations when the underlying graph G is a star, path, cycle, complete bipartite graph, or rooted binary tree. We also provide some observations when G is any (arbitrary) tree.

4.1 Stars

Consider the star graph $K_{1,n-1}$, which has a single central node and $n - 1$ “spokes” connected to the central node but not to each other.

THEOREM 4.1. *If G is the star $K_{1,n-1}$, then the minimum envy allocation π^* under identical valuations corresponds to:*

- for odd n , putting the unique median value in the center of the star, and all the houses on the spokes in any order; the value of the envy is $\sum_{i > (n-1)/2+1} v(h_i) - \sum_{i \leq (n-1)/2} v(h_i)$.
- for even n , putting either of the medians in the center of the star, and all other houses on the spokes in any order; the value of the envy for either median is $\sum_{i > (n+1)/2} v(h_i) - \sum_{i < (n+1)/2} v(h_i)$.

PROOF. The proof is a restatement of the well-known fact that in any multiset of real numbers, the sum of the L_1 -distances is minimized by the median of the multiset. It is easy to verify that for even n , both medians yield the same value. \square

4.2 Paths

Consider the path graph P_n .

THEOREM 4.2. *If G is the path graph P_n , then the minimum envy allocation π^* under identical valuations attains a total envy of $v(h_n) - v(h_1)$, is unique (up to reversing the values along the path), and corresponds to placing the houses in sorted order along P_n .*

PROOF SKETCH. The houses h_1 and h_n have to go to two of the vertices. The subpath between these two vertices must have envy at least $v(h_n) - v(h_1)$. The rest follows. \square

4.3 Cycles


Now, consider the cycle graph C_n .

THEOREM 4.3. *If G is the cycle graph C_n , then any minimum envy allocation π^* under identical valuations attains a total envy of $2(v(h_n) - v(h_1))$, and corresponds to the following: place h_1 and h_n arbitrarily on any two vertices of the cycle, and then place the remaining houses so that each of the two paths from h_1 to h_n along the cycle consists of houses in sorted order.*

PROOF SKETCH. In any allocation, there are two internally vertex-disjoint subpaths from h_1 to h_n on a cycle. Each of those subpaths has an envy of at least $v(h_n) - v(h_1)$. \square

Corollary 4.4. *For $n \geq 3$, the number of optimal allocations along the cycle C_n is 2^{n-3} , up to rotations and reversals.*

Perhaps slightly non-obviously, the proofs of Theorems 4.2 and 4.3 can be seen as purely geometric arguments using the valuation interval. To see this, consider the path P_n , and take any allocation π that does not satisfy the form stated in Theorem 4.2, and consider how the allocation looks on the valuation interval. First, observe that every sub-interval of the valuation interval between consecutive houses needs to be covered by some line segment from the allocation. Otherwise, there would be no edge with a house from the left to a house from the right of the sub-interval, which is impossible, as P_n is connected. But the only way to meet this lower bound of one line segment for each sub-interval of the valuation interval is to sort the houses along the path. The allocation looks as follows on the valuation interval.



The geometric argument for cycles is similar (with an allocation illustrated below).



4.4 Complete Bipartite graphs

Let us start with the complete bipartite graph $K_{r,r}$ ($r \geq 1$) where both parts have equal size.




Figure 2: Examples of characterized connected graphs

THEOREM 4.5. *When G is the graph $K_{r,r}$, the minimum envy allocation π^* has the following property: for every $i \in [r]$ the houses $\{h_{2i-1}, h_{2i}\}$ cannot be allocated to agents in the same side of the bipartite graph. Moreover, all allocations which satisfy this property have the same (optimal) envy.*

PROOF SKETCH. Consider an allocation π where the stated property does not hold and consider the smallest i where the property is violated. In this case we can find a simple swap that can improve envy. Assume h_{2i-1} and h_{2i} are allocated to one part. If k is the least index greater than $2i$ such that h_k is not allocated to an agent in the same part, then swapping h_k and h_{k-1} leads to a reduction in envy. The exact calculation is quite technical; see the full version of the paper [17]. \square

This also implies a straightforward linear-time algorithm to compute a minimum envy allocation for the graph $K_{r,r}$.

We can now generalize this result to complete bipartite graphs where the two parts have unequal size. Due to the similarity of the two proofs, we omit the proof sketch.

THEOREM 4.6. *When G is the graph $K_{r,s}$ ($r > s$), the minimum envy allocation π^* has the following property:*

- If $r - s =: 2m$ is even, then the first and last m houses are allocated to the larger part, and for all $i \in [s]$, the houses h_{m+2i-1} and h_{m+2i} are allocated to different parts.
- If $r - s =: 2m + 1$ is odd, then the first m and last $m + 1$ houses are allocated to the larger part. For all $i \in [s]$, the houses h_{m+2i-1} and h_{m+2i} are allocated to the larger and smaller parts respectively.

Moreover, all allocations which satisfy this property have the same (optimal) envy.

Corollary 4.7. *For any complete bipartite graph $K_{r,s}$ ($r \geq s$),*

- If $r - s$ is even, there are 2^s optimal allocations;
- If $r - s$ is odd, there is exactly one optimal allocation,

up to permutations over allocations to the same side of the graph.

It is easy to see that our linear time algorithm generalizes to general complete bipartite graphs as well. Theorem 4.6 generalizes Theorem 4.1. When the number of spokes in the star is odd, there are two possible houses that can be allocated to the center in an optimal allocation. However, when the number of spokes is even, any optimal allocation allocates a unique house to the central node.

4.5 Rooted Binary Trees

In this section, we consider binary trees. A *binary tree* T is defined as a rooted tree where each node has either 0 or 2 children.

Our main result is a structural property characterizing at least one of the optimal allocations for any instance where the graph G corresponds to a binary tree. We call this the *local median property*.

Definition 4.8 (Local Median Property). *An allocation on a binary tree satisfies the local median property if, for any internal node, exactly one of its children is allocated a house with value less than that of the node. In other words, the value allocated to any internal node is the median of the set containing the node and its children.*

The proofs in this section will use the following lemma. We define the *inverse* of a valuation function v as a valuation function v^{inv} such that $v^{\text{inv}}(h) = -v(h)$ for all $h \in H$ (appropriately shifted so that all values are nonnegative). We note that any allocation has the same envy along any edge with respect to the inverted valuation and the original valuation, whose straightforward proof we relegate to the full version [17].

Lemma 4.9. *The envy along any edge of the graph G under an allocation π with respect to the valuation v is equal to the envy along the same edge of the graph G under the allocation π with respect to the valuation v^{inv} .*

We will now show that at least one minimum envy allocation satisfies the local median property. More formally, we show the following: given a binary tree T and any allocation π , there exists an allocation that satisfies the local median property and has equal or lower total envy. The proof relies on the following lemma.

Lemma 4.10. *Given a binary tree T and an allocation π , let i be some internal node which does not satisfy the local median property. Then, there exists an allocation π' such that*

- (a) *For the subtree T' rooted at i , we have that $\text{Envy}(\pi(T'), T') > \text{Envy}(\pi'(T'), T')$;*
- (b) *For any other subtree T'' not contained by T' , we have that $\text{Envy}(\pi(T''), T'') \geq \text{Envy}(\pi'(T''), T'')$.*

PROOF SKETCH. Consider an internal vertex not satisfying the local median property; by Lemma 4.9, it suffices to consider the case when this violating vertex has a value that is less than both its children. We can now “push” the value allocated to this vertex down the tree by a series of swaps until that value reaches a leaf or satisfies the local median property. It can be shown that the total envy at the end satisfies the two criteria stated. A full proof is given in the full version [17]. \square

Lemma 4.10 immediately gives rise to the following corollary, which we state as a theorem.

THEOREM 4.11. *For any binary tree T , at least one minimum envy allocation satisfies the local median property.*

Unfortunately, the local median property is too weak to exploit for a polynomial-time algorithm. Ideally, we would like to use the property to show that some minimum envy allocation satisfies an even stronger property called the *global median property*.

Definition 4.12 (Global Median Property). *An allocation on a binary tree satisfies the global median property if, for every internal node, all the houses in one subtree of the node have value less than the house allocated to the node, and all the houses in the other subtree have value greater than the house allocated to the node.*

Empirically, it seems like minimum envy allocations on rooted binary trees do indeed satisfy the global median property. If this turned out to be true, we would be able to devise an algorithm that would significantly reduce the search space for a minimum envy allocation from $\Theta(n!)$ to $\Theta(2^d)$, for a tree T of maximum depth d . This would give us an $O(2^d)$ algorithm for computing minimum envy allocations on such rooted binary trees.

Conjecture 4.13. *There is an algorithm that computes an optimal allocation on a rooted binary tree of maximum depth d in time $O(2^d)$. In particular, this algorithm runs in polynomial time on (nearly) balanced trees.*

4.6 General Trees

How do we take the approaches for rooted binary trees and build towards arbitrary trees? Note that one consequence of Theorem 4.11 is that in an optimal allocation on a rooted binary tree, the minimum and the maximum must both appear on leaves.

In the minimum linear arrangement problem, it is known [29] that when the underlying graph is a tree, some optimal allocation assigns both the minimum and maximum values to leaves, and furthermore, the (unique) path from this minimum to the maximum consists of monotonically increasing values. This characterization is used crucially in designing the polynomial time algorithm on trees [9].

Empirically, this same property for trees seems to hold for non-uniformly spaced values as well. The proof technique used in [29] does not extend to our setting, but testing the problem on 200 randomly generated trees and uniformly random values on the interval $[0, 100]$ always gave us both these properties on trees: the minimum and maximum values both end up on leaves, with a monotonic path between them.

We believe that graphical house allocation, unlike minimum linear arrangements, is NP-hard on trees. It would be remarkable if the structural characterization holds for our problem, but the polynomial-time algorithm does not work. We relegate answering this to future work.

5 DISCONNECTED GRAPHS

In this section, we consider disconnected graphs, starting with a structural characterization, and then using that to obtain upper bounds for several natural classes of disconnected graphs that all had lower bounds in Section 3.




Figure 3: For the valuation interval on top, the optimal allocation to $P_2 + C_3$ is to give the two low-valued houses to the edge, and to give the three high-valued houses to the triangle. This is the only allocation where the envy is negligible. For the valuation interval on the bottom, the optimal allocation to $P_2 + C_3$ is to give the two extreme-valued houses to the edge, and the cluster in the middle to the triangle. Any other allocation has to count one of the long halves of the interval multiple times, and is therefore strictly suboptimal. This is an instance where we see one of the connected components being “split” by another in the valuation interval.

5.1 A Structural Characterization

We start by remarking that Proposition 2.2 is *false* in our setting, and so we can no longer assume our graph is connected without loss of generality. For instance, consider an instance when the underlying graph G is a disjoint union of an edge and a triangle. The two valuation intervals in Figure 3 yield very different optimal structures for this same instance.

We remark that this is a major departure from the linear arrangement problem, as the spacing of the values along the valuation interval becomes a key factor in the structure of optimal allocations. Recall from Proposition 2.2 that in the minimum linear arrangement problem, in an optimal allocation, the connected components of a graph take the houses in contiguous subsets along the valuation interval. The example in Figure 3 shows that this is not necessarily true in our problem, and therefore serves as a motivation to classify disconnected graphs according to whether they have this property or not. We call the relevant property *separability*, defined as follows.

Definition 5.1 (Splitting). *Let $G_1 = (N_1, E_1)$ and $G_2 = (N_2, E_2)$ be two of the connected components of $G = (N, E)$, and fix an arbitrary allocation π . We say G_1 splits G_2 in π if the values of $\pi(G_1)$ form a contiguous subset of the values in $\pi(G_1) \cup \pi(G_2)$.*

Definition 5.2 (Separability). *Let G be a disconnected graph with connected components G_1, \dots, G_k . Then,*

- (1) G is separable if there exists an ordering G_1, \dots, G_k of the components where, for all valuation intervals, there is an optimal allocation where for all $1 \leq i < j \leq k$, G_i splits G_j .
- (2) G is strongly separable if, in addition, G_j also splits G_i . Note that this is only possible if an optimal allocation assigns contiguous subsets of values to all connected components.
- (3) G is inseparable if there is a valuation interval where for each optimal allocation π , there are components G_1 and G_2 with $u, u' \in \pi(G_1)$ and $v, v' \in \pi(G_2)$ such that $u < v < u' < v'$.

A class \mathcal{A} of graphs is separable (resp. strongly separable) if every graph in it is separable (resp. strongly separable). Conversely, \mathcal{A} is inseparable if it contains an inseparable graph.

We note that a graph G is inseparable if and only if it is not separable. Furthermore, it is strongly separable only if it is separable.

For minimum linear arrangements, all disconnected graphs are strongly separable, by Proposition 2.2. In contrast, for our problem, Figure 3 already provides an example of a graph that is not separable. We discuss several examples of strongly separable graphs in our problem in Section 5.2; in particular, disjoint unions of paths (respectively, cycles or stars) satisfy strong separability.

Our formulation of separability and strong separability has an immediate algorithmic consequence.

Proposition 5.3. *Suppose G is strongly separable and has k connected components. If we can find a minimum envy allocation for each component in time $O(\text{poly}(n))$, then we can find a minimum envy allocation on G in time $O(\text{poly}(n) \cdot k!)$. If G is separable, and we can find a minimum envy allocation for each component in time $O(\text{poly}(n))$, then we can find a minimum envy allocation on G in time $n^{O(k)}$.*

The proof follows straightforwardly from the definitions of (strong) separability.

It is not immediately obvious that there are separable graphs that are not strongly separable. Figure 3 shows an example of such a graph (Theorem 5.15 proves separability). We will see more examples of this later, but we remark that there are even separable forests that are not strongly separable (Figure 4). We state this formally here, relegating the proof once again to the full version [17].

Proposition 5.4. *There exists a separable forest that is not strongly separable.*




Figure 4: Example of a separable forest that is not strongly separable. The forest is trivially separable. For the bottom valuation line, an optimal allocation must allocate the extreme clusters in the interval to the larger connected component. See the full version [17] for a detailed proof.

Less obviously, *inseparable forests* exist, as shown by the following proposition (and Figure 5), whose proof is in the full version [17].

Proposition 5.5. *There exists an inseparable forest.*

5.2 Disjoint Unions of Paths, Cycles, and Stars

We now move on to algorithmic approaches and characterizations of minimum envy allocations, and start with the setting where G is a disjoint union of paths. Suppose $G = P_{n_1} + \dots + P_{n_r}$. What does an optimal allocation on G look like?




Figure 5: Example of an inseparable forest. Suppose $s_1 < s_2 < s_3 < s_4$, and they satisfy for all i, j , $|s_i - s_j| \geq 3$, and for all i, j, k , $s_i + s_j > s_k + 2$. Then, an optimal allocation on this instance must allocate the entire cluster of size $s_i + 1$ on the valuation interval to the corresponding star-like cluster of the given forest. See the full version [17] for a detailed proof.

THEOREM 5.6. *Let G be a disjoint union of paths, $P_{n_1} + \dots + P_{n_r}$. Then, G is strongly separable. Furthermore, in any optimal allocation, within each path, the houses appear in sorted order.*

PROOF SKETCH. In any allocation π , we may assume by Theorem 4.2 that the houses allocated to each path appear in sorted order along that path. Next, in an allocation π , if there is any overlap between two of the paths, then we can remove the overlap by reassigning the houses among just those two paths, and obtain an allocation with strictly less envy. \square

The following corollary shows an FPT algorithm on the disjoint union of paths, parameterized by the number of different paths.

Corollary 5.7. *We can find an optimal allocation for an instance on an undirected n -agent graph G that is the disjoint union of paths in time $O(nr!)$, where r is the number of paths.*

PROOF SKETCH. By Theorem 5.6, it suffices to find the optimal ordering of the paths. There are $r!$ such orderings, and for each, we can test the envy in linear time. \square

If G is a disjoint union of cycles, say $G = C_{n_1} + \dots + C_{n_r}$, the same theorems characterizing optimal allocations go through, using Theorem 4.3. We omit the proofs, but state the results formally.

THEOREM 5.8. *Let G be a disjoint union of cycles, $C_{n_1} + \dots + C_{n_r}$. Then G is strongly separable. Furthermore, in any optimal allocation, within each cycle, the houses appear in the form characterized in Theorem 4.3.*

Corollary 5.9. *We can find an optimal house allocation for an instance on an undirected n -agent graph G that is the disjoint union of cycles in time $O(nr!)$, where r is the number of cycles.*

We remark here that if t is the number of different path (or cycle) lengths, then a dynamic programming algorithm (details in the full version [17]) computes the minimum envy allocation in time $O(n^{t+1})$. Therefore, combining the two approaches, we have a time complexity of $O(\min(nr!, n^{t+1}))$. An immediate application of this dynamic programming algorithm is for graphs with degree at most

one. These graphs are special case of the disjoint union of paths where the path length can either be 0 or 1.

Corollary 5.10. *We can find an optimal house allocation for an instance on an undirected n -agent graph G with maximum degree 1 in time $O(n^3)$.*

Perhaps remarkably, there is no particularly elegant characterization when the underlying graph G is a disjoint union of paths and cycles, even when there is only one path and one cycle. This is a consequence of Figure 3.

Finally, a similar result holds for disjoint unions of stars, though the proof is somewhat different. We omit the proof of Corollary 5.12, which follows from Theorem 5.11.

THEOREM 5.11. *Let G be a disjoint union of stars, $K_{1,n_1} + \dots + K_{1,n_r}$. Then G is strongly separable. Furthermore, in any optimal allocation, within each star, the houses appear in the form characterized in Theorem 4.1.*

PROOF SKETCH. We can “separate” any two stars while improving our objective. Given any allocation π , suppose two of the stars K_{1,n_1} and K_{1,n_2} overlap in their allocated values. Suppose their centers receive houses $h_i < h_j$ respectively. We show that exchanging K_{1,n_1} ’s most-valuable house for K_{1,n_2} ’s least-valuable house strictly decreases envy. We apply this repeatedly until the two stars split each other. We apply Theorem 4.1 to every star at the beginning of the process and after every swap to maintain the invariant that each star has a median house at its center. \square

Corollary 5.12. *We can find an optimal house allocation for an instance on an undirected n -agent graph G that is the disjoint union of stars in time $O(nr!)$, where r is the number of stars.*

5.3 Disjoint Unions of Cliques

We now turn our attention to disjoint unions of cliques. We first demonstrate that when all cliques have the same size, we maintain strong separability.

THEOREM 5.13. *Let G be a disjoint union of cliques with equal sizes, $K_{n/r}^1 + \dots + K_{n/r}^r$. Then, G is strongly separable.*

PROOF SKETCH. Without loss of generality, consider two cliques, say K and K' , on $n/2$ vertices each, and consider an arbitrary instance with n values. Suppose in some allocation π , K receives $A \cup A'$ and K' receives $B \cup B'$, where $A \cup B$ form the $n/2$ lower-valued houses, and $A' \cup B'$ form the $n/2$ higher-valued houses. We can show that we improve the envy by assigning them $A \cup B$ and $A' \cup B'$ respectively. \square

Because the cliques are all of equal sizes and agents have identical valuations, Theorem 5.13 implies that there is a trivial algorithm for assigning houses to agents. We can assign the first n/r houses to one clique, the next n/r houses to the next clique, and so on.

Corollary 5.14. *We can find an optimal house allocation for an instance on an undirected n -agent graph G that is the disjoint union of equal-sized cliques in time $O(n)$.*

We now turn our attention to the case when the cliques are not all of the same size.

As we saw in Figure 3, strong separability must be ruled out when cliques have different sizes. We will show that separability still holds. We show further that the largest clique splits all other cliques, the second largest clique splits all cliques except possibly the largest one, and so on. The proof is heavily technical, and we relegate the casework and details to the full version [17].

THEOREM 5.15. *Let G be a disjoint union of cliques with arbitrary sizes, $K_{n_1} + \dots + K_{n_r}$, where $n_1 \geq \dots \geq n_r$. Then, G is separable (but not strongly separable if the n_i ’s are not all equal). In particular, for all $1 \leq i < j \leq r$, in every optimal allocation, K_{n_i} splits K_{n_j} .*

PROOF SKETCH. Without loss of generality, consider two cliques K and K' with $|K| > |K'|$. We will show that K receives a contiguous subset among the houses received by $K \cup K'$. The basic approach is that if there are three values $v(h_i) < v(h_j) < v(h_k)$ with h_j going to K' but h_i, h_k going to K , we can swap houses around and obtain a better allocation by a counting argument. \square

Theorem 5.15 implies an XP algorithm for finding a minimum envy allocation on unions of cliques.

Corollary 5.16. *We can find an optimal house allocation for an instance on an undirected n -agent graph G that is the disjoint union of cliques in time $O(n^{r+2})$, where r is the number of cliques.*

There seems to be a separation between unions of differently-sized cliques and unions of stars, cycles, and paths. We suspect the problem may be W[1]-hard for unions of arbitrary cliques.

6 CONCLUSION AND DISCUSSIONS

We investigate a generalization of the classical house allocation problem where the agents are on the vertices of a graph representing the underlying social network. We wish to allocate the houses to the agents so as to minimize the aggregate envy among neighbors. Even for identical valuations, we show that the problem is computationally hard and structurally rich. Furthermore, our structural insights facilitate algorithmic results for several natural and well-motivated graph classes.

There are a few natural questions for future research. We might consider other fairness objectives such as *minimizing the maximum envy* present on any edge of the graph. For evenly-spaced valuations, this corresponds to the classical graph theoretic property of *bandwidth*, which is also known to be NP-complete for general graphs, and quite hard to approximate as well [6, 26]. It would be interesting to know whether trees admit polynomial time characterizations of the minimum envy, or—more remarkably—whether they are NP-complete but admit the structural similarities to the minimum linear arrangement problem discussed in Section 4.6. We might hope to completely characterize all strongly separable graphs in terms of their graph theoretic structure. Another important future direction would be to extend some of these results for non-identical valuations.

ACKNOWLEDGMENTS

We are grateful to the anonymous reviewers at AAMAS 2023 for helpful comments. We are indebted to Andrew McGregor for pointing us to the linear arrangement problem, and for suggesting multiple useful lines of inquiry. We would also like to thank Cameron Musco for extremely helpful discussions. Finally, we would like to thank Yair Zick for being involved in many of the early discussions, and providing guidance throughout. RS acknowledges support from NSF grant no. CCF-1908849. RV acknowledges support from DST INSPIRE grant no. DST/INSPIRE/04/2020/000107 and SERB grant no. CRG/2022/002621. HH acknowledges support from NSF IIS grants #2144413 and #2107173.

REFERENCES

- [1] Atila Abdulkadiroğlu and Tayfun Sönmez. 1999. House Allocation with Existing Tenants. *Journal of Economic Theory* 88, 2 (1999), 233–260.
- [2] Rediet Abebe, Jon Kleinberg, and David C Parkes. 2017. Fair Division via Social Comparison. In *Proceedings of the 16th International Conference on Autonomous Agents and MultiAgent Systems*. 281–289.
- [3] Elad Aigner-Horev and Erel Segal-Halevi. 2022. Envy-Free Matchings in Bipartite Graphs and their Applications to Fair Division. *Information Sciences* 587 (2022), 164–187.
- [4] Aurélie Beynier, Yann Chevaleyre, Laurent Gourvès, Julien Lesca, Nicolas Maudet, and Anaëlle Wilczynski. 2018. Local Envy-Freeness in House Allocation Problems. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 292–300.
- [5] Aurélie Beynier, Nicolas Maudet, and Anastasia Damamme. 2018. Fairness in Multiagent Resource Allocation with Dynamic and Partial Observations. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*. 1868–1870.
- [6] Gunter Blache, Marek Karpiński, and Juergen Wirtgen. 1997. *On Approximation Intractability of the Bandwidth Problem*. Inst. für Informatik.
- [7] Robert Bredereck, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2022. Envy-Free Allocations Respecting Social Networks. *Artificial Intelligence* (2022).
- [8] Thang Nguyen Bui and Curt Jones. 1992. Finding Good Approximate Vertex and Edge Partitions is NP-Hard. *Inform. Process. Lett.* 42, 3 (1992), 153–159.
- [9] F.R.K. Chung. 1984. On Optimal Linear Arrangements of Trees. *Computers & Mathematics with Applications* 10, 1 (1984), 43–60.
- [10] John A Doucette, Alan Tsang, Hadi Hosseini, Kate Larson, and Robin Cohen. 2019. Inferring True Voting Outcomes in Homophilic Social Networks. *Autonomous Agents and Multi-Agent Systems* 33, 3 (2019), 298–329.
- [11] Eduard Eiben, Robert Ganian, Thekla Hamm, and Sebastian Ordyniak. 2020. Parameterized Complexity of Envy-Free Resource Allocation in Social Networks. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence*. 7135–7142.
- [12] S. Even and Y. Shiloach. 1978. NP-Completeness of Several Arrangements Problems. *Technical Report, TR-43 The Technicon* (1978), 29. <https://www.cs.technion.ac.il/users/wwwb/cgi-bin/tr-info.cgi/1975/CS/CS0043>
- [13] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. 2019. Envy-Freeness in House Allocation Problems. *Mathematical Social Sciences* 101 (2019), 104–106.
- [14] M.R. Garey, D.S. Johnson, and L. Stockmeyer. 1976. Some Simplified NP-Complete Graph Problems. *Theoretical Computer Science* 1, 3 (1976), 237–267.
- [15] Laurent Gourvès, Julien Lesca, and Anaëlle Wilczynski. 2017. Object Allocation via Swaps along a Social Network. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 213–219.
- [16] Umberto Grandi. 2017. Social Choice and Social Networks. *Trends in Computational Social Choice* (2017), 169–184.
- [17] Hadi Hosseini, Justin Payan, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. 2023. Graphical House Allocation. *arXiv preprint arXiv:2301.01323* (2023).
- [18] Ayumi Igarashi and Edith Elkind. 2016. Hedonic Games with Graph-Restricted Communication. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. 242–250.
- [19] Takehiro Ito, Yuni Iwamasa, Naonori Kakimura, Naoyuki Kamiyama, Yusuke Kobayashi, Yuta Nozaki, Yoshio Okamoto, and Kenta Ozeki. 2022. Reforming an Envy-Free Matching. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence*. 5084–5091.
- [20] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. 2013. Bin Packing with Fixed Number of Bins Revisited. *J. Comput. System Sci.* 79, 1 (2013), 39–49.
- [21] Naoyuki Kamiyama. 2021. The Envy-Free Matching Problem with Pairwise Preferences. *Inform. Process. Lett.* 172 (2021), 106158.
- [22] Naoyuki Kamiyama, Pasin Manurangsi, and Warut Suksompong. 2021. On the Complexity of Fair House Allocation. *Operations Research Letters* 49, 4 (2021), 572–577.
- [23] Subhash Khot. 2006. Ruling Out PTAS for Graph Min-Bisection, Dense k-Subgraph, and Bipartite Clique. *SIAM J. Comput.* 36, 4 (2006), 1025–1071.
- [24] Pascal Lange and Jörg Rothe. 2019. Optimizing Social Welfare in Social Networks. In *Proceedings of the 6th International Conference on Algorithmic Decision Theory*. 81–96.
- [25] Sagar Massand and Sunil Simon. 2019. Graphical One-Sided Markets. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*. 492–498.
- [26] Christos Papadimitriou. 1976. The NP-Completeness of the Bandwidth Minimization Problem. *Computing* 16, 3 (1976), 263–270.
- [27] Dominik Peters. 2016. Graphical Hedonic Games of Bounded Treewidth. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. 586–593.
- [28] Alvin E Roth, Tayfun Sönmez, and M Utku Ünver. 2004. Kidney Exchange. *The Quarterly Journal of Economics* 119, 2 (2004), 457–488.
- [29] M. A. Seidvasser. 1970. The Optimal Number of the Vertices of a Tree. *Diskref. Anal.* 19 (1970), 56–74.
- [30] Lloyd Shapley and Herbert Scarf. 1974. On Cores and Indivisibility. *Journal of Mathematical Economics* 1, 1 (1974), 23–37.
- [31] Lars-Gunnar Svensson. 1999. Strategy-Proof Allocation of Indivisible Goods. *Social Choice and Welfare* 16, 4 (1999), 557–567.
- [32] Alan Tsang, John A Doucette, and Hadi Hosseini. 2015. Voting with Social Influence: Using Arguments to Uncover Ground Truth. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*. 1841–1842.