

Graphical Cake Cutting via Maximin Share

Edith Elkind¹, Erel Segal-Halevi² and Warut Suksompong³

¹Department of Computer Science, University of Oxford

²Department of Computer Science, Ariel University

³School of Computing, National University of Singapore

elkind@cs.ox.ac.uk, erelsgl@gmail.com, warut@comp.nus.edu.sg

Abstract

We study the recently introduced cake-cutting setting in which the cake is represented by an undirected graph. This generalizes the canonical interval cake and allows for modeling the division of road networks. We show that when the graph is a forest, an allocation satisfying the well-known criterion of maximin share fairness always exists. Our result holds even when separation constraints are imposed; however, in the latter case no multiplicative approximation of proportionality can be guaranteed. Furthermore, while maximin share fairness is not always achievable for general graphs, we prove that ordinal relaxations can be attained.

1 Introduction

Cake cutting is an old and famous problem in resource allocation, with the cake serving as a metaphor for a heterogeneous divisible resource that is supposed to be fairly divided among interested agents. While the problem has long enjoyed substantial attention from mathematicians and economists, it has also attracted ongoing interest from computer scientists, not least those working in artificial intelligence [Balkanski *et al.*, 2014; Li *et al.*, 2015; Brânzei *et al.*, 2016; Alijani *et al.*, 2017; Menon and Larson, 2017; Bei *et al.*, 2018; Arunachaleswaran *et al.*, 2019; Goldberg *et al.*, 2020; Hosseini *et al.*, 2020]. Indeed, as Procaccia [2013] aptly put it, cake cutting is more than just child’s play.

The cake in cake cutting is typically assumed to be a one-dimensional interval. Even though the linear representation is appropriate for modeling the division of time (for instance, usage of a jointly-owned facility), or space in a hallway, it is too simplistic to capture more complex resources such as road networks. This consideration has led Bei and Suksompong [2021] to introduce a more general *graphical cake* model, in which the resource comes in the form of a connected undirected graph. In parallel, Segal-Halevi [2021] addressed the case of graphs given by a disjoint union of intervals. In contrast to the single interval cake, for general graphs it is not always possible to find a connected *proportional* allocation, that is, an allocation that gives every agent a connected subset of the cake worth at least $1/n$ of the agent’s value for the entire cake, where n denotes the number of agents among

whom the cake is divided. Nevertheless, Bei and Suksompong showed that more than half of this guarantee can be recovered: for any connected graph, it is possible to ensure that every agent obtains at least $1/(2n - 1)$ of her total value, and this factor is tight in general (but can be improved for certain graphs). For the union-of-intervals case, Segal-Halevi established an approximation factor of $1/(m + n - 1)$, where m denotes the number of intervals.

In this paper, we study graphical cake cutting with respect to another prominent fairness notion, *maximin share fairness* [Budish, 2011]. An allocation satisfies this notion if it assigns to every agent a bundle worth at least her maximin share, i.e., the largest value that she can get if she is allowed to partition the cake into n connected parts and take the least valuable part. Maximin share fairness is a robust notion, which can naturally take into account the features and constraints arising in various settings. Indeed, this robustness enables us to derive positive results in three settings where approximate proportionality fails. First, we allow the graph to be arbitrary—it may be disconnected, and each of its connected components may have an arbitrary topology. Second, we allow agents to have general monotone valuations—unlike most of the cake cutting literature, we do not assume that valuations must be additive. With disconnected graphs or non-additive valuations, providing an approximate proportionality guarantee solely in terms of n is impossible. Third, we also consider a recently introduced setting of Elkind *et al.* [2021c] in which the shares of different agents should be sufficiently separated from one another; this allows us to model space between segments of roads with different owners, for example transition or buffer zones. In the presence of separation constraints, obtaining any multiplicative approximation of proportionality is again infeasible, as the value of all agents may be entirely concentrated in the same tiny portion. This observation motivated Elkind *et al.* to use the maximin share benchmark when separation is imposed. They showed that with separation, a maximin allocation always exists for a path graph, but not for a cycle graph.

1.1 Our Contributions

We begin in Section 3 by addressing the basic case where no separation constraints are imposed. In this case, for proportionality approximations it can be crucial whether pieces of different agents are allowed to share a finite number of points

[Bei and Suksompong, 2021]. We show that this modeling choice is also crucial with respect to maximin share fairness. In particular, if points can be shared, then a maximin allocation may not exist even when the graph is a star, whereas if sharing is not allowed, the existence of such an allocation can be guaranteed if the graph is acyclic (i.e., a disjoint union of trees, also known as a forest). Our results complement those of Bei and Suksompong, who observed that approximate proportionality cannot be ensured even for trees under the no-sharing assumption. Moreover, our guarantees degrade gracefully for graphical cakes with cycles: we attain a 1-out-of- $(n+r)$ maximin allocation, where r is the *feedback vertex set number* of the cake (that is, the smallest number of vertices whose removal would make the cake acyclic).

In Section 4, we consider the more general case where the pieces of any two agents must be separated by distance at least a given (positive) parameter. Our main technical result shows that a maximin allocation exists whenever the graph is acyclic—this significantly generalizes the existence result of Elkind *et al.* [2021c] for paths and complements their non-existence result for cycles. As with paths, our proof uses the following high-level idea: Given the maximin partitions of the agents, we find a part in one agent’s partition such that allocating the part to that agent rules out at most one part in each remaining agent’s partition; this allows us to recurse on the remaining agents and cake. While in the case of paths the desired part can be found by simply scanning the path from left to right, in an arbitrary forest there is no ‘left’ or ‘right’, so new techniques are needed. We develop auxiliary lemmas related to *real trees*—metric spaces defined by tree graphs—which may be of independent interest. As in the case of no separation, we obtain a 1-out-of- $(n+r)$ maximin allocation for general graphs.

For the case of positive separation, we show that, in general, the factor $n+r$ cannot be improved: for every $r \geq 0$, if n is sufficiently large, there exists a graph with feedback vertex set number r and a set of n agents with no 1-out-of- $(n+r-1)$ maximin allocation. However, better guarantees can be attained for smaller n and specific classes of graphs.

1.2 Additional Related Work

As mentioned earlier, cake cutting is a popular topic among researchers of several disciplines—see the classic books of Brams and Taylor [1996] and Robertson and Webb [1998], as well as a more recent survey by Procaccia [2016] offering a computer scientist’s perspective.

For a discussion of the work relevant to graphical fair division and separation constraints, see the related work sections in the papers of Bei and Suksompong [2021] and Elkind *et al.* [2021c]; we highlight here some important aspects of our study. First, we assume that each agent must receive a connected piece of cake. This assumption is often made in order to ensure that agents do not end up with a collection of crumbs—indeed, a bundle made up of tiny stretches of road in different parts of the network is unlikely to be of much use. Second, the connectivity requirement is imposed not only on the allocation, but also in the definition of the maximin share. This is consistent with previous work on maximin share fairness in constrained settings [Bouveret *et al.*, 2017;

Biswas and Barman, 2018; Lonc and Truszczynski, 2020; Bei *et al.*, 2021; Elkind *et al.*, 2021c]. Bouveret *et al.* [2017] proved that for *indivisible* items lying on a tree, a maximin allocation exists. However, as we discuss in Section 5, the ‘‘last diminisher’’ approach that they used for this proof does not work in our setting with separation. Recently, Igarashi and Zwicker [2021] studied envy-freeness in graphical cake cutting under the assumption that agents cannot share individual points, while Elkind *et al.* [2021b] investigated land division with separation constraints.

In the papers above, as in our paper, the resource to be divided forms a graph. A complementary line of work studies fair division scenarios in which the *agents’* relationship is captured by a graph [Abebe *et al.*, 2017; Bei *et al.*, 2017; Aziz *et al.*, 2018].

2 Preliminaries

There is a set of agents $\mathcal{N} = [n]$, where $[k] := \{1, 2, \dots, k\}$ for any positive integer k . The cake is represented by a finite undirected graph $G = (V, E)$, which may be connected or not. Each agent has a nonnegative, monotone, and continuous valuation function v_i , which is not necessarily additive. In particular, continuity implies that the vertices in V have zero value. This model captures the classic cake cutting setting as well as the circular cake (a.k.a. pie) studied by Elkind *et al.* [2021c]: an interval cake corresponds to taking G to be a path graph, while a pie is equivalent to a cycle graph. A *piece of cake* is a finite union of intervals from one of more edges in E . The piece is said to be *connected* if for any points x, y in it, one can get from x to y along the graph G while staying within this piece of cake. We assume that each agent must receive a connected piece of cake.

There is a separation parameter $s \geq 0$. When $s > 0$, the edge lengths play an important role. We measure distance along the edges of G . For any two points $x, y \in G$, we denote by $\text{DIST}^G(x, y)$ the length of a shortest path from x to y along the edges of G ; if x and y belong to different connected components of G , we set $\text{DIST}^G(x, y) = \infty$. For two pieces of cake $X, Y \subseteq G$, we denote by $\text{DIST}^G(X, Y)$ the shortest distance between a point in X and a point in Y along the edges of G , i.e., $\text{DIST}^G(X, Y) = \inf_{x \in X, y \in Y} \text{DIST}^G(x, y)$; if Y consists of a single point y , we simply write $\text{DIST}^G(X, y)$.

A *partition* of the cake is a set $\mathbf{P} = \{P_1, \dots, P_n\}$, where each P_i is a connected piece of cake, and the pieces are pairwise disjoint: $P_i \cap P_j = \emptyset$ for all $i \neq j$. When $s = 0$, we will consider, in addition to the disjoint-pieces setting, an alternative setting in which $P_i \cap P_j$ may contain finitely many points. An *allocation* is defined similarly, except that we have a vector $\mathbf{A} = (A_1, \dots, A_n)$ instead of a set, where piece A_i is allocated to agent i . A partition \mathbf{P} is said to be *s-separated* if $\text{DIST}^G(P_i, P_j) \geq s$ for all $i \neq j$; an analogous definition holds for an allocation. We assume that partitions and allocations are required to be *s-separated*. Observe that for $s > 0$, in any *s-separated* partition or allocation, some of the cake necessarily remains unallocated. Moreover, since any two pieces are separated by a positive distance, we assume without loss of generality in this case that the pieces contain

closed intervals only. Denote by $\Gamma_{n,s}$ the set consisting of all s -separated partitions.

The main fairness notion of our paper is the following:

Definition 2.1. The *maximin share* of agent i , denoted by $\text{MMS}_i^{n,s}$, is defined as $\sup_{P \in \Gamma_{n,s}} \min_{j \in [n]} v_i(P_j)$.

We omit the superscript s when it is clear from the context. Similarly to the interval cake [Elkind *et al.*, 2021c], the supremum in Definition 2.1 can be replaced with a maximum. In other words, there is always a maximizing partition, which we refer to as a *maximin partition*. An allocation in which every agent receives at least her maximin share is said to be a *maximin allocation*.

All omitted proofs can be found in the full version of our paper [Elkind *et al.*, 2021a].

3 No Separation

In this section, we address the basic case where there is no separation constraint imposed on the allocation, i.e., $s = 0$. When $s > 0$, the pieces of any two agents cannot be adjacent to each other, so we can assume without loss of generality that all pieces consist of closed intervals only and the pieces have empty intersections. For $s = 0$, however, this is not true: there are essential differences between the empty-intersection setting and the *finite-intersection* setting, in which pieces may overlap in finitely many points. This observation was made by Bei and Suksompong [2021] with respect to approximate proportionality for the case of a star graph. Specifically, in the empty-intersection setting, $n - 1$ agents do not receive the center of the star and therefore can receive cake from at most one edge; as a consequence, one obtains strong negative results. On the other hand, in the finite-intersection setting, non-trivial welfare guarantees can be obtained.

As we will see, the distinction between empty intersection and finite intersection is crucial with respect to maximin share fairness too. Note that we use the same constraints when selecting an allocation and when computing the benchmark: e.g., in the finite-intersection setting, when computing the agents' maximin shares, we optimize over all partitions where parts may intersect in a finite number of points. Thus, allowing intersections results in a more demanding benchmark.

First, we show that in the finite-intersection setting, a maximin allocation may not exist.

Proposition 3.1. *Assume that the allocated pieces are allowed to intersect in a finite number of points. There exists an instance with $n = 3$ agents and a star cake in which no maximin allocation exists.*

Proof. The proof follows the celebrated Theorem 2.1 of Kurokawa *et al.* [2018], which shows that a maximin allocation of indivisible objects may not exist for $n = 3$ agents.

In their instance, there are 12 objects, indexed by $j \in [3]$ and $k \in [4]$. Each agent $i \in [3]$ values each object (j, k) by:

$$v_i(j, k) = 1000000 + 1000 \cdot T_{j,k} + E_{j,k}^{(i)},$$

where $T, E^{(1)}, E^{(2)}, E^{(3)}$ are carefully chosen 3×4 matrices with all values smaller than 100. Kurokawa *et al.* proved that every agent can partition the objects into 3 subsets of 4 objects

each, in such a way that the sum of values in each subset is exactly 4055000; this value is therefore the maximin share of all agents. These authors then showed that no allocation gives every agent at least this value.

In our instance, there is a star graph with 12 edges connected to a single center vertex c . The edges are indexed by pairs (j, k) with $j \in [3], k \in [4]$. Agent $i \in [3]$ has value $v_i(j, k)$ for the edge (j, k) , and this value is spread uniformly across the edge. Since the pieces may intersect in a finite number of points, the maximin share of each agent in our instance is also 4055000, and corresponds to partitioning the set of edges into 3 subsets of 4 edges each, intersecting in c .

If an agent's piece is contained in a single edge, then her value is clearly less than 2000000. Hence, in a maximin allocation, no edge is shared among two or more agents. We may therefore assume without loss of generality that each agent receives a piece containing two or more whole edges. But the same argument as that of Kurokawa *et al.* [2018] shows that no such allocation can be a maximin allocation. \square

Next, we show that in the empty-intersection setting, a maximin allocation always exists when the cake is a forest.

Theorem 3.2. *Let G be a forest and $s = 0$. Assume that all allocated pieces must be completely disjoint. For agents with arbitrary monotone valuations, a maximin allocation exists.*

Intuitively, given the n maximin partitions of the agents, we want to choose a part in one agent's partition that overlaps at most one part in each remaining agent's partition—this will allow us to recurse on the remaining agents and their leftover partitions. To this end, we introduce the following definition. Given a graph G and a family $\mathbf{X} = (X_1, \dots, X_k)$ of connected pieces of G , we say that a piece X_j is *0-good* if for all $j_1, j_2 \in [k]$ the following holds: If $X_{j_1} \cap X_{j_2} \neq \emptyset$ and $X_{j_2} \cap X_{j_1} \neq \emptyset$, then $X_{j_1} \cap X_{j_2} \neq \emptyset$.

Lemma 3.3. *Let G be a tree and let $\mathbf{X} = (X_1, \dots, X_k)$ be a family of connected subsets of G , for some $k \in \mathbb{N}$. For some $j^* \in [k]$, the piece X_{j^*} is 0-good.*

In order to prove this lemma, we must handle both open and closed pieces.¹ Indeed, for $n = 3$ and a star graph with three edges of equal value, a maximin partition contains two open pieces and one closed piece.²

Proof of Lemma 3.3. Let m be the number of edges in G . The proof is by induction on m .

For the base case $m = 1$, G is an interval; assume without loss of generality that it is the interval $[0, 1]$. Each X_j is also an interval which may be open, half-open, or closed. For each $j \in [k]$, let ℓ_j and r_j be the left and right endpoint of X_j , respectively. Choose $j^* \in [k]$ such that r_{j^*} is smallest; if there is a tie, break it in favor of a piece that is right-open, i.e., does not contain r_{j^*} . If X_{j^*} contains its right endpoint

¹We are grateful to Alex Ravsky for the proof idea.

²The finite-intersection analogue of Lemma 3.3 does not hold. For example, if G is a star graph with center c and edges e_1, e_2, e_3, e_4 , and $\mathbf{X} = (e_1 \cup c \cup e_2, e_3 \cup c \cup e_4, e_1 \cup c \cup e_3, e_2 \cup c \cup e_4)$, there is no X_j with the property that if $X_{j_1} \cap X_{j_2}$ is infinite and $X_{j_2} \cap X_{j_1}$ is infinite, then $X_{j_1} \cap X_{j_2}$ is infinite. This explains why Theorem 3.2 does not extend to the finite-intersection case.

(i.e., it is right-closed), then any piece X_j intersecting X_{j^*} must have $\ell_j \leq r_{j^*}$ and must contain the point r_{j^*} . On the other hand, if X_{j^*} does not contain its right endpoint (i.e., it is right-open), then any piece X_j intersecting X_{j^*} must have $\ell_j < r_{j^*}$ and must contain an interval $(r_{j^*} - \varepsilon, r_{j^*})$ for some sufficiently small $\varepsilon > 0$. In both cases, any two such pieces X_j intersect, so X_{j^*} is 0-good. This concludes the analysis for base case.

For the inductive step, let $m \geq 2$, assume that the statement holds for graphs with at most $m - 1$ edges, and consider a graph G with m edges. Since G is a tree, it contains an edge e with endpoints u and w such that w is a leaf of G . Let G^- be the graph G without the vertex w and the internal points of edge e ; note that u is a vertex of G^- . We consider two cases.

Case 1: At least one piece of \mathbf{X} is contained in e (so it is an interval). We can then view e as the interval $[0, 1]$ (with $w = 0, u = 1$), and use the same approach as for $m = 1$, i.e., pick the interval X_{j^*} with the smallest right endpoint, breaking ties in favor of right-open intervals. Then X_{j^*} is 0-good by the same argument as in the base case $m = 1$.

Case 2: No piece of \mathbf{X} is fully contained in e . This means that all pieces of \mathbf{X} intersect G^- . Let $\mathbf{X}' := (X'_1, \dots, X'_k)$, where $X'_j := X_j \cap G^-$ for each $j \in [k]$. By the inductive assumption applied to G^- , at least one piece in \mathbf{X}' , say X'_{j^*} , is 0-good with respect to \mathbf{X}' . It suffices to show that X_{j^*} is also 0-good with respect to \mathbf{X} .

We claim that if X_{j^*} intersects some other piece X_j , then X'_{j^*} also intersects X'_j . To see this, consider a point $z \in X_{j^*} \cap X_j$. If $z \in G^-$, then X'_{j^*} intersects X'_j and we are done. If $z \in e$, then both X_{j^*} and X_j intersect e . On the other hand, we assume that all pieces of \mathbf{X} intersect G^- . Thus, both X_{j^*} and X_j contain the point u , which is the unique point connecting e and G^- . Hence $u \in X'_{j^*} \cap X'_j$, so again X'_{j^*} intersects X'_j . This establishes the claim.

We now show that X_{j^*} is 0-good with respect to \mathbf{X} . Suppose that X_{j^*} intersects two other pieces in \mathbf{X} , say X_{j_1} and X_{j_2} . By the claim in the previous paragraph, X'_{j^*} intersects both X'_{j_1} and X'_{j_2} . Since X'_{j^*} is 0-good with respect to \mathbf{X}' , it must be that X'_{j_1} intersects X'_{j_2} . In particular, X_{j_1} intersects X_{j_2} . Hence, X_{j^*} is 0-good with respect to \mathbf{X} . \square

With Lemma 3.3 in hand, we can now show that, in the empty-intersection setting, a maximin allocation exists whenever the cake is a forest.

Proof of Theorem 3.2. For each agent, consider her maximin partition. Every part of the partition is contained in some tree of the forest. Let $T \subseteq G$ be a tree that contains at least one part from the maximin partition of at least one agent.

For every agent $i \in \mathcal{N}$, let k_i be the number of parts of i 's maximin partition that are contained in T , and denote the parts by $T_{i,1}, \dots, T_{i,k_i}$. By Lemma 3.3, there exists some $i \in \mathcal{N}$ and $j \in [k_i]$ such that $T_{i,j}$ is 0-good. Allocate the part $T_{i,j}$ to agent i , and divide the remaining cake recursively among the remaining agents.

The remaining cake is still a forest. By definition of a 0-good subset, for every other agent, at most one part of her maximin partition overlaps the allocated piece $T_{i,j}$. Hence, for each of the $n - 1$ remaining agents, at least $n - 1$ parts from

her maximin partition remain intact. Therefore the recursive call indeed returns a maximin allocation. \square

As we have seen, the seemingly minor distinction of whether individual points can be shared among allocated pieces makes a decisive difference in relation to maximin share fairness. Which assumption is more realistic depends on the use case, for example whether road intersections can only be owned by one agent or shared by multiple agents. Bei and Suksompong [2021] showed that nontrivial egalitarian welfare can be obtained only when sharing is allowed. Thus, our results complement theirs: we establish that, even when sharing is infeasible, a reasonable fairness guarantee can still be provided in terms of the maximin share.

We now proceed to general graphs. We consider an ordinal relaxation called *1-out-of- k maximin share*, denoted by $\text{MMS}_i^{k,s}$, or simply MMS_i^k when s is clear from the context. The idea is that instead of taking partitions into n parts as in the canonical maximin share, we allow partitions into k parts, where $k > n$ is a given parameter. For each graph G , let $\text{FVSNUM}(G)$ be the *feedback vertex set number* of G , that is, the minimum number of vertices whose removal makes the graph acyclic.³

Theorem 3.4. *Let $s = 0$, and assume that all allocated pieces must be completely disjoint. For any graph G and any n agents with arbitrary monotone valuations, there exists an allocation of G in which each agent i receives a connected piece with value at least $\text{MMS}_i^{n+\text{FVSNUM}(G)}$.*

Proof. Let $r := \text{FVSNUM}(G)$. Pick a subset of r vertices such that after deleting these vertices the remaining graph is a forest, and delete them (while keeping their adjacent edges intact as open intervals). For each agent, consider her 1-out-of- $(n+r)$ maximin partition. Since all pieces of each partition are pairwise disjoint, each vertex deletion harms at most one part in each partition. Thus, once the graph becomes a forest, for every agent, at least n parts remain. By Theorem 3.2, there is an allocation in which every agent i gets at least one of her n parts, and therefore value at least MMS_i^{n+r} . \square

For every graph G , let $\text{MMSRANK}(G)$ be the smallest integer $r \geq 0$ such that for any integer $n \geq 1$ and any n agents with arbitrary monotone valuations, there exists an allocation of G in which each agent i receives a connected piece with value at least MMS_i^{n+r} . Theorem 3.4 shows that $\text{MMSRANK}(G) \leq \text{FVSNUM}(G)$; we do not know if this inequality is tight.

When agents' valuations are *additive*, Theorem 3.4 is *not* tight for some graphs. In particular, for the cycle graph C we have $\text{FVSNUM}(C) = 1$, but MMS_i^n can be guaranteed to all agents (by transforming C into an interval, for which a proportional allocation exists). Defining $\text{MMSRANKADD}(G)$ analogously to $\text{MMSRANK}(G)$, but for additive valuations only, we thus obtain $\text{MMSRANKADD}(C) = 0$.

³Computing $\text{FVSNUM}(G)$ is NP-hard [Karp, 1972], but here we only use it for existence proofs. Note that $\text{FVSNUM}(G)$ is upper-bounded by the *circuit rank* of G , that is, the minimum number of edges whose removal makes the graph acyclic. The circuit rank of a graph $G = (V, E)$ with c connected components is $|E| - |V| + c$.

4 Positive Separation

In this section, we consider the case where a separation constraint is imposed, that is, $s > 0$.

4.1 Cutting Forests

First, we assume that G is a forest. We start with several lemmas about the DIST^G metric when G is a tree.⁴

By definition of a tree, for any two points $x, y \in G$, there is a unique (simple) path between x and y . Denote this unique path by $\text{PATH}^G(x, y)$ or $x \rightarrow y$, and observe that the length of this path is $\text{DIST}^G(x, y)$. We say that two subsets of G are *essentially-disjoint* if they intersect in at most a single point.

Lemma 4.1. *Let G be a tree, $X \subseteq G$ a closed connected subset, and $r \in G \setminus X$ a point. There exists a unique point $x_* \in X$ (a function of X and r) with the following properties:*

(a) *The path from any point in X to r passes through x_* . That is, for any $y \in X$ it holds that $x_* \in \text{PATH}^G(y, r)$.*

(b) *x_* is closer to r than any other point in X is. That is, $\text{DIST}^G(x_*, r) < \text{DIST}^G(y, r)$ for all $y \in X \setminus \{x_*\}$.*

(c) *For any $y \in X$ it holds that $\text{DIST}^G(y, r) = \text{DIST}^G(y, x_*) + \text{DIST}^G(x_*, r)$.*

Denote the unique point x_* guaranteed by Lemma 4.1 by $\text{NEAREST}(X, r)$.⁵ The lemma can be generalized as follows:

Lemma 4.2. *Let G be a tree and $X, R \subseteq G$ be closed connected subsets with $X \cap R = \emptyset$. There exists a unique point $x_* \in X$ (a function of X and R) satisfying the following properties:*

(a) *The path from any point in X to any point in R passes through x_* .*

(b) *For every point $r \in R$, x_* is closer to r than any other point in X is.*

For non-intersecting subsets $X, R \subseteq G$, denote the unique point x_* guaranteed by Lemma 4.2 by $\text{NEAREST}(X, R)$. Note that $\text{NEAREST}(X, R) \neq \text{NEAREST}(R, X)$: the former is in X while the latter is in R . We now define “ s -good” pieces similarly to 0-good pieces in Lemma 3.3. Given a graph G and a family $\mathbf{X} = (X_1, \dots, X_k)$ of closed connected pieces of G , a piece X_j is called *s -good* provided that for all $j_1, j_2 \in [k]$ the following holds: If $\text{DIST}^G(X_{j_1}, X_j) < s$ and $\text{DIST}^G(X_{j_2}, X_j) < s$, then $\text{DIST}^G(X_{j_1}, X_{j_2}) < s$.

Lemma 4.3. *Let G be a tree and $\mathbf{X} = (X_1, \dots, X_k)$ a family of closed connected subsets of G , for some integer $k \geq 1$. If $s > 0$, then for some $j^* \in [k]$, the piece X_{j^*} is s -good.*

Proof. Fix an arbitrary point $r \in G$ as the tree root. For every $j \in [k]$, let $x_j := \text{NEAREST}(X_j, r)$ and $d_j := \text{DIST}^G(x_j, r)$. Let $j^* \in \arg \max_{j \in [k]} d_j$, so that X_{j^*} is a piece in \mathbf{X} farthest from r ; abusing notation slightly, we refer to this piece as X_0

⁴In this case, the metric space DIST^G is known as a *real tree* or an \mathbb{R} -tree [Bestvina, 2001].

⁵ $\text{NEAREST}(X, r)$ is closely related to the concept of *median* in a tree. Given three points x, y, z of a tree graph, there is a unique point in $\text{PATH}^G(x, y) \cap \text{PATH}^G(y, z) \cap \text{PATH}^G(z, x)$; this point is called the *median* of x, y, z . More generally, any graph with this uniqueness property is called a *median graph*; such graphs have been studied in voting theory [Nehring and Puppe, 2007].

and let $x_0 := x_{j^*}$. We claim that X_0 is s -good. To prove this, we need several auxiliary claims concerning X_0 ; see the full version of our paper [Elkind *et al.*, 2021a] for the proofs.

Claim 1. For each $j \in [k]$, if X_0 intersects X_j , then $x_0 \in X_j$.

Claim 2. For each $j \in [k]$, if X_0 does not intersect X_j , then $x_0 = \text{NEAREST}(X_0, X_j)$.

For every $j \in [k]$ such that X_0 and X_j are disjoint, let $y_j := \text{NEAREST}(X_j, X_0) = \text{NEAREST}(X_j, x_0)$ and $z_j := \text{NEAREST}(\text{PATH}^G(x_0, y_j), r)$, i.e., z_j is the point at which the path from x_0 to r meets the path from y_j to r . Our next auxiliary claim is:

Claim 3. For each $j \in [k]$ such that X_0 does not intersect X_j , $\text{DIST}^G(y_j, z_j) \leq \text{DIST}^G(x_0, z_j)$.

By Claims 1 and 2 and the definition of y_j , we get the following useful formula for the distance between X_0 and X_j , for every $j \in [k]$:

$$\text{DIST}^G(X_0, X_j) = \text{DIST}^G(x_0, X_j) = \text{DIST}^G(x_0, y_j), \quad (1)$$

where the second equality holds if X_0 and X_j are disjoint. In other words, to measure the distance between the sets X_0 and X_j , we can consider the distance between the single point $x_0 \in X_0$ (the point closest to r) and X_j . If $X_0 \cap X_j = \emptyset$, so that y_j is defined, then we can consider the distance between x_0 and the single point $y_j \in X_j$ (the point closest to X_0).


We are finally ready to prove that X_0 is s -good. Consider two arbitrary pieces of \mathbf{X} , say X_1 and X_2 , such that $\text{DIST}^G(X_1, X_0) < s$ and $\text{DIST}^G(X_2, X_0) < s$. We have to prove that $\text{DIST}^G(X_1, X_2) < s$.

If $X_0 \cap X_1 \neq \emptyset$, then $x_0 \in X_1$ by Claim 1, so

$$\text{DIST}^G(X_1, X_2) \leq \text{DIST}^G(x_0, X_2) = \text{DIST}^G(X_0, X_2) < s,$$

where the equality holds by (1). An analogous claim holds if $X_0 \cap X_2 \neq \emptyset$. If $X_1 \cap X_2 \neq \emptyset$, then obviously $\text{DIST}^G(X_1, X_2) = 0 < s$. So from now on suppose that $X_0 \cap X_1 = X_0 \cap X_2 = X_1 \cap X_2 = \emptyset$. Then y_1 and y_2 are defined, and by (1) we have $\text{DIST}^G(y_1, x_0) < s$ and $\text{DIST}^G(y_2, x_0) < s$.

By definition of z_1 and z_2 , the path $x_0 \rightarrow r$ must pass through both z_1 and z_2 . Without loss of generality, suppose that z_1 comes no later than z_2 on this path, so $\text{DIST}^G(x_0, z_1) \leq \text{DIST}^G(x_0, z_2)$, as in the illustration below:



Consider the path $y_1 \rightarrow z_1 \rightarrow z_2 \rightarrow y_2$. The length of this path is at most

$$\begin{aligned} & \text{DIST}^G(y_1, z_1) + \text{DIST}^G(z_1, z_2) + \text{DIST}^G(z_2, y_2) \\ & \leq \text{DIST}^G(x_0, z_1) + \text{DIST}^G(z_1, z_2) + \text{DIST}^G(z_2, y_2) \\ & = \text{DIST}^G(x_0, y_2) = \text{DIST}^G(X_0, X_2) < s; \end{aligned}$$

the first inequality holds by Claim 3 and the last equality by (1). We have demonstrated a path of length shorter than s from a point $y_1 \in X_1$ to a point $y_2 \in X_2$; this proves that $\text{DIST}^G(X_1, X_2) < s$. It follows that X_0 is s -good. \square

With Lemma 4.3 in hand, we can now establish the existence of a maximin allocation for forests by using similar arguments as in Theorem 3.2. In particular, when we allocate an s -good part $T_{i,j}$ to agent i , we remove all portions of the tree that are within distance s of $T_{i,j}$, and divide the remaining cake recursively among the remaining agents.

Theorem 4.4. *Let G be a forest and $s > 0$. For agents with arbitrary monotone valuations, a maximin allocation exists.*

4.2 Cutting General Graphs

We now proceed to general graphs. Even when the graph is a simple cycle, Elkind *et al.* [2021c] showed that a maximin allocation does not necessarily exist; however, the 1 -out-of- $(n + 1)$ maximin share can be guaranteed. We present here a more general theorem analogous to Theorem 3.4. However, our theorem requires the assumption that the length of each edge is at least s . We remark that the separation parameter s is generally small in our motivating applications such as transition or buffer zones, so this assumption is realistic. Nevertheless, it is an interesting question whether the result continues to hold without this assumption.

Theorem 4.5. *Let $s > 0$. Let G be a graph in which the length of each edge is at least s . For any n agents with arbitrary monotone valuations, there exists an s -separated allocation in which each agent i receives a connected piece with value at least $\text{MMS}_i^{n+\text{FVSNUM}(G)}$.*

Proof. Let $r := \text{FVSNUM}(G)$. Let u_1, \dots, u_r be a set of vertices such that, when they are deleted from G , the remaining graph is a forest. For each $j \in [r]$, remove an open interval of length $s/2$ from each edge adjacent to u_j . Consider the 1 -out-of- $(n + r)$ maximin partition of each agent. For each $j \in [r]$, the diameter of the set removed due to the vertex u_j is less than s . Hence, each removed set overlaps at most one part of each agent’s partition. Therefore, once the graph becomes a forest, for every agent, at least n parts remain. By the argument in the proof of Theorem 4.4, there exists an s -separated allocation in which every agent i gets at least one of the remaining parts of her s -separated maximin partition, and therefore value at least MMS_i^{n+r} . Then, reconstruct the original graph by putting back the removed intervals of length $s/2$. The allocation remains s -separated. \square

As in the case $s = 0$, we do not know whether the factor $n + \text{FVSNUM}(G)$ is tight in general. Below, we present a class of graphs for which we can obtain a tight bound. In particular, we consider the family of graphs such that the feedback vertex set number of each connected component is at most 1 (that is, every connected component is either a tree, or can be made acyclic by removing a single vertex). Set $N := \min(n + \text{FVSNUM}(G), 2n - 1)$.

Theorem 4.6. *The following statements hold for any real number $s > 0$ and integer $n \geq 1$:*

- (a) *Let G be a vertex-disjoint union of graphs with FVS number ≤ 1 , such that the length of each edge is at least s . For any n agents with monotone valuations, there is an allocation in which each agent i receives value at least MMS_i^N .*
- (b) *For any integer $r \geq 1$, there exists a graph G with $\text{FVSNUM}(G) = r$ (specifically, a union of r cycles and zero or more trees), and n agents with additive valuations, such that no allocation gives each agent i at least MMS_i^{N-1} .*

Theorem 4.6 shows that Theorem 4.5 is tight for unions of trees and cycles when the number of cycles is less than n . However, exact bounds for other graphs are not known.

5 Discussion

In this work, we have studied the division of a graphical cake using the maximin share notion, both with and without separation constraints. Our most technically challenging result shows that a maximin allocation exists for positive separation whenever the graph is acyclic. A tempting approach to simplify this proof is by using the “last diminisher” method, wherein each agent can trim a proposed piece as long as the value of the piece that remains after trimming is at least the agent’s maximin share. Indeed, this method was used by Bouveret *et al.* [2017] to establish the existence result for trees in the context of indivisible goods, and the algorithm of Elkind *et al.* [2021c] for interval cakes can also be seen as a version of last diminisher. We remark here that the approach does not work in our setting with separation. Indeed, consider the subtree in Figure 1, where two of Alice’s parts in her maximin partition are bold, while one of Bob’s parts is dashed. The last diminisher method may allocate Bob’s part to him; however, when we take the separation requirement into account, we cannot allocate either of Alice’s parts in its entirety. Note that the dashed piece is not s -good. This example precisely demonstrates why, in the proof of Theorem 4.4, we need to reason carefully about real trees in order to guarantee the existence of an s -good piece.




Figure 1: An example subtree in which the last diminisher method fails to produce a maximin allocation.

Throughout the paper, we have formulated a number of open questions that stem naturally from our work. More generally, our work builds upon an active line of research that incorporates realistic constraints in fair division problems. We believe that identifying and studying such considerations will lead to technically intriguing questions as well as practically useful fairness guarantees.

Acknowledgments

This work was partially supported by the Israel Science Foundation under grant number 712/20 and by an NUS Start-up Grant. We would like to thank the anonymous reviewers for their valuable comments.

References

- [Abebe *et al.*, 2017] Rediet Abebe, Jon M. Kleinberg, and David C. Parkes. Fair division via social comparison. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 281–289, 2017.
- [Alijani *et al.*, 2017] Reza Alijani, Majid Farhadi, Mohammad Ghodsi, Masoud Seddighin, and Ahman S. Tajik. Envy-free mechanisms with minimum number of cuts. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 312–318, 2017.
- [Arunachaleswaran *et al.*, 2019] Eshwar Ram Arunachaleswaran, Siddharth Barman, and Nidhi Rathi. Fair division with a secretive agent. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1732–1739, 2019.
- [Aziz *et al.*, 2018] Haris Aziz, Sylvain Bouveret, Ioannis Caragiannis, Ira Giagkousi, and Jérôme Lang. Knowledge, fairness, and social constraints. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 4638–4645, 2018.
- [Balkanski *et al.*, 2014] Eric Balkanski, Simina Brânzei, David Kurokawa, and Ariel D. Procaccia. Simultaneous cake cutting. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 566–572, 2014.
- [Bei and Suksompong, 2021] Xiaohui Bei and Warut Suksompong. Dividing a graphical cake. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [Bei *et al.*, 2017] Xiaohui Bei, Youming Qiao, and Shengyu Zhang. Networked fairness in cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 3632–3638, 2017.
- [Bei *et al.*, 2018] Xiaohui Bei, Guangda Huzhang, and Warut Suksompong. Truthful fair division without free disposal. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 63–69, 2018.
- [Bei *et al.*, 2021] Xiaohui Bei, Ayumi Igarashi, Xinhang Lu, and Warut Suksompong. The price of connectivity in fair division. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [Bestvina, 2001] Mladen Bestvina. \mathbb{R} -trees in topology, geometry, and group theory. In R. J. Daverman and R. B. Sher, editors, *Handbook of Geometric Topology*, chapter 2, pages 55–91. Elsevier, 2001.
- [Biswas and Barman, 2018] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.
- [Bouveret *et al.*, 2017] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 135–141, 2017.
- [Brams and Taylor, 1996] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [Brânzei *et al.*, 2016] Simina Brânzei, Ioannis Caragiannis, David Kurokawa, and Ariel D. Procaccia. An algorithmic framework for strategic fair division. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pages 418–424, 2016.
- [Budish, 2011] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- [Elkind *et al.*, 2021a] Edith Elkind, Erel Segal-Halevi, and Warut Suksompong. Graphical cake cutting via maximin share. *arXiv preprint arXiv:2105.04755*, 2021.
- [Elkind *et al.*, 2021b] Edith Elkind, Erel Segal-Halevi, and Warut Suksompong. Keep your distance: Land division with separation. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- [Elkind *et al.*, 2021c] Edith Elkind, Erel Segal-Halevi, and Warut Suksompong. Mind the gap: Cake cutting with separation. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [Goldberg *et al.*, 2020] Paul W. Goldberg, Alexandros Hollender, and Warut Suksompong. Contiguous cake cutting: Hardness results and approximation algorithms. *Journal of Artificial Intelligence Research*, 69:109–141, 2020.
- [Hosseini *et al.*, 2020] Hadi Hosseini, Ayumi Igarashi, and Andrew Searns. Fair division of time: Multi-layered cake cutting. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 182–188, 2020.
- [Igarashi and Zwicker, 2021] Ayumi Igarashi and William S. Zwicker. Fair division of graphs and of tangled cakes. *arXiv preprint arXiv:2102.08560*, 2021.
- [Karp, 1972] Richard M. Karp. Reducibility among combinatorial problems. In *Proceedings of a Symposium on the Complexity of Computer Computations*, pages 85–103, 1972.
- [Kurokawa *et al.*, 2018] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 64(2):8:1–8:27, 2018.
- [Li *et al.*, 2015] Minming Li, Jialin Zhang, and Qiang Zhang. Truthful cake cutting mechanisms with externalities: Do not make them care for others too much! In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 589–595, 2015.
- [Lonc and Truszczynski, 2020] Zbigniew Lonc and Miroslaw Truszczynski. Maximin share allocations on cycles. *Journal of Artificial Intelligence Research*, 69:613–655, 2020.
- [Menon and Larson, 2017] Vijay Menon and Kate Larson. Deterministic, strategyproof, and fair cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 352–358, 2017.
- [Nehring and Puppe, 2007] Klaus Nehring and Clemens Puppe. The structure of strategy-proof social choice — Part I: General characterization and possibility results on median spaces. *Journal of Economic Theory*, 135(1):269–305, 2007.
- [Procaccia, 2013] Ariel D. Procaccia. Cake cutting: Not just child’s play. *Communications of the ACM*, 56(7):78–87, 2013.
- [Procaccia, 2016] Ariel D. Procaccia. Cake cutting algorithms. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13, pages 311–329. Cambridge University Press, 2016.
- [Robertson and Webb, 1998] Jack Robertson and William Webb. *Cake-Cutting Algorithms: Be Fair if You Can*. Peters/CRC Press, 1998.
- [Segal-Halevi, 2021] Erel Segal-Halevi. Fair multi-cake cutting. *Discrete Applied Mathematics*, 291:15–35, 2021.