4.1 Trees and Planar Graphs

Recall that we presented two approximation guarantees for trees, $O(n)$ (Proposition 3.1) and $O(\Delta \log n)$ (Corollary 3.5). Both of these results are $\widetilde{O}(n)$ in the worst case.

Theorem 4.2. *For any constant $\epsilon > 0$, there is no efficient $O(n^{1-\epsilon})$ approximation algorithm for Graphical House Allocation on depth-2 trees unless P = NP.*

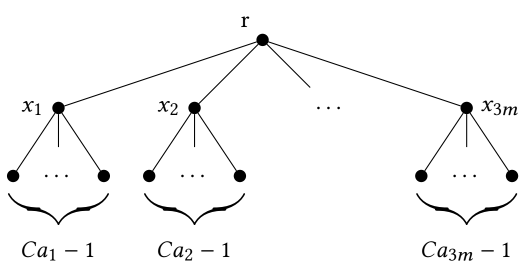

Figure 2: Mapping a Unary 3-Partition instance to a tree.

Proof Sketch. Given a Unary 3-Partition instance, we construct a graph according to Figure 2 where C is some positive integer we will decide later. The multiset of house values consists of CT houses with value j for each $j \in [m]$, and one house with value 0.

If there is a valid 3-partition, we can construct an allocation with envy at most $3m^2$. If there is no 3-partition, any allocation must have envy at least C. We can now set C appropriately. □

4.2 General and Bounded-Degree Graphs

In this section, we generalize the arguments from Section 4.1 to other classes of graphs. The main technique is similar to that of Theorem 4.2, so we just present ideas for the graph construction in each of these proofs, with the details in Hosseini et al. [19].

We first match the $O(n^2)$ upper bound for connected graphs (Proposition 3.1 and Corollary 3.5).

Theorem 4.3. *For any constant $\epsilon > 0$, there is no efficient $O(n^{2-\epsilon})$ approximation algorithm for Graphical House Allocation on connected graphs unless P = NP.*

Proof Sketch. We replace the Ca_i-sized stars in Figure 2 with Ca_i-sized cliques. The rest of the proof is similar to Theorem 4.2. □

So far in our two lower bounds (Theorems 4.2 and 4.3), we were able to use simple counting techniques, because counting edges with non-zero envy in stars and cliques is straightforward. Our next results will require much more careful analysis.

We will start with bounded-degree planar graphs. Our reduction uses grid graphs instead of stars and cliques, and so we will need a technical lemma to help us with estimating the number of edges with nonzero envy.

Lemma 4.4. *Let $G = Grid(r, c)$ be a grid graph with r rows and c columns such that $r \le c$. Let $A \subseteq V$ be any set of nodes in this graph such that $|A| \le rc/2$. Then, $\delta_G(A) \ge \min\{\sqrt{|A|}, r/2\}$.*

Armed with Lemma 4.4, we can now present our lower bound on bounded-degree planar graphs.

Theorem 4.5. *For any constant $\epsilon > 0$, no efficient $O(n^{0.5-\epsilon})$ approximation algorithm exists for Graphical House Allocation on bounded-degree planar graphs unless P = NP.*

Proof Sketch. We replace the stars of size Ca_i in Figure 2 with grid graphs containing C rows and Ca_i columns. The rest of the proof flows similarly to Theorem 4.2. Lemma 4.4 helps in estimating the envy blow-up if there is no 3-partition. □

Note that Theorem 4.5 matches the $O(\sqrt{n})$ upper bound from Corollary 3.5.

Our next lower bound applies to arbitrary graphs with bounded degree, and matches the $O(n)$ upper bound from Proposition 3.1 and Corollary 3.5. In this reduction, we use a recent polynomial-time algorithm [8] to compute bipartite Ramanujan multigraphs for any even number m of vertices, and any degree $d \ge 3$. At a high level, we replace the star gadgets from the proof of Theorem 4.2 with these Ramanujan graphs and use the expansion properties of Ramanujan graphs to prove a lemma similar to (and stronger than) Lemma 4.4.

Theorem 4.6. *For any constant $\epsilon > 0$, there is no efficient $O(n^{1-\epsilon})$ approximation algorithm for the Graphical House Allocation problem on bounded-degree graphs unless P = NP.*

4.3 Bounded-Degree Trees

Our final lower bound shows that Graphical House Allocation is NP-hard even when the underlying graph is a bounded degree tree. We still use Unary 3-Partition in our reduction but this proof is significantly different from the previous ones. Our reduction will use a gadget we call the *flower.*[4]

Definition 4.7. *The flower $F(n, k)$ is a rooted tree with n nodes and maximum degree $k + 1$, defined recursively as follows: for any $k \ge 1$, $F(1, k)$ is simply an isolated vertex which is the root node. For $n > 1$, $F(n, k)$ consists of a root node connected to the root nodes of d other flowers $F(n_1, k), \ldots, F(n_d, k)$ such that*

(a) $\sum_{i=1}^{d} n_i = n - 1$,
(b) if $n - 1 \ge k$, then $d = k$ if n and k have different parities, and $d = k - 1$ otherwise,
(c) each n_i is odd,
(d) for any $i, j \in [d]$, $|n_i - n_j| \le 2$.

To ensure consistency with floral terminology, we refer to the root node of the flower $F(n, k)$ as its pistil *and the (recursively smaller) flowers $F(n_1, k), \ldots, F(n_d, k)$ as its* petals.

Before we use flowers, we show that they are well-defined and efficiently constructible.

Lemma 4.8. *For any $n \ge 1$ and $k \ge 3$, the flower $F(n, k)$ exists and can be constructed in poly(n, k) time.*

The reason we build flowers is because they satisfy the two following useful properties.

Lemma 4.9. *Let $F(n, k)$ be a flower on the set of vertices N, and suppose $n \ge 10k$, and n and k have different parities. Then, $F(n, k)$ satisfies the following properties:*

[4] To the best of our knowledge, our specific flower graph is novel but it is possible (likely even) that the term "flower" has appeared before in the graph theory literature.

(i) For any $A \subseteq N$ such that $|A|$ is even and A does not contain the pistil, $\delta(A) \geq 2$.
(ii) Each petal of $F(n,k)$ has size in the interval $\left[\frac{4n}{5k}, \frac{6n}{5k}\right]$.

These simple properties are all we need to show the hardness of Graphical House Allocation on bounded-degree trees.

Theorem 4.10. *Graphical House Allocation is NP-hard on bounded-degree trees.*

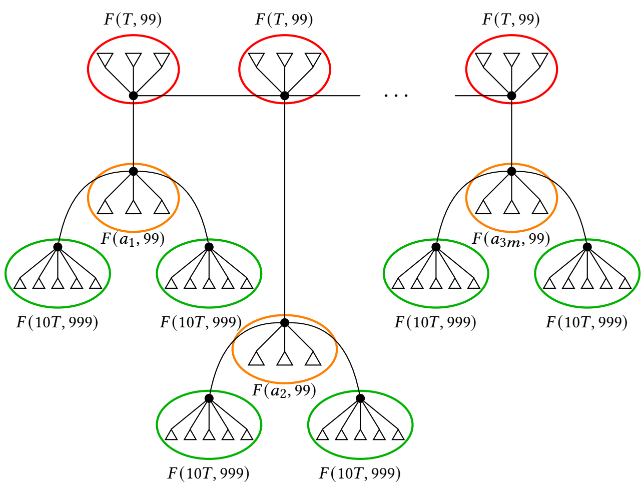

Figure 3: Mapping a Unary 3-Partition instance to a bounded degree tree. Here, the orange, red, and green circles correspond to small, medium, and large flowers respectively.

Proof Sketch. Given a Unary 3-Partition instance, we construct a graph according to Figure 3; the shaded circles correspond to pistils and the white triangular blocks correspond to petals. The house values are defined as follows: we have $4m+1$ unique values such that the gaps between these values are exponentially decreasing. That is, the gap between the least and the second least value is significantly larger than the gap between the second least and the third least value and so on. For each unique value, there are T houses with that value in the multiset H, with the exception of the largest value which has enough houses (with that value) to ensure the total size of the multiset H is equal to the number of nodes in the graph.

We can show that in any optimal allocation, the first $3m$ clusters must be allocated to flowers of the form $F(T,99)$. The next m clusters must be allocated in a way that creates a 3-partition to minimize envy. That is, each of these values must be allocated to three flowers of the form $F(a_i,99)$ such that the total size of these three flowers sums up to T. If it is not possible to do this, the minimum envy of the allocation is strictly higher. This allows us to separate instances with a valid 3-Partition. □

5 THE CURIOUS CASE OF COMPLETE BINARY TREES

In this section, we investigate Graphical House Allocation on instances where the underlying graph is a complete binary tree B_k. Recall that such a tree has depth k, and $2^{k+1}-1$ vertices in total, of which 2^k are leaves. All leaves, furthermore, are at the same depth.

In Hosseini et al. [20, Theorem 4.11], it was shown that for any binary tree (complete or otherwise), at least one optimal allocation satisfies the *local median property*: the value at every internal node is the median among the values given to that node and its two children. The same authors surmised that, for any binary tree, at least one optimal allocation satisfies the stronger *global median property*: for every internal node v, either its left subtree gets strictly lower-valued houses and its right subtree gets strictly higher-valued houses, or the other way round. Note that if true, this would lead to a straightforward recursive polynomial-time algorithm that would compute an optimal allocation on (nearly) balanced binary trees.

We now give a refutation of this conjecture. We illustrate an instance on a complete binary tree of depth 3, in which no optimal allocation satisfies the global median property. This is a quite surprising result that shows that the general problem on complete binary trees may be much harder than expected.

Example 5.1. *Consider the instance (B_3, H), where*

$$H = \{0,0,0,0,0,0,0,1,1,1,2,3,3,3,3\}.$$

See Figure 4. The top shows the only allocation satisfying the global median property (up to re-ordering). The total non-negligible envy incurred by this assignment comes out of the thick red edges of the B_3, which incur a total envy of 6. However, the bottom shows an allocation with an envy of 5 (incurred by the thick red edges), showing that the global median is strictly sub-optimal.

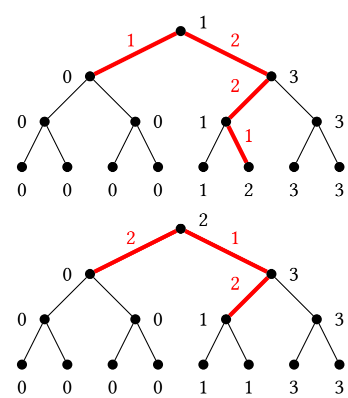

Figure 4: Refutation of the global median property on complete binary trees.

Fix an arbitrary instance of Graphical House Allocation on the complete binary tree B_k on $n = 2^{k+1}-1$ vertices, and consider the valuation interval. There are n values on the interval. Of particular interest to us is the size of the *smallest* $(i, n-i)$-cut, i.e., $\delta_{B_k}(i)$. Since $\delta_{B_k}(i) = \delta_{B_k}(n-i)$, we can WLOG take $i \leq \lceil n/2 \rceil$. We now need a definition.

Definition 5.2 (Repunit Representation and Elegance)**.** *For any $m \geq 1$, let a* repunit representation *of m be any finite sequence*

m	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
elegance(m)	1	2	1	2	3	2	1	2	3	2	3	2	3	2	1	2	3	2	3	4

Table 2: List of elegance(m) **for** $1 \leq m \leq 20$**.**

$(a_1, \ldots, a_r) \in \mathbb{Z}^r$ *satisfying*

$$m = \sum_{i=1}^{r} \text{sgn}(a_i) \cdot (2^{|a_i|} - 1)$$

where $\text{sgn}(a_i)$ *is* 1 *(resp.* -1*) if* $a_i \geq 0$ *(resp.* $a_i < 0$*). Note that every* $m \geq 1$ *has a repunit representation (e.g., the length-m sequence of all ones). We define* elegance(m) *as the smallest r for which m has a repunit representation* $(a_1, \ldots, a_r)$ *of length r.*

The intuition behind this definition is to capture the most "efficient" way to write m as the sum or difference of binary repunits, i.e., numbers of the form $11\ldots1$. For instance, elegance$(10) = 2$, because $10 = (2^3 - 1) + (2^2 - 1)$, and there is no shorter repunit representation. Similarly, elegance$(12) = 2$, as $12 = (2^4 - 1) - (2^2 - 1)$. Note that 12 cannot be written as the *sum* of two repunits. Table 2 summarizes the elegance of all numbers up to 20.

The following proposition relates elegance to the size of the smallest $(i, n - i)$-cut in a complete binary tree, namely $\delta_{B_k}(i)$.

Proposition 5.3. *Let* B_k *be the complete binary tree on* $n = 2^{k+1} - 1$ *vertices. Then for* $i \leq 2^k - 1$, elegance$(i) - 1 \leq \delta_{B_k}(i) \leq$ elegance(i).

Proof Sketch. For each edge going across a cut in B_k, one of its endpoints is the root of a binary subtree, and it contributes a term in a repunit representation (possibly along with an extra additive term). Conversely, any repunit representation gives rise to a cut. Therefore, cuts correspond to repunit representations up to a single additive term. Minimizing both sides yields the result. □

We note that if $i \ll n$, then in fact $\delta_{B_k}(i) =$ elegance(i). Therefore, elegance(i) actually characterizes the size of the minimum $(i, n - i)$ cut in any sufficiently large binary tree.

Consider a value-agnostic algorithm for complete binary trees. Such an algorithm would need to assign the house values in any instance in some fixed order $(v_1, \ldots, v_n)$ to the vertices of B_k. The following proposition shows that doing this cannot simultaneously achieve the optimal cut on all smallest subintervals, and this leads to a lower bound on the approximability.

Proposition 5.4. *There is no value-agnostic algorithm for complete binary trees that attains an approximation better than* $(5/3) \approx 1.67$.

The counterexample in Proposition 5.4 and the failure of the global median property (Example 5.1) may seem to suggest that, even for complete binary trees, *any* constant approximation ratio is unattainable. Remarkably, the following result shows that this is not the case: there is a *value-agnostic* algorithm attaining a constant approximation on any complete binary tree. Indeed, ordering the vertices of B_k in the standard in-order traversal and allocating the (sorted) values in that order yields a 3.5-approximation.

Theorem 5.5. *Let* B_k *be the complete binary tree on* $n = 2^{k+1} - 1$ *vertices. Then, on any house allocation instance on* B_k*, assigning the houses in increasing order to the vertices of* B_k *in the standard in-order traversal gives us a total envy at most* 3.5 *times the optimal value.*

It is instructive to check why this technique does not hold for arbitrary binary trees. Proposition 5.3 does not hold in general for non-complete binary trees. A complete binary tree ensures that there is always a binary subtree of the size given by a repunit representation to include on one side of the cut, but we lose this guarantee for non-complete trees.

We leave it as an open problem to construct either value-agnostic deterministic algorithms that achieve an approximation ratio better than 3.5, or to obtain any polynomial-time algorithm (which cannot be value-agnostic) to obtain any approximation ratio better than 1.67 for complete binary trees. We believe there should be an exact algorithm for this very special class of graphs, and hope that this will instigate future research into this problem.

6 CONCLUSIONS

We explored the approximability of Graphical House Allocation, presenting tight approximation algorithms for several classes of connected graphs, to our knowledge the first such results in the area. In particular, we gave polynomial-time algorithms exploiting graph structures to approximate the optimal envy on general graphs, trees, planar graphs, bounded-degree graphs, bounded-degree planar graphs, and bounded-degree trees; for each of these classes, we also gave a matching lower bound. Our algorithms were value-agnostic, i.e., they took into account only the input graph and the ordering among the house values but not the values themselves. We showed that any allocation on a random graph is a $(1 + o(1))$-approximation, and also gave a value-agnostic algorithm to show a 3.5-approximation on all instances on complete binary trees.

The main question we leave for future work is the complexity of Graphical House Allocation on complete binary trees. We know by the results in Section 5 that no exact algorithm can be value agnostic, but there seems to be no obvious way of leveraging the values, on even such a structured class of graphs.

Conjecture 6.1. *Graphical House Allocation is polynomial-time solvable on complete binary trees.*

ACKNOWLEDGMENTS

We wish to thank Justin Payan for many helpful discussions during the early stages of this work. We also wish to thank Paul Seymour for some discussions related to the material in Section 5. Rohit Vaish acknowledges support from SERB grant no. CRG/2022/002621 and DST INSPIRE grant no. DST/INSPIRE/04/2020/000107. Andrew McGregor and Rik Sengupta acknowledge support from NSF grant CCF-1934846. Hadi Hosseini acknowledges support from NSF IIS grants #2144413 and #2107173.

REFERENCES

[1] 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.

[2] Sanjeev Arora, Alan Frieze, and Haim Kaplan. 1996. A New Rounding Procedure for the Assignment Problem with Applications to Dense Graph Arrangement Problems. In *Proceedings of 37th Conference on Foundations of Computer Science*. 21–30.

[3] 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.

[4] Arpita Biswas, Justin Payan, Rik Sengupta, and Vignesh Viswanathan. 2023. *The Theory of Fair Allocation Under Structured Set Constraints*. 115–129. https://doi.org/10.1007/978-981-99-7184-8_7

[5] Robert Bredereck, Andrzej Kaczmarczyk, and Rolf Niedermeier. 2022. Envy-Free Allocations Respecting Social Networks. *Artificial Intelligence* 305 (2022), 103664.

[6] 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.

[7] F.R.K. Chung. 1984. On Optimal Linear Arrangements of Trees. *Computers & Mathematics with Applications* 10, 1 (1984), 43–60.

[8] Michael B. Cohen. 2016. Ramanujan Graphs in Polynomial Time. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*. 276–281.

[9] Hristo N Djidjev and Imrich Vrt'o. 2006. Crossing Numbers and Cutwidths. *Journal of Graph Algorithms and Applications. v7* (2006), 245–251.

[10] 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.

[11] Guy Even, Joseph Seffi Naor, Satish Rao, and Baruch Schieber. 2000. Divide-and-Conquer Approximation Algorithms via Spreading Metrics. *J. ACM* 47, 4 (2000), 585–616.

[12] S. Even and Y. Shiloach. 1978. NP-Completeness of Several Arrangements Problems. *Technical Report, TR-43 The Technicon* (1978), 29.

[13] Uriel Feige and James R Lee. 2007. An Improved Approximation Ratio for the Minimum Linear Arrangement Problem. *Inform. Process. Lett.* 101, 1 (2007), 26–29.

[14] Andreas Emil Feldmann. 2012. Fast Balanced Partitioning Is Hard Even on Grids and Trees. In *Proceedings of the 37th International Symposium on Mathematical Foundations of Computer Science (MFCS)*. 372–382.

[15] Andreas Emil Feldmann and Luca Foschini. 2015. Balanced Partitions of Trees and Applications. *Algorithmica* 71, 2 (feb 2015), 354–376.

[16] Jiarui Gan, Warut Suksompong, and Alexandros A. Voudouris. 2019. Envy-Freeness in House Allocation Problems. *Mathematical Social Sciences* 101 (2019), 104–106.

[17] M.R. Garey, D.S. Johnson, and L. Stockmeyer. 1976. Some Simplified NP-Complete Graph Problems. *Theoretical Computer Science* 1, 3 (1976), 237–267.

[18] M. R. Garey and D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness* (first edition ed.). W. H. Freeman.

[19] Hadi Hosseini, Andrew McGregor, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. 2023. Tight Approximations for Graphical House Allocation. arXiv:2307.12482 [cs.DS]

[20] Hadi Hosseini, Justin Payan, Rik Sengupta, Rohit Vaish, and Vignesh Viswanathan. 2023. Graphical House Allocation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 161–169.

[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] Ephraim Korach and Nir Solel. 1993. Tree-Width, Path-Width, and Cutwidth. *Discrete Applied Mathematics* 43, 1 (1993), 97–101.

[24] Tom Leighton and Satish Rao. 1999. Multicommodity Max-Flow Min-Cut Theorems and Their Use in Designing Approximation Algorithms. *J. ACM* 46, 6 (1999), 787–832.

[25] Jayakrishnan Madathil, Neeldhara Misra, and Aditi Sethia. 2023. The Complexity of Minimizing Envy in House Allocation. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2673–2675.

[26] Burkhard Monien and Ivan Hal Sudborough. 1988. Min Cut is NP-Complete for Edge Weighted Trees. *Theoretical Computer Science* 58, 1-3 (1988), 209–229.

[27] Satish Rao and Andréa W Richa. 2005. New Approximation Techniques for Some Linear Ordering Problems. *SIAM J. Comput.* 34, 2 (2005), 388–404.

[28] M. A. Seidvasser. 1970. The Optimal Number of Vertices of a Tree. *Diskref. Anal.* 19 (1970), 56–74.

[29] Lloyd Shapley and Herbert Scarf. 1974. On Cores and Indivisibility. *Journal of Mathematical Economics* 1, 1 (1974), 23–37.

[30] Lars-Gunnar Svensson. 1999. Strategy-Proof Allocation of Indivisible Goods. *Social Choice and Welfare* 16, 4 (1999), 557–567.

[31] Mihalis Yannakakis. 1985. A Polynomial Algorithm for the Min-Cut Linear Arrangement of Trees. *J. ACM* 32, 4 (1985), 950–988.