

Coalition Games on Interaction Graphs: A Horticultural Perspective

Nicolas Bousquet^{1,4}, Zhentao Li³, and Adrian Vetta^{1,2}

¹Department of Mathematics and Statistics, McGill University

² School of Computer Science, McGill University

³ École Normale Supérieure de Paris.

⁴ Group for Research in Decision Analysis (GERAD), HEC Montréal

March 2, 2015

Abstract

We examine cooperative games where the viability of a coalition is determined by whether or not its members have the ability to communicate amongst themselves independently of non-members. This necessary condition for viability was proposed by Myerson in [14] and is modeled via an interaction graph $G = (V, E)$; a coalition $S \subseteq V$ is then viable if and only if the induced graph $G[S]$ is connected. The non-emptiness of the core of a coalition game can be tested by a well-known covering LP. Moreover, the integrality gap of its dual packing LP defines exactly the multiplicative least-core and the relative cost of stability of the coalition game. This gap is upper bounded by the packing-covering ratio which, for graphical coalition games, is known to be at most the treewidth of the interaction graph plus one [13].

We examine the packing-covering ratio and integrality gaps of graphical coalition games in more detail. We introduce the thicket parameter of a graph, and prove it precisely measures the packing-covering ratio. It also approximately measures the primal and dual integrality gaps. The thicket number provides an upper bound of both integrality gaps. Moreover we show that for any interaction graph, the primal integrality gap is, in the worst case, linear in terms of the thicket number while the dual integrality gap is polynomial in terms of it. At the heart of our results, is a graph theoretic minmax theorem showing the thicket number is equal to the minimum width of a vine decomposition of the coalition graph (a vine decomposition is a generalization of a tree decomposition). We also explain how the thicket number relates to the VC-dimension of the set system produced by the game.

1 Introduction

At the heart of cooperative game theory is the problem of how a group of agents should share the wealth that they collectively create. Its foremost concept is the *core* whose roots date back to Edgeworthian bargaining and cooperative improvement ([5]; see also [10]). It was first formalized by Gillies in [7, 8] via a *coalition game* $\mathcal{G} = (I, v)$ with a set I of agents and a valuation function $v : 2^I \rightarrow \mathbb{N}$. The *core* of the coalition game is the set of feasible solutions to:

$$\begin{aligned}
\sum_{i:i \in I} x_i &= v(I) \\
\sum_{i:i \in S} x_i &\geq v(S) \quad \forall S \subset I \\
x_i &\geq 0
\end{aligned}$$

Informally, we are allocating x_i to agent i and $v(S)$ represents the amount of wealth that the coalition S can generate by itself. Consequently, the coalition S will *block* any distribution scheme that does not allocate its members at least $v(S)$ in total. Thus the wealth $v(I)$ of the *grand coalition* must be distributed in such a fashion that no coalition wishes to block the allocation. The core is the set of vectors of payoffs that have this property.

This definition immediately prompts two questions: (i) What processes enable the formation of coalitions? (ii) Even if coalitions can form and negotiate, do core solutions exist? Concerning the former question, it is unrealistic to assume that every subset of agents has the ability to act as a collective. Indeed, Myerson [15] argued that feasible coalitions require structural properties that enable them to function. Clearly, one necessary property is that “communication” is possible between members of the coalition and [14] formalized this ability using an *interaction (communication) graph* $G = (I, E)$. Here a pair of agents induces an edge in G if they are able to interact and a coalition S is feasible if S induces a connected subgraph of G . Observe that two members of a feasible coalition do not need to be able to interact directly but they must be able to communicate indirectly via chains consisting of other members of the coalition. Thus, a coalition S is *viable* if and only if the induced subgraph $G[S]$ is connected – in particular $v(S) = 0$ when $G[S]$ is disconnected. Such *graphical coalition games* are the focus of this paper.¹

For the latter question, the core is often empty. Indeed, it is straightforward to verify that the core of the game \mathcal{G} is non-empty *if and only if* the following primal linear program has an optimal (fractional) solution whose value $\kappa^f(\mathcal{G})$ equals $v(I)$.

$$\begin{aligned}
\text{Covering-LP: } \min \quad & \sum_{i \in I} x_i \\
\text{s.t.} \quad & \sum_{i:i \in S} x_i \geq v(S) \quad \forall S \subseteq I \\
& x_i \geq 0
\end{aligned}$$

Interestingly, there is an elegant graphical characterization for when a graphical coalition game has a non-empty core for all possible valuation functions v . Namely, a graphical coalition game is *strongly balanced* if and only if the interaction graph G is a forest [11].

1.1 The Least-Core and the Relative Cost of Stability

Given the possible emptiness of the core, it is natural to consider solutions where the core constraints in this **Covering-LP** are relaxed. Specifically, for each feasible coalition S , given $\alpha \geq 1$, there is a constraint $\alpha \cdot \sum_{i:i \in S} x_i \geq v(S)$. These constraints imply that a coalition S will not block an allocation unless it can unilaterally improve its total wealth by more than an α factor. The set of feasible solutions then form the α -*core*. The minimum α for which the α -core is non-empty arises when $\alpha^* = \frac{\kappa^f(\mathcal{G})}{v(I)}$. The α^* -core is called the (*multiplicative*)

¹See [6] for a general introduction to cooperative games on networks.

least-core.² Moreover, we can also determine the least-core by considering the dual of the primal linear program.

$$\begin{aligned} \text{Packing-LP: } \max \quad & \sum_{S:S \subseteq I} v(S) \cdot y_S \\ \text{s.t.} \quad & \sum_{S \subseteq I: i \in S} y_S \leq 1 \quad \forall i \in I \\ & y_S \geq 0 \quad \forall S \subseteq I \end{aligned}$$

Let $\rho^f(\mathcal{G})$ be the optimal fractional solution to this Packing-LP, and let $\rho(\mathcal{G})$ be the optimal integral solution. Then α^* is exactly equal to the *dual integrality gap* $\frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})}$. To verify this, observe that $\rho^f(\mathcal{G}) = \kappa^f(\mathcal{G}) = \alpha^* \cdot v(I)$, by strong duality. But the optimal integral solution to the dual has value $\rho(\mathcal{G}) = v(I)$; simply set $y_I = 1$ and $y_S = 0$ for every $S \neq I$. Here we are making the standard assumption in the literature that, for any coalition game, the valuation function v is superadditive. In particular, property reflects the simple observation that any collective has the option to voluntarily partition itself into subgroups to generate wealth. Thus, $\alpha^* = \frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})}$.

This paper will study least-cores in coalition games over interaction graphs. Interestingly, multiplicative least-cores are equivalent to the concept of the *relative cost of stability*. Bachrach et al. [2] asked how much it would cost (an external authority) to stabilize a coalition game; *i.e.* what is the minimum total payment required such that no coalition can benefit by blocking the allocation. The relative cost of stability [12] is then defined to be the ratio between this minimum payment and the total wealth the grand coalition can generate, namely $\frac{\kappa^f(\mathcal{G})}{v(I)} = \alpha^*$. Thus, the relative cost of stability is also given by the dual integrality gap.

1.2 Our Results

So to develop an understanding of coalition games, we must study the primal and dual linear programs. In particular, we will focus on the graphical coalition games of [14]. Specifically, we are interested in how the primal and dual integrality gaps vary with the topology of the interaction graph.

As inferred by our nomenclature, the primal and dual form a pair of packing and covering linear programs. Thus a natural starting point is to consider the *packing-covering ratio* of a game \mathcal{G} ; this is the ratio $\frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})}$ between the values of the optimal integral solutions to the primal and the dual. Observe that, by strong duality, the packing-covering ratio is the product of the primal integrality gap $\frac{\kappa(\mathcal{G})}{\kappa^f(\mathcal{G})}$ and the dual integrality gap $\frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})}$. Consequently, the packing-covering ratio trivially upper bounds both integrality gaps. Packing-covering ratios have been studied extensively in graph theory. Special attention has focused on problems with the Erdős-Pósa property, where the ratio is a function of the packing number and is otherwise independent of the graph. Interestingly, whilst graphical coalition games do not have the Erdős-Pósa property, the packing-covering ratio can be bounded by an important parameter of the interaction graph, namely, treewidth. Indeed, Meir et al. [13] proved that, for any valuation function (game \mathcal{G}) over an interaction graph G , the packing-covering ratio is at most the treewidth $\omega(G)$ plus one.

²The least-core is often defined with respect to an additive (ε), not a multiplicative (α), guarantee. Since additive guarantees are not scale invariant, it is preferable here to focus upon multiplicative guarantees.

We extend the work of Meir et al. in several ways. First, we show that structurally treewidth is not the most appropriate invariant in understanding the packing-covering ratio. The topological parameter that corresponds exactly to the packing-covering ratio is a concept we term the *thicket number* of the graph. Specifically, in Section 3, we show that for every coalition game \mathcal{G} over a graph G the packing-covering ratio is at most the thicket number, $\tau(G)$, of the graph. Conversely, for *every* graph G there exists a coalition game \mathcal{G} for which the packing-covering ratio is at least the thicket number.

Theorem 1.1. *For any interaction graph G , the packing-covering ratio satisfies:*

$$\tau(G) \leq_{\exists} \frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} \tau(G)$$

Observe that, in order to concisely formulate our results, we use the notation \leq_{\exists} and \leq_{\forall} . Here \leq_{\exists} means that *there exists* a game \mathcal{G} over the interaction graph G such that the inequality is satisfied, and \leq_{\forall} means that *for every* game \mathcal{G} over G the inequality is satisfied.³

Theorem 1.1 relies on a graphical minmax result that we prove in Section 2. Specifically, we show that thickets have a dual notion called *vine decompositions*. These decompositions can be viewed as a “thin” relative of tree decompositions. In particular, the vine width of a graph is at most the treewidth plus one, and is typically smaller.

In principle, the primal and dual integrality gaps could be much less than the thicket number. However, we prove the thicket number is (approximately) the correct measure for these integrality gaps as well. Specifically, in Section 6 we prove

Theorem 1.2. *For any interaction graph G , the primal integrality gap satisfies:*

$$\frac{1}{4}\tau(G) \leq_{\exists} \frac{\kappa(\mathcal{G})}{\kappa^f(\mathcal{G})} \leq_{\forall} \tau(G)$$

Interestingly, unlike for the packing-covering ratio, the upper and lower bounds cannot be closed completely for the primal integrality gap. Indeed, for any graph G there is a constant a_G such that $a_G \cdot \tau(G) \leq_{\exists} \frac{\kappa(\mathcal{G})}{\kappa^f(\mathcal{G})} \leq_{\forall} a_G \cdot \tau(G)$. However a_G really does vary with the graph. In particular, we show that $a_G \rightarrow 1$ for the family of graphs that correspond to the powers of paths. On the other hand, we prove that $a_G \leq \frac{1}{2}$ for cliques. It follows that the constant 1 in the upper bound in Theorem 1.3 cannot be decreased, whilst the constant in the lower bound cannot be increased above $\frac{1}{2}$.

Next consider the dual integrality gap. In Section 5 we prove

Theorem 1.3. *There exist c and δ such that for any interaction graph G , the dual integrality gap satisfies:*

$$c \cdot \tau(G)^\delta \leq_{\exists} \frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} \tau(G)$$

Again, it is not possible to close the upper and lower bounds for the dual integrality gap completely. Even more interestingly, the polynomial range between the upper and lower bounds is necessary. This is since for the family of grid graphs, the exponent of $\tau(G)$ in the

³We remark that inequalities of the form \geq_{\exists} and \geq_{\forall} are not interesting from a game-theoretic perspective. Indeed, for any graph G , we have that $\frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})} \geq_{\forall} 1$, by weak duality. Furthermore, the packing-covering ratio of any coalition game with one viable coalition trivially equals 1. So $1 \geq_{\exists} \frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})}$.

lower bound must be 1, whereas for the family of cliques the exponent of $\tau(G)$ in the upper bound is $\frac{1}{2}$. It follows that the exponent 1 in the upper bound of Theorem 1.2 cannot be decreased, whilst the exponent δ in the lower bound cannot be increased above $\frac{1}{2}$. We remark that the value of δ in Theorem 1.2 relies on the existence of a grid minor of polynomial size in the treewidth of the graph, a deep result of [4]. Determining the best possible order of the polynomial in the lower bound of Theorem 1.2 (between δ and $\frac{1}{2}$) is an interesting open question.

Finally, in Section 4, we show how the VC-dimension may also be used to bound the packing-covering ratio and the integrality gaps. Given our previous discussion, the resultant bounds must be weaker than those obtainable via the thicket number. Indeed, we show how these VC-dimension bounds may be derived from the thicket number bounds.

2 Thickets and Vines

Recall that [13] show that the packing-covering ratio can be bounded in terms of the treewidth of the interaction graph G . To understand this, we begin with a brief review of treewidth. We will then introduce a better fitting parameter for analyzing coalition games.

2.1 Tree Decompositions and Brambles

Treewidth provides a measure of how closely a graph shares some structural separation properties possessed by trees. Formally, given an undirected graph $G = (V, E)$ we may represent it by a tree $T = (N, L)$ and a labeling $\ell : N \rightarrow 2^V$.⁴ The labeling assigns to each node $t \in T$ a subset $\ell(t) = V_t$ of vertices of G . For each $v \in V$ we denote by T_v the set of nodes in T for which v is included in the label, *i.e.* $T_v = \{t : v \in V_t\}$. We say that a tree and labeling, (T, ℓ) , is a *tree decomposition* of G if:

- (i) For each vertex v of G , the set T_v is a **non-empty** and **connected** subgraph of T .
- (ii) For each edge $e = (u, v)$ in G , the subtrees T_u and T_v **intersect** in T .

The *width* of a tree decomposition (T, ℓ) of G is the size of the largest label of a node in T minus one.⁵ The *treewidth*, $\omega(G)$, is the minimum width of a tree decomposition of G . Meir et al. [13] show that treewidth relates to coalition games via the following bound.

Theorem 2.1. *Meir et al. [13] For any interaction graph G , the packing-covering ratio satisfies:*

$$\frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} \omega(G) + 1$$

To delve further into this topic, it is important to note that there are combinatorial structures called brambles that provide a dual notion for tree decompositions. A *bramble* is a collection $\mathcal{F} = \{F_1, F_2, \dots, F_p\}$ of sets such that:

- (i) Each $F_i \subseteq V$ induces a connected subgraph of G , and
- (ii) Every pair F_i and F_j in \mathcal{F} *either* intersect (share a vertex) *or* are adjacent (there is an edge with one endpoint in F_i and one endpoint in F_j).

⁴For clarity, we will refer to vertices and edges in G and nodes and links in T .

⁵The decision to subtract one was made to ensure trees have treewidth one. Unfortunately, as is apparent, this choice leads to an unaesthetic “plus one” in many theorems concerning treewidth.

The *hitting size* $\beta(\mathcal{F})$ of a bramble \mathcal{F} is the minimum size of a subset of vertices that intersects each set F_i in \mathcal{F} . The *bramble number* $\beta(G) = \max_{\mathcal{F}} \beta(\mathcal{F})$ is the maximum hitting size of any bramble in G . Seymour and Thomas [18] proved the following minmax theorem:

Theorem 2.2. [18] *The bramble number $\beta(G)$ is equal to the treewidth $\omega(G) - 1$.*

Brambles (rather than tree decompositions) directly relate to coalition games. Moreover, the relationship is actually through combinatorial structures we call thickets.

2.2 Thickets

Let $G = (V, E)$ be an undirected graph. A *thicket* $\mathcal{H} = \{H_1, H_2, \dots, H_p\}$ is a collection of sets such that:

- (i) Each $H_i \subseteq V$ induces a connected subgraph of G , and
- (ii) Every pair H_i and H_j in \mathcal{H} intersects.

Observe that thickets differ from brambles in that they must pairwise intersect – adjacency is not sufficient. The *hitting size* $\tau(\mathcal{H})$ of a thicket \mathcal{H} is the minimum size of a vertex set that intersects each set H_i in \mathcal{H} . The *thicket number* $\tau(G) = \max_{\mathcal{H}} \tau(\mathcal{H})$ is the maximum hitting size of any thicket in G .

Intuitively, thickets are indeed the objects that directly correspond to the packing-covering ratio. The packing number of a thicket \mathcal{H} is exactly one, but its covering number is $\tau(\mathcal{H})$. We will formalize this intuition in Section 3. First, let's see three simple classes of graphs that will illustrate the concept of thickets, and which will also be very useful in the technical results that follow.

Example 1: Trees. Let $G = T_n$ be a tree on n vertices. A thicket $\mathcal{H} = \{H_1, \dots, H_p\}$ on T_n then consists of a collection of pairwise intersecting subtrees. It is well-known, by the *Helly Property* of trees, that such a collection must contain a common vertex. Thus the thicket number, $\tau(T_n)$, of a tree T_n is at most 1.

Example 2: Cliques. Let $G = K_n$ be a clique on n vertices. A thicket $\mathcal{H} = \{H_1, \dots, H_p\}$ on K_n then consists of a collection of pairwise intersecting subcliques. Suppose the smallest of these cliques, say H_1 , has cardinality at most $\lceil \frac{1}{2}n \rceil$. Then, as H_1 itself intersects all of the sets in \mathcal{H} , we have a hitting set of cardinality $\lceil \frac{1}{2}n \rceil$. Otherwise all the cliques have cardinality at least $\lceil \frac{1}{2}n \rceil + 1$. But then any set X of cardinality $\lceil \frac{1}{2}n \rceil$ is a hitting set for \mathcal{H} . Thus $\tau(G) \leq \lceil \frac{1}{2}n \rceil$.

Example 3: Grids. Let $G = R_k$ be a $k \times k$ grid graph – the planar graph formed by a grid of k rows and k columns. Consider the thicket \mathcal{H} defined as follows. We have a set $H_{R,C} = R \cup C$, for each row R and each column C in the grid. Clearly each set is connected. Moreover, each pair of sets intersect. So \mathcal{H} is a thicket. Now take any vertex set X of cardinality less than $k - 1$. Since there are k rows, X must miss some row \hat{R} ; similarly it must miss some column \hat{C} . Hence, X is not a hitting set for \mathcal{H} as it does not intersect $H_{\hat{R},\hat{C}}$. So the hitting size $\tau(\mathcal{H})$ is at least k . Consequently, the thicket number, $\tau(R_k)$, of the grid R_k is at least k .

2.3 Vine Decompositions

Brambles are dual to tree decompositions, and thickets also have dual structures. Since the definition of a thicket is more stringent than that of a bramble, it must be the case that the definition of its dual structures is more relaxed than that of a tree decomposition. In particular, its dual will be a **thin** tree (let's call it a *vine*!).

Formally, given an undirected graph $G = (V, E)$, we construct a (vine) tree $T = (N, L)$. The labeling assigns to each node $t \in T$ a subset $\ell(t) = V_t$ of vertices of G . For each $v \in V$ we denote by T_v the set of nodes in T for which v is included in the label, *i.e.* $T_v = \{t : v \in V_t\}$. We say that a tree and labeling, (T, ℓ) , is a *vine decomposition* of G if:

- (i) For each vertex v of G , the set T_v is a **non-empty** and **connected** subgraph of T .
- (ii) For each edge $e = (u, v)$ in G , the subtrees T_u and T_v **intersect or are adjacent** in T .

The *width* of a vine decomposition (T, ℓ) of G is the size of the largest label of a node in T . The *vinewidth*, $\nu(G)$, is the minimum width of a vine decomposition of G . The main structural result of the paper is that the thicket number $\tau(G)$ is equal to the *vinewidth* $\nu(G)$. Before proving this result in Section 2.4, we will develop some understanding of vine decompositions. First, let's return to the simple examples of trees, cliques and grids.

Example 1: Trees. Observe that a tree $G = T_n$ gives a trivial vine decomposition of itself, that is $T = (T_n, \ell)$ where $\ell(v) = \{v\}$. This is a vine decomposition as for each edge (u, v) in G the two vertices are clearly still adjacent in T . Thus the *vinewidth* of a tree is at most 1.

Example 2: Cliques. Let $G = K_n$ be a clique on n vertices. The clique has *vinewidth* at most $\lceil \frac{1}{2}n \rceil$. The corresponding vine decomposition has two nodes, each containing (roughly) half the vertices. This is a vine decomposition as for each edge (u, v) in G the two vertices are either in the same node of T or in adjacent nodes.

Example 3: Grids. Let G be a $k \times k$ grid graph. Let the (vine) tree T be a path on k nodes. Let the i th node in the path satisfy $\ell(i) = C_i$, where C_i is the set of vertices in the i th column of G . This is a vine decomposition. For each $v \in V$ the set T_v is a singleton node in T , and is thus non-empty and connected. For each edge $e = (u, v)$ in G , either $T_u = T_v$ if u and v are in the same column of G , or T_u and T_v are adjacent nodes in T if u and v are in the same row of G . Clearly the width of this vine decomposition is k and, thus, the *vinewidth* $\nu(G)$ is at most k .

Now treewidth and *vinewidth* are at most a multiplicative factor two apart.

Theorem 2.3. *Vinewidth and treewidth are related by $\nu(G) - 1 \leq \omega(G) \leq 2\nu(G) - 1$. Moreover there exist graphs for which the lower and upper bounds are tight.*

Proof. Observe that a tree decomposition of G is a vine decomposition of G . Thus, $\nu(G) \leq \omega(G) + 1$. Note that this bound is almost tight for the grid R_k ; we have $\nu(R_k) = k$ (this follows as the thicket number equals the *vinewidth* (see Theorem 2.5 below)) and it is well-known that $\omega(R_k) = k$. In fact, graphs can be constructed for which this lower bound is exactly tight; we omit the details.

On the other hand, given a vine decomposition (T, ℓ) we can create a tree-decomposition $(\hat{T}, \hat{\ell})$ by augmenting it as follows. We replace each link in T by a path of length two in \hat{T} . For each new node $t \in \hat{T}$ in the middle of the path that replaced the link $(t_1, t_2) \in T$ we set $\hat{\ell}(t) = \ell(t_1) \cup \ell(t_2)$. It is easy to verify that this is a tree-decomposition. Furthermore the width of this tree-decomposition is at most $2\nu(G) - 1$ (since each label has size at most $2\nu(G)$). This upper bound is tight for cliques as K_n has treewidth $n - 1$ and vinewidth $\lceil \frac{n}{2} \rceil$. \square

As with tree decompositions, an important property of vine decompositions is that nodes in the (vine) tree correspond to separators in the original graph.

Lemma 2.4. *Let (T, ℓ) be a vine decomposition of G , and let t be an internal node in T . Then V_t is a separator in G .*

Proof. Take any node $t \in T$ with degree $r \geq 2$. Let T_1, T_2, \dots, T_r be the subtrees formed by the removal of t . Let $C_i = \bigcup_{x \in T_i} V_x \setminus V_t$. We claim that there is no edge between C_i and C_j for $i \neq j$ in $G \setminus V_t$. To see this, take any pair of vertices $x \in C_i$ and $y \in C_j$ where $i \neq j$. Since neither x nor y are in V_t their corresponding subtrees T_x and T_y can neither intersect nor be adjacent in T , because $T_x \subseteq T_i$ and $T_y \subseteq T_j$. Thus (x, y) is not an edge of G . \square

2.4 The Thicket-Vinewidth Duality Theorem

Here we present the thicket-vinewidth duality theorem.

Theorem 2.5. *The thicket number $\tau(G)$ is equal to the vinewidth $\nu(G)$.*

To prove this we apply the approach used by [18] to bound the bramble number. In particular we prove the following stronger result, which characterizes when a thicket can be extended to create a thicket with hitting size k .

Lemma 2.6. *For any thicket \mathcal{H} in G , exactly one of the following holds:*

- (a) *There is a thicket \mathcal{H}' with hitting size k such that $\mathcal{H} \subseteq \mathcal{H}'$.*
- (b) *There is a vine decomposition (T, ℓ) of G such that for any node $s \in T$ with $|V_s| \geq k$:

 - (i) s is a leaf in T , and
 - (ii) V_s is not a hitting set for \mathcal{H} .*

Before proving Lemma 2.6, let's see why it does give Theorem 2.5.

of Theorem 2.5. First we show that $\nu(G) \geq \tau(G)$. Take a thicket $\mathcal{H} = \{H_1, H_2, \dots, H_p\}$ and a vine decomposition $T = (N, L)$ of G . We will show the subtree $H[T_i]$ of T corresponding to each element of \mathcal{H} pairwise intersect, apply the Helly Property to find a common intersection for all $H[T_i]$ and see that the label of this node at the intersection is a hitting set for \mathcal{H} .

More precisely, as H_i is a connected subgraph of G , there is a tree R_i in G spanning the vertices of H_i . Now consider the subgraph of T induced by H_i ; that is $T[H_i] = \bigcup_{w \in H_i} T_w$. We claim that $T[H_i]$ is a tree (i.e., $T[H_i]$ is connected). Take any edge $(v_r, v_s) \in R_i$. As T is vine decomposition, we have that $T_{v_r} \cup T_{v_s}$ induces a connected subtree in T . Extending this argument over every edge of R_i implies $T[H_i]$ is connected.

Now take any pair H_i and H_j in \mathcal{H} , $1 \leq i < j \leq p$. Since \mathcal{H} is a thicket they both contain some vertex u in G . It follows that $T_u \subseteq T[H_i] \cap T[H_j]$. In particular, the $T[H_i]$, $1 \leq i \leq p$,

are pairwise intersecting subtrees of T . Therefore, by the Helly Property of trees, there is a node $t \in \bigcap_{i=1}^p T[H_i]$. We claim that the vertex set V_t in G is a hitting set for \mathcal{H} . To see this, take any $H_i \in \mathcal{H}$. We have $t \in T[H_i]$ and, thus, there is some $T_w \subseteq T[H_i]$ where $w \in V_t$. Since $w \in H_i \cap V_t$, the claim follows. Since $|V_t| \leq \tau(G)$, the thicket number is at most the vinewidth: $\tau(G) \leq \nu(G)$.

Next we must prove that $\nu(G) \leq \tau(G)$. This follows from Lemma 2.6. To see this, let $\mathcal{H} = \emptyset$ be the empty thicket and $k = \nu(G)$. Then either (a) there is a thicket \mathcal{H}' with hitting size k or (b) there is a vine decomposition (T, ℓ) such that if $|V_s| \geq k$ then s is a leaf in T and V_s is not a hitting set for \mathcal{H} . If (a) holds, then \mathcal{H} is a thicket of hitting size $\nu(G)$ and so by definition of $\tau(G)$, $\tau(G) \geq \nu(G)$. We now show (b) cannot hold. For all s , V_s is a hitting set since \mathcal{H} is empty. Thus we have this restatement of (b): there is a vine decomposition (T, ℓ) such that $|V_s| < k$ for all $s \in T$, i.e. the vinewidth is at most $k - 1 = \nu(G) - 1$, a contradiction to the definition of $\nu(G)$. \square

So now we must prove Lemma 2.6. Our proof is based upon an interpretation by Reed in [16] of the bramble-treewidth duality theorem. Before proving Lemma 2.6, let us state two lemmas.

Lemma 2.7. *Given a thicket \mathcal{H} with hitting size h . Let X_1 and X_2 be hitting sets for \mathcal{H} . Then any separator for X_1 and X_2 contains at least h vertices.*

Proof. Take a set $H_i \in \mathcal{H}$. The set H_i is connected and intersects X_1 and X_2 at respectively v_1 and v_2 . Thus there is a path P_i in H_i between the v_1 and v_2 . A set S that separates X_1 and X_2 must disconnect this path. This applies for every set $H_i \in \mathcal{H}$, so S must hit every set in \mathcal{H} and is, thus, a hitting set. Therefore $|S| \geq h$. We remark that, in fact, there are h vertex disjoint paths from X_1 to X_2 in G . \square

Lemma 2.8. *Let $(\hat{T}, \hat{\ell})$ be a vine decomposition of G . Let X be a separator of G with a component C of $G \setminus X$. Suppose there is a node $t \in \hat{T}$ such that $\hat{V}_t \cap C = \emptyset$ and every separator for X and \hat{V}_t contains at least $|X|$ vertices. Then there is a vine decomposition $(T = \hat{T}, \ell)$ of $G[X \cup C]$ with*

(1) $V_t = X$. (2) $V_s \subseteq \hat{V}_s$, for each leaf $s \neq t$ in T . (3) $|V_s| \leq |\hat{V}_s|$, for all $s \in T$.

Lemma 2.8 allows us to restrict any tree decomposition on G to one on a subset of the vertices of G that contains a special label. This subset is the union of a cutset X and a component C of $G - X$. The special label is X .

Proof. Let A and B be two subsets of vertices. Menger's theorem ensures that the maximum number of internally vertex disjoint path from A to B is equal to the minimum size of a separator for A and B . Thus, there are $|X| = k$ vertex disjoint paths from \hat{V}_r to X . Let these paths be $\{P_1, P_2, \dots, P_k\}$ where the endpoint of P_i in X is x_i . Because \hat{V}_t is disjoint from C , it follows that all the P_i are also disjoint from C .

Now consider the vine tree \hat{T} . Let Q_i be the path in \hat{T} from from \hat{T}_{x_i} to t (not including vertices of \hat{T}_{x_i}). We claim the desired vine decomposition $(T = \hat{T}, \ell)$ on $G[X \cup C]$ is given by taking

$$V_s = (\hat{V}_s \cap (X \cup C)) \cup \{x_i : s \in Q_i\} = (\hat{V}_s \cap (X \cup C)) \cup \{x_i : P_i \cap \hat{V}_s \neq \emptyset\}$$

Note that $s \in Q_i$ if and only if P_i intersects \hat{V}_s because the node s separates \hat{V}_{x_i} from t in \hat{T} . Now, first let's verify that this is a vine decomposition on $G[X \cup C]$. Take any vertex $v \in C$. Then $T_v = \hat{T}_v$, a subtree of $T = \hat{T}$. On the other hand, take any vertex $x_i \in X$. Then $T_v = \hat{T}_v \cup Q_i$ which, again, is a subtree of $T = \hat{T}$. Furthermore, for any edge (u, v) in $G[X \cup C]$ we know that \hat{T}_u and \hat{T}_v either intersect or are adjacent in \hat{T} . Therefore, T_u and T_v either intersect or are adjacent in $T = \hat{T}$. Thus we have a vine decomposition.

Now let's show that the three required properties hold. (1) Each path P_i ends in \hat{V}_t . Thus, we have $x_i \in V_t$. Since \hat{V}_t is disjoint from C , we obtain $V_t = X$. (2) Take a leaf $s \neq t$ of T . The paths Q_i do not contain leaves and so $V_s \subseteq \hat{V}_s$. (3) Take a non-leaf s of T . If P_i intersects \hat{V}_s then we have $x_i \in V_s$. Suppose $x_i \notin \hat{V}_s$. But then there is a $y_i \in P_i \cap \hat{V}_s$ where $y_i \neq x_i$. By disjointness, y_i is not in P_j for any $j \neq i$. Moreover y_i is not in C since P_i is disjoint from C . Hence, $y_i \in \hat{V}_s \setminus V_s$. It follows that $|V_s| \leq |\hat{V}_s|$. \square

We now have all the tools we need to prove Lemma 2.6.

of Lemma 2.6. Assume (b) holds. So suppose there is a vine decomposition where any node $s \in T$ with $|V_s| \geq k$ is a leaf **and** does not provide a hitting set for \mathcal{H} . For every $\mathcal{H} \subseteq \mathcal{H}'$, the set V_s is not a hitting set for \mathcal{H} either. But, as we saw when proving that $\nu(G) \geq \tau(G)$ in the proof of Theorem 2.5, the vine decomposition must contain a node t that provides a hitting set V_t for \mathcal{H}' . But, by (ii), such a node has $|V_t| < k$. Thus, \mathcal{H} is a thicket with hitting size at most $k - 1$, and then (a) does not hold.

We must now show that at least one of (a) or (b) holds. We prove this by contradiction. In particular, take a counter-example \mathcal{H} with the fewest number of hitting sets of size at most $k - 1$.

Since (a) does not hold, every thicket \mathcal{H}' containing \mathcal{H} has hitting size at most $k - 1$. In particular, \mathcal{H} itself has at least one hitting set X of cardinality at most $k - 1$. By assumption, no vine decomposition exists with property (b). Hence, $X \neq V(G)$ otherwise the trivial vine decomposition satisfies (b). Now let $\{C_1, C_2, \dots, C_r\}$ be the connected components of $G \setminus X$. We will find vine decompositions (T^i, ℓ^i) of $G[X \cup C_i]$, for $1 \leq i \leq r$, that satisfy

- (1) There is a node t_i of T^i with $V_{t_i}^i = X$.
- (2) Any node $s \in T^i$ with $|V_s^i| \geq k$ is a leaf and V_s is not a hitting set for \mathcal{H} .

Since the $V_{t_i}^i$ are identical we may combine the T^i -s by merging together all t_i into a single node whose label is the union of labels of nodes we merged. Observe that this will create a vine decomposition T for G satisfying (i) and (ii). The theorem then follows. So we need to show that (1) and (2) hold for $G[X \cup C_i]$.

First suppose C_i does not intersect some thicket element H_i in \mathcal{H} . Then we can define T^i to be a tree with two nodes s and t_i , where $V_{t_i}^i = X$ and $V_s^i = C_i$. Trivially this is a valid vine decomposition. Clearly, $|V_{t_i}^i| = |X| < k$ thus (1) holds. Now V_s^i is not a hitting set for \mathcal{H} , thus (2) holds.

Therefore, we may assume that C_i is a hitting set for \mathcal{H} . By the connectedness of C_i , we have that $\hat{\mathcal{H}} = \mathcal{H} \cup \{C_i\}$ is also a thicket. Since (a) does not hold for \mathcal{H} , we know that $\hat{\mathcal{H}}$ has a hitting set of size at most $k - 1$. Moreover, X is not a hitting set for $\hat{\mathcal{H}}$ since X is disjoint from C_i . Thus $\hat{\mathcal{H}}$ has fewer hitting sets of size at most $k - 1$ than \mathcal{H} . Consequently, $\hat{\mathcal{H}}$ is not a counterexample. But (a) cannot hold for $\hat{\mathcal{H}}$ otherwise it holds for $\mathcal{H} \subset \hat{\mathcal{H}}$, a contradiction. Hence (b) holds for $\hat{\mathcal{H}}$.

So take a vine decomposition $(\hat{T}, \hat{\ell})$ of G satisfying (b) for $\hat{\mathcal{H}}$. There must be a leaf $t \in \hat{T}$ with $|\hat{V}_t| \geq k$ such that \hat{V}_t is a hitting set for \mathcal{H} , and $\hat{V}_t \cap C_i = \emptyset$. If not, $(\hat{T}, \hat{\ell})$ is a vine decomposition of G satisfying (b) with respect to \mathcal{H} , contradicting the definition of G . We will transform \hat{T} into a vine decomposition for $G[X \cup C_i]$ satisfying both (1) and (2). Recall X is a minimum hitting set for \mathcal{H} . Furthermore, \hat{V}_t is a hitting set for \mathcal{H} . Thus, by Lemma 2.7, every separator for X and \hat{V}_t contains at least $|X|$ vertices. Therefore, we may apply Lemma 2.8 with $C = C_i, t = t$ to give a vine decomposition of $G[X \cup C_i]$. This vine decomposition satisfies (1) and (2). To see (1), note that t_i is a leaf with $V_{t_i}^i = X$. Now $|V_s^i| \leq |\hat{V}_s^i|$ and $|\hat{V}_s^i| \geq k$ only if s is a leaf. Thus, s must be a leaf if $|V_s^i| \geq k$ and, hence, $V_s^i \subseteq \hat{V}_s^i$. But \hat{V}_s^i is not a hitting set for $\hat{\mathcal{H}}$. Consequently, either \hat{V}_s^i is not a hitting set for \mathcal{H} or $\hat{V}_s^i \cap C_i = \emptyset$. In the latter case, we then have that $\hat{V}_s^i \subseteq X$ and so $|\hat{V}_s^i| \leq |X| \leq k - 1$. Thus (2) holds. \square

3 The Packing-Covering Ratio

In this section, we prove that the packing-covering ratio is given exactly by the thicket number of the interaction graph (Theorem 1.1). To prove this, let's consider the lower and upper bounds separately.

Theorem 3.1. *For any interaction graph G , the packing-covering ratio satisfies*

$$\tau(G) \leq_{\exists} \frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})}$$

Proof. Given G , let $\mathcal{H} = \{H_1, H_2, \dots, H_p\}$ be a thicket with maximum hitting size $\tau(G)$. We define a game \mathcal{G} using the following valuation function:

$$v(S) = \begin{cases} 1 & \text{if } S \in \mathcal{H} \\ 0 & \text{otherwise} \end{cases}$$

Recall each set $H_i \in \mathcal{H}$ is connected. Thus v is a valid valuation function for a coalition game over the interaction graph G . Furthermore, the sets $\{H_1, H_2, \dots, H_p\}$ are pairwise-intersecting. Thus any partition \mathcal{S} of the agents can include at most one set from \mathcal{H} . So $\rho(\mathcal{G}) = 1$. On the other hand, with integral payoffs, we must provide a dollar to at least one agent in each coalition in \mathcal{H} . The cheapest way to do this is to give a dollar to each agent in a minimum hitting set for \mathcal{H} . Thus $\kappa(\mathcal{G}) = \tau(\mathcal{H}) = \tau(G)$ and the packing-covering ratio is exactly $\tau(G)$. \square

Theorem 3.2. *For any interaction graph G , the packing-covering ratio satisfies*

$$\frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} \tau(G)$$

Proof. Take any game \mathcal{G} with valuation function v over an interaction graph G . Let $T = (N, L)$ be a vine decomposition of G . We may assume each label in the vine decomposition has size $\tau(G)$; if not, simply add vertices from the labels of adjacent nodes. Root the (vine) tree T at an arbitrary node r . In turn, we may now consider each subtree of T to be rooted at its (unique) node closest to the root r . We claim that, for each coalition Q , the nodes $T(Q) = \bigcup_{v \in Q} T_v$ induce a connected graph in T . Indeed, by viability of the coalition, we

know $G[Q]$ is connected. Thus, for every edge (u, v) in the subtree $G[Q]$, we have that $T_u \cup T_v$ is connected, by definition of a vine decomposition. Then, since connectivity is transitive, $T(Q)$ induces a subtree in T .

Hence, we may define the *root of a coalition*, t_Q , to be the root of $T(Q)$. We are ready now to describe a payment allocation \mathbf{x} that proves the theorem. To simplify the analysis, instead of allocating values to agents directly, we have an allocation $x_{i,t}$ for each agent i and each node t in T_v . The total allocation for agent i is then simply $x_i = \sum_{t \in T_i} x_{i,t}$. We work bottom-up from the leaves to the root, and allocate to all vertices in a label $\ell(t)$ in turn. At a node t , for each coalition Q whose root is t , we compute the total amount $x(Q, t)$ allocated to Q in descendants of t and the residual value $r(Q, t) = \max(v(Q) - x(Q, t), 0)$ we still need to add to Q to create a valid allocation. If Q_t^* is the coalition of maximum residual value, we set $x_{v,t} = r(Q_t^*, t)$ for every agent i in $\ell(t)$.

By the choice of Q_t^* , we can conclude (simply by looking only at allocations for t and its descendants) that $x(Q) \geq v(Q)$, for every coalition Q with root t . Furthermore, if $r(Q_t^*)$ is positive then $x(Q^*) = v(Q^*)$. Thus, since every coalition has a root, $x(Q) \geq v(Q)$ for every feasible coalition Q . Now let's bound the total cost of the allocation. By construction, the cost of our allocation is $\sum_t r(Q_t^*, t) \cdot |\ell(t)|$. Therefore, since all labels have size $|\ell(t)| = \tau(G)$, we have that

$$\kappa(\mathcal{G}) \leq \tau(G) \cdot \sum_t r(Q_t^*, t)$$

Therefore, it suffices to prove that

$$\sum_t r(Q_t^*, t) \leq \rho(\mathcal{G})$$

To do this we construct an integral packing \mathcal{Q} of coalitions as follows. Consider the nodes of T from root to leaves (in a postorder traversal). Initially no node is marked as deleted. At a node t of T , if t is not marked as deleted and the residual $r(Q_t^*, t)$ is positive, add Q_t^* to \mathcal{Q} and mark all nodes in $T(Q_t^*)$ as deleted. Otherwise, mark t as deleted. We bound the packing value of \mathcal{Q} inductively using a potential function. This potential is defined as the total allocation in remaining nodes (*i.e.*, nodes that are not marked as deleted). Initially this potential is $\tau(G) \cdot \sum_t r(Q_t^*, t)$. When we add an element to \mathcal{Q} the potential drops by $\tau(G) \cdot x(Q_t^*) = \tau(G) \cdot v(Q_t^*)$, whilst the value of the packing \mathcal{Q} increases by $v(Q_t^*)$. By the end, every node is marked as deleted. Consequently, our potential function is zero and, so, the value of our packing has increased by exactly $\frac{1}{\tau(G)} \cdot \tau(G) \cdot \sum_t r(Q_t^*, t) = \sum_t r(Q_t^*, t)$. Thus \mathcal{Q} has the desired value. It remains to verify the coalitions in \mathcal{Q} are disjoint. So take any two coalitions, say $Q_{t_1}^*$ and $Q_{t_2}^*$. If t_1 and t_2 do not form an ancestor-descendant pair in T then $Q_{t_1}^*$ and $Q_{t_2}^*$ are disjoint since the nodes containing vertices of $Q_{t_1}^*$ are in the tree rooted in t_1 while those of $Q_{t_2}^*$ are in the tree rooted in t_2 . So, without loss of generality, assume that t_2 is an ancestor of t_1 . It follows that $t_1 \notin T(Q_{t_2}^*)$, otherwise $Q_{t_1}^*$ would not have been selected. Thus $Q_{t_1}^*$ and $Q_{t_2}^*$ are again disjoint. \square

We conclude this section with a brief discussion on computational implications. Computing the treewidth of a graph G is an NP-hard problem [1]. Whilst no formal proof is given here, the NP-completeness arguments can be extended to vine width. The existence of a constant-factor approximation algorithm for the treewidth running in polynomial time

remains an important open problem. There does, however, exist a constant-factor approximation algorithm for the treewidth in FPT time parameterized by the treewidth [3]. Since, by Theorem 2.3, the vinewidth is within a factor two of the treewidth, this provides a constant-factor approximation algorithm for the vinewidth of the graph.

4 VC-Dimension

A ubiquitous measure of the complexity of a set-family is its VC-dimension [19]. As we have a set-family derived from the interaction graph G , it is natural to ask whether we can relate the packing-covering ratio to the VC-dimension. We explore this question in this section.

First, recall the definition of VC-dimension. Given a ground set I and a collection $\mathcal{R} = \{R_1, \dots, R_m\}$ of subsets of I , we say that $X \subseteq I$ is *shattered* by \mathcal{R} if, for all $Y \subseteq X$, there exists some $R_j \in \mathcal{R}$ such that $Y = X \cap R_j$. The *VC-dimension* of (I, \mathcal{R}) is then the maximum cardinality of a shattered set.

Interestingly, for any set-family, in a simple game⁶ the primal integrality gap can also be upper bounded by the VC-dimension. Specifically,

Theorem 4.1. *Haussler, Welzl [9] Let \mathcal{R} be a set-family with VC-dimension d . Then the primal integrality gap of any simple game \mathcal{G} whose viable coalitions are \mathcal{R} satisfies*

$$\frac{\kappa(\mathcal{G})}{\kappa^f(\mathcal{G})} \leq d \cdot \log d$$

We can strengthen this result when the family of coalitions is induced by an interaction graph $G = (V, E)$. In particular, let the *graphical set family* $\mathcal{S} = \{S_1, \dots, S_r\}$ of G be the set of *all* connected induced subgraphs of G . We then define the *VC-dimension* of the graph G to be the VC-dimension of the graphical set family of G .

Theorem 4.2. *Let G be a graph with VC-dimension d . Then, restricting to simple games, the packing-covering ratio satisfies*

$$\frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} d + 1$$

Proof. Let \mathcal{G} be a coalition game over interaction graph G . Let \mathcal{R} be the set of coalitions of \mathcal{G} . Note that \mathcal{R} is a subset for the graphical set family \mathcal{S} of G . Let X be a minimum hitting set of \mathcal{R} . We want to show that $\rho(\mathcal{G}) \geq \frac{1}{d+1} \cdot |X|$. Now by the minimality of X there is a justifying coalition $R_x \in \mathcal{R}$ for each vertex $x \in X$; specifically, there exists a coalition R_x such that $R_x \cap X = \{x\}$. Let $\mathcal{J} \subseteq \mathcal{R}$ be the set of justifying coalitions. So $|\mathcal{J}| = |X|$.

First, assume there is a set $R^* \in \mathcal{J}$ that intersects k of the other justifying coalitions. Let $\mathcal{J}' = \{R_{x_1}, \dots, R_{x_k}\}$ be the collection of k justifying coalitions that intersect R^* . We now show $Y = \{x_1, \dots, x_k\}$ is shattered for the graphical set family \mathcal{S} . We claim that for any $Y' \subseteq Y$, the set $R' = R^* \cup \bigcup_{x \in Y'} R_x$ intersects Y in exactly Y' . The set R' is connected since every R_x is connected and, by construction, each of them intersects the connected set R^* . Thus R' is in the graphical set family \mathcal{S} and is connected. Since $R^* \cap Y$ is empty, we have $R' \cap Y = Y'$, as desired. So Y can be shattered (for the graphical set family \mathcal{S}), and thus $k \leq d$ since the VC-dimension of G is at most d .

⁶A *simple game* is a coalition game where the value of every coalition is either 0 or 1.

Consequently, we may assume that every set \mathcal{J} intersects at most d of the other justifying coalitions. Then we can easily obtain a disjoint packing of $\frac{1}{d+1} \cdot |X|$ coalitions in \mathcal{J} . Simply select any coalition R in \mathcal{J} ; then remove R and the (at most) d coalitions it intersects from \mathcal{J} , and recurse. The theorem follows. \square

But, we have seen that the thicket number gives the packing-covering ratio exactly. Thus we should be able to upper bound the thicket number of any graph by a function of the VC-dimension. Indeed this is the case.

Theorem 4.3. *Let G be a graph of with VC-dimension d . Then the thicket number $\tau(G)$ is at most d .*

Proof. Let $\mathcal{H} = \{H_1, \dots, H_p\}$ be a thicket with hitting number $\tau(G)$, and let X be a minimum hitting set for \mathcal{H} . We claim that X is shattered. To see this take any $x \in X$. There exists some justifying set $H_x \in \mathcal{H}$ such that $H_x \cap X = \{x\}$, otherwise X is not minimal. Now take any $Y \subseteq X$. Without loss of generality, let $Y = \{x_1, x_2, \dots, x_r\}$. Now each H_{x_i} induces a connected graph. Moreover, the sets H_{x_i} pairwise-intersect. Thus $W = \cup_{i=1}^r H_{x_i}$ also induces a connected graph. Thus W is in the graphical set family \mathcal{S} . Since $H_{x_1} \cap X = \{x_1\}$ we have that

$$X \cap W = X \cap (\cup_{i=1}^r H_{x_i}) = \cup_{i=1}^r (X \cap H_{x_i}) = \cup_{i=1}^r x_i = Y$$

Thus X is shattered. \square

Combining Theorem 4.3 with Theorem 1.1, we obtain the following strengthening of Theorem 4.2, which also applies to non-simple games.

Corollary 4.4. *Let G be a graph with VC-dimension d . Then the packing-covering ratio satisfies*

$$\frac{\kappa(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} d$$

Of course, Corollary 4.4 must give a weaker bound than Theorem 1.1. Indeed, the VC-dimension of a graph G can be arbitrarily larger than its thicket number. To see this, consider a star graph with n edges. Since a star is a tree, the vinewidth of G is equal to 1 by Theorem 2.3. But the VC-dimension of the star is n since any subset of leaves of the star can be shattered using the graphical set family of G .

5 The Dual Integrality Gap

We now consider the integrality gaps of the primal and dual linear programs. Clearly, these gaps are at most the packing-covering ratio, and thus at most the thicket number $\tau(G)$. It is conceivable, however, that the integrality gaps could be much smaller than the thicket number. In Section 6 we will consider the primal integrality gap. In this section, we examine the dual integrality gap. Recall that this gap determines the multiplicative least-core and measures the relative cost of stability.

Before quantifying the dual integrality gap in graphical coalition games, it is informative to give practical interpretations of the integral and fractional packing numbers. Suppose there is a one unit interval of time, and at any point in time an agent can choose to work for any coalition it belongs to. Thus, each agent partitions the interval into sub-intervals

associated with assorted coalitions. Now consider two different determinants for whether a coalition is productive. First, suppose that a coalition can only function if its members meet together – thus the coalition can generate wealth only if its members are working for it simultaneously. In this setting, at any point in time t , the functioning coalitions are disjoint. It then follows, by superadditivity, that it is best if the grand coalition meets at time t . Since this argument holds for any t , the optimal solution is that the grand coalition meets for the entire unit of time (*i.e.*, each agent contributes all its time to the grand coalition). Thus, we obtain the integral packing solution $\rho(\mathcal{G})$. Second, instead suppose that it is not necessary for a coalition to convene simultaneously. The productivity of a coalition is then determined by the minimum time contribution of one of its members. So coalition members may work at different times and, in this setting, it is possible that more wealth can be generated if the agents fractionally allocate their time amongst multiple coalitions. Indeed the optimal solution now is the fractional packing solution $\rho^f(\mathcal{G})$.

The main result of this section is the following:

Theorem 5.1. *There exist $c, \delta > 0$ such that for any interaction graph G , the dual integrality gap satisfies:*

$$c \cdot \tau(G)^\delta \leq_{\exists} \frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})} \leq_{\forall} \tau(G)$$

The upper bound follows from Theorem 1.1. Thus it remains for us to prove the lower bound. To do this, we will apply some important results from graph minor theory. A graph H is a *minor* of a connected graph G if the vertices of G can be partitioned into $|V(H)|$ non-empty connected subgraphs such that if (u, v) is an edge of H then there exists an edge of G with one endpoint in the subgraph corresponding to u and one endpoint in the subgraph corresponding to v . Robertson and Seymour [17] proved that every graph of treewidth k admits a grid minor of size $f(k)$, that is, the $f(k) \times f(k)$ grid is a minor of every graph of treewidth at least k . Their bound f was improved several times over the years, but only recently was a polynomial bound obtained by Chekuri and Chuzhoy [4]. They prove that there exist c' and $\delta > 0$ such that every graph G has a grid minor of size at least $c' \cdot \omega(G)^\delta$. Since $\tau(G) \leq 2\omega(G)$ by Theorem 2.3, this implies that there exist c and δ such that every graph has a grid minor of size $c \cdot \tau(G)^\delta$. Now, before completing the proof of Theorem 5.1, let us learn more about the dual integrality gap of grids.

Lemma 5.2. *The dual integrality gap of an $n \times n$ grid R_n satisfies*

$$\frac{1}{2}\tau(R_n) \leq_{\exists} \frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})}$$

Proof. Take the grid R_n and define a thicket $\mathcal{H} = \{H_1, H_2, \dots, H_n\}$ as follows. The set H_i is the set of vertices in the i th row or in the i th column. Thus each H_i is viable, and H_i and H_j intersect at two vertices in the grid – namely, the vertices with grid coordinates (i, j) and (j, i) . We create a simple game \mathcal{G} over the grid R_n by assigning value one to those coalitions in \mathcal{H} . Clearly $\rho(\mathcal{G}) = 1$ as the sets in \mathcal{H} are pairwise intersecting. On the other hand, $\rho(\mathcal{G}) \geq \frac{1}{2}n = \frac{1}{2}\tau(R_n)$. This is easy to see because we may assign a fractional value of $\frac{1}{2}$ to each set in \mathcal{H} . Every grid vertex is in exactly two sets of \mathcal{H} , so this is a valid fractional packing. \square

We are now ready to complete the proof of Theorem 5.1.

of Theorem 5.1. Let G be a graph of vinewidth $\nu(G) = \tau(G)$. We may assume that G is connected. By [4], there exist c and δ such that G admits a grid minor R_k of size at least $k = c \cdot \tau(G)^\delta$. We wish to apply Lemma 5.2 to show the existence of a game with dual integrality gap at least $\frac{1}{2}k$. We define the thicket $\mathcal{H} = \{H_1, H_2, \dots, H_k\}$ as before, except now each node in the grid minor R_k corresponds to a connected subgraph of the original graph G . Formally, denote by $u_{i,j}$ the vertices of the grid of size $k \times k$ (where i refers to the row and j to the column of the vertex $u_{i,j}$) and by $X_{i,j}$ the connected subset of G which is assigned to $u_{i,j}$. For every $i \leq k$, let $H_i = \bigcup_{j=1}^k X_{i,j} \cup \bigcup_{j=1}^k X_{j,i}$. The coalition H_i is viable by the definition of a minor. Again, we take the simple game \mathcal{G} where only the coalitions in \mathcal{H} are given value 1. For every $i \neq j$, the sets H_i and H_j intersect since both $Y_i \cap Y_j = X_{i,j} \cup X_{j,i}$. Thus $\rho(\mathcal{G}) = 1$ and $\rho^f(\mathcal{G}) \geq \frac{k}{2}$. \square

Of course, any subsequent improvement in the polynomial function of [4] will give an improved lower bound for 5.1. Nonetheless, this grid-minor method *cannot* provide a linear lower bound. We prove this by considering the class of clique graphs.

Lemma 5.3. *The dual integrality gap of the clique K_n satisfies*

$$\sqrt{\frac{\tau(K_n)}{2}} - 1 \leq \exists \frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})} \leq \forall \sqrt{8} \cdot \sqrt{\tau(K_n)}$$

To prove this lemma, we will use the same idea as Lemma 5.1 for the lower bound and the following lemma for the upper bound.

Lemma 5.4. *For any interaction graph G on n vertices, the dual-integrality gap satisfies*

$$\frac{\rho^f(\mathcal{G})}{\rho(\mathcal{G})} \leq \forall 2\sqrt{n}$$

Proof. Take any game \mathcal{G} on G . Since $\rho^f(\mathcal{G}) = \kappa^f(\mathcal{G})$ by strong duality, it suffices to show that $\kappa^f(\mathcal{G}) \leq 2\sqrt{n} \cdot \rho(\mathcal{G})$. Call a coalition *large* if it contains at least \sqrt{n} agents and *small* otherwise. Now greedily select a packing of small coalitions $\{S_1, S_2, \dots, S_k\}$ as follows. Let S_1 be the small coalition of maximum value. Then recursively, let S_{i+1} be the small coalition of maximum value that is disjoint from $\{S_1, \dots, S_i\}$. Allocate $v(S_j)$ to every agent in S_j . In addition, we allocate $\frac{1}{\sqrt{n}} \cdot v^*$ to every agent, where $v^* = \max_S v(S)$. We claim that this allocation \mathbf{x} is a feasible solution to the primal. To see this take any coalition S . If S is large then $\sum_{i \in S} x_i = \frac{|S|}{\sqrt{n}} \cdot v^* \geq v(S)$, because $|S| \geq \sqrt{n}$. Suppose S is small. If S was selected in the greedy packing then $\sum_{i \in S} x_i > \sum_{i \in S} v(S_i) = |S_i| \cdot v(S) \geq v(S)$. Otherwise, let S_j be the lowest index set in the packing that intersects S . By the greedy selection mechanism we then have $v(S_j) \geq v(S)$. Therefore, $\sum_{i \in S} x_i > v(S_j) \geq v(S)$ as each agent in $S \cap S_j$ is allocated at least $v(S_j)$ (and there is at least one such agent). So we have $\kappa^f(\mathcal{G}) \leq \sum_{i \in I} x_i$ and, furthermore,

$$\begin{aligned} \sum_{i \in I} x_i &= n \cdot \frac{1}{\sqrt{n}} v^* + \sum_{j=1}^k |S_j| \cdot v(S_j) \\ &\leq \sqrt{n} \cdot v^* + \sqrt{n} \cdot \sum_{j=1}^k v(S_j) \\ &\leq \sqrt{n} \cdot \rho(\mathcal{G}) + \sqrt{n} \cdot \rho(\mathcal{G}) \\ &\leq 2\sqrt{n} \cdot \rho(\mathcal{G}) \end{aligned}$$

The lemma follows. \square

of Lemma 5.3. First recall that the vinewidth of the clique K_n is $\lceil \frac{n}{2} \rceil$. Moreover, every coalition over interaction graph K_n induces a connected subgraph and then is viable. For the lower bound, we label the vertices of K_n by coordinates i, j that vary between 0 and \sqrt{n} (i.e., we place the vertices in a grid). Let H_i consist of those vertices in the i th row or i th column. Now, consider the game where each H_i is a coalition of value 1. Since any two element of H_i intersect, $\rho = 1$. Since we may fractionally choose $\frac{1}{2}$ of each H_i , $\rho^f \geq \frac{1}{2}\sqrt{n} = \frac{1}{2}\sqrt{n} \cdot \rho$. This proves the lower bound.

Since $n \leq 2\tau(K_n)$, Lemma 5.4 tells us the dual integrality gap is at most $2\sqrt{n} \leq \sqrt{8\tau(G)}$. This proves the upper bound. \square

6 The Primal Integrality Gap

To conclude, we consider the primal integrality gap. This measures the maximum ratio in the cost between paying the agents in integral amounts and paying in fractional amounts. Our first result is that the thicket number does quantify the primal integrality gap to within a constant factor, namely Theorem 1.2. The upper bound follows from Theorem 1.1, so it suffices to show the lower bound. (Due to space considerations, all the proofs for this section are deferred to the appendix.)

Theorem 6.1. *For any interaction graph G , the primal integrality gap satisfies:*

$$\frac{1}{4}\tau(G) \leq \exists \frac{\kappa(G)}{\kappa^f(G)}$$

Proof. Take a graph $G = (V, E)$ with thicket number $\tau(G)$ and let $\mathcal{H} = \{H_1, \dots, H_p\}$ be a thicket with a minimum hitting set of size $\tau(G)$. Let X be a minimum cardinality hitting set for \mathcal{H} . Now consider the following coalition game \mathcal{G} over the interaction graph G . For a coalition $S \subseteq V$, we set $v(S) = 1$ if there exists a family $\mathcal{H}' \subseteq \mathcal{H}$ such that $S = \cup_{H \in \mathcal{H}'} H$ and $|S \cap X| \geq \lceil \frac{1}{2}\tau(G) \rceil$. Since, the sets in the thicket are connected and pairwise intersecting we have that S is viable; thus v is a valid valuation function.

The linear program has a solution with value $\kappa^f(\mathcal{G}) \leq 2$. To see this, consider the solution where each agent of X is allocated $\frac{2}{\tau(G)}$ and all other agents are allocated 0. Because each coalition with value 1 has at least $\lceil \frac{1}{2}\tau(G) \rceil$ members in X , this is a feasible fractional solution.

To prove the primal integrality gap is at least $\frac{1}{4}\tau(G)$, we will now show that the optimal integral solution to the primal has value at least $\lfloor \frac{1}{2}\tau(G) \rfloor + 1 \geq \frac{1}{2}\tau(G)$. Suppose not. Then the set \mathcal{S} of coalitions with value 1 must have a hitting set Y of cardinality at most $\lfloor \frac{1}{2}\tau(G) \rfloor$. Now let $\hat{\mathcal{H}}$ be the sets in the thicket \mathcal{H} that are disjoint from Y . Then the minimum size of a hitting set Z for the family $\hat{\mathcal{H}}$ is at least $\lceil \frac{1}{2}\tau(G) \rceil$. Otherwise $Y \cup Z$ is a hitting set of \mathcal{H} of size less than $\tau(G)$. But then $\hat{S} = \cup_{H \in \hat{\mathcal{H}}} H$ satisfies $|\hat{S} \cap X| \geq \lceil \frac{1}{2}\tau(G) \rceil$. Thus, by definition, \hat{S} is a coalition in \mathcal{S} . However Y is a hitting set of \mathcal{S} and $Y \cap \hat{S} = \emptyset$, a contradiction. \square

So the primal integrality gap (for the worst game) is within a factor 4 for any pair of interaction graphs with the same thicket number. Recall that the packing-covering ratio is the same (for the worst game) for every pair of interaction graphs with the same thicket number by Theorem 1.1. Is this also the case for the primal integrality gap? The answer is

no. There are graphs whose primal integrality gaps differ. In particular we will show that the integrality gap for a clique is equal to $\frac{1}{2}\tau(G)$, up to an additive constant, whereas the class of graphs that are “powers of a path” have integrality gaps that tend to $\tau(G)$. The latter inequality ensures that the upper bound of Theorem 1.2 cannot be improved in general using thicket number.

Lemma 6.2. *Let $G = K_n$ be a clique on n vertices. Then the primal integrality gap of K_n satisfies*

$$\frac{1}{2}\tau(K_n) \leq \exists \frac{\kappa(\mathcal{G})}{\kappa^f(\mathcal{G})} \leq \forall \frac{1}{2}\tau(K_n) + 1$$

Proof. Recall that the vinewidth and thicket number of the clique K_n equal $\lceil \frac{1}{2}n \rceil$. We will prove that, for any coalition game \mathcal{G} on the clique, the integrality gap is at most $\frac{1}{4}n + 1 \leq \frac{1}{2}\tau(G) + 1$. Observe that every coalition S over interaction graph K_n induces a connected subgraph and is, thus, viable. So, in what follows, we need not verify the viability of any coalition. Let \mathcal{S} be the set of coalitions in the game \mathcal{G} with value 1.

Let \mathbf{x} be an optimal fractional solution to the primal. Thus, for any coalition $S \in \mathcal{S}$, we have $\sum_{i \in S} x_i \geq 1$. So $\kappa^f(\mathcal{G}) = \sum_{i \in V} x_i$. Now write $\kappa(\mathcal{G})$ as $(1 - \alpha)n + 1$ by choosing the appropriate $\alpha > 0$ (if there does not exist such an α , the result is proved). Thus, there is no hitting set for \mathcal{S} with cardinality at most $(1 - \alpha)n$. We claim that $\kappa^f(\mathcal{G}) \geq 1/\alpha$. To see, this take any set X of $(1 - \alpha)n$ agents. Since X is not a hitting set, there exists a coalition $S_X \in \mathcal{S}$ contained in its complement $\bar{X} = V \setminus X$. Thus $\sum_{i \in \bar{X}} x_i \geq 1$. This holds for *every* set of agents X of cardinality $(1 - \alpha)n$ and there are $\binom{n}{(1-\alpha)n} = \binom{n}{\alpha n}$ such sets. Each agent appears in a α -fraction of the complements of these sets. Consequently, $\alpha \cdot \binom{n}{\alpha n} \cdot \sum_{i \in V} x_i \geq \binom{n}{\alpha n}$. Therefore, $\kappa^f(\mathcal{G}) = \sum_{i \in V} x_i \geq 1/\alpha$ as claimed. As $\kappa(\mathcal{G}) = (1 - \alpha)n + 1$, the primal integrality gap is at most

$$\alpha(1 - \alpha)n + \alpha \leq 2\alpha(1 - \alpha)\tau(K_n) + \alpha \leq 2\alpha(1 - \alpha)\tau(K_n) + 1.$$

This is maximized when $\alpha = \frac{1}{2}$. Thus

$$\frac{\kappa(G)}{\kappa^f(G)} \leq \forall \frac{1}{2}\tau(G) + 1$$

This upper bound is tight, we have a matching lower bound

$$\frac{1}{2}\tau(G) \leq \exists \frac{\kappa(G)}{\kappa^f(G)}$$

To see this, consider the game where any coalition of size $\lceil \frac{n}{2} \rceil$ has value 1 and any other coalition has value 0. Then we have $\kappa^f \leq 2$ because allocating $\frac{2}{n}$ to each agent, gives a feasible fractional solution. On the other hand, $\kappa = \lfloor \frac{n}{2} \rfloor + 1$, otherwise some coalition will block the allocation. \square

There are, however, graphs for which the upper bound of $\tau(G)$ on the primal integrality gap is obtained. Specifically, this bound is obtained for power graphs of a path. Let P be a path on n vertices. The r -th power of P , denoted by $G = P^r$, is formed by connecting any pair of vertices whose distance is at most r in P .

Theorem 6.3. *Let $G = P^r$ be the r -th power of a path on n vertices where $n \geq 3r$. Then $\tau(G) = r$ and, provided $n \geq k^2(r+1)$, the primal integrality gap satisfies*

$$\left(1 - \frac{2}{k}\right) \cdot \tau(G) \leq \exists \frac{\kappa(\mathcal{G})}{\kappa^f(\mathcal{G})}$$

Theorem 6.3 ensures that the constant 1 in the upper bound of Theorem 1.2 cannot be improved. However, as observed in Lemma 6.2, this upper bound cannot be reached for every graph as there are graphs where the primal integrality gap is upper bounded by around half the vinewidth.

Proof. To show that P^r has thicket number $\tau(P^r) = r$, let us first prove that $\tau(P^r) \leq r$. We create a vine decomposition $T = (N, L)$ as follows. Label the vertices in order along the path as $\{1, 2, \dots, n\}$. Then we let T be a path on $\lceil \frac{n}{r} \rceil$ nodes, where each node corresponds to a set of vertices of the form $\{qr + 1, qr + 2, \dots, qr + r = (q+1)r\}$, where $0 \leq q \leq \lceil \frac{n}{r} \rceil - 1$. Thus $|T_v| = 1$ and so T_v is trivially connected, for each vertex $v \in P^r$. Moreover, for each edge (u, v) in $G = P^r$, either u and v are in the same node of T or are in adjacent nodes. So this is a vine decomposition and $\tau(P^r) = \nu(P^r) \leq r$.

Let us now show that $\tau(P^r) \geq r$. Consider the restriction of P^r to the first $3r$ vertices, and construct a thicket \mathcal{H} as follows. Let $A = \{1, \dots, r\}$, $B = \{r+1, \dots, 2r\}$ and $C = \{2r+1, \dots, 3r\}$. We put a set H in \mathcal{H} if it induces a connected subgraph, it contains at least one vertex from each of A, B and C , and it contains more than $r/2$ elements from at least two of A, B, C . Note that any pair of sets in \mathcal{H} intersect. Assume by contradiction that X is a hitting of size less than r . Since $|X| < r$, none of A, B, C are completely contained in X . Furthermore, X contains less than half of the vertices from at least two of these three sets. Thus the complement \bar{X} of X contains at least one vertex in all of A, B, C and more than half of the vertices of two of these three sets. Moreover \bar{X} is connected since X does not contain r consecutive vertices of P . So $\bar{X} \in \mathcal{H}$, contradicting the fact that X is a hitting set.

We now show the lower bound by constructing a coalition game \mathcal{G} on P^r . For a coalition $S \subseteq V$, we set $v(S) = 1$ if S is a connected subgraph of P^r of cardinality at least $\frac{n}{k}$, and set $v(S) = 0$ otherwise. We have that $\kappa^f(\mathcal{G}) \leq k$, since allocating $\frac{k}{n}$ to each agent gives a feasible fractional solution. We claim that $\kappa(\mathcal{G}) \geq r(k-1)$.

Let \mathcal{S} be the coalitions of value 1. Now any hitting set X for \mathcal{S} must contain at least r agents from amongst $\{1, \dots, \lceil \frac{n}{k} \rceil + r\}$; otherwise, there is a connected subgraph of cardinality $\lceil \frac{n}{k} \rceil$ missed by X . Similarly, X must contain r agents amongst $\{\lceil \frac{n}{k} \rceil + r + 1, \dots, 2(\lceil \frac{n}{k} \rceil + r)\}$ and, in general, r agents amongst $\{a(\lceil \frac{n}{k} \rceil + r) + 1, \dots, (a+1)(\lceil \frac{n}{k} \rceil + r)\}$ for $0 \leq a \leq \lfloor \frac{n}{\lceil \frac{n}{k} \rceil + r} \rfloor - 1$. This implies that the number of agents in X is at least

$$\left\lfloor \frac{n}{\lceil \frac{n}{k} \rceil + r} \right\rfloor \cdot r \geq (k-2)r$$

Here the inequality holds if we select k such that $n \geq k^2(r+1)$. Hence, the primal integrality gap is at least $(1 - \frac{2}{k}) \cdot \tau(G)$. \square

Acknowledgements. The third author is grateful to Reshef Meir for introducing him to this problem.

References

- [1] S. Arnborg, D. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM J. Algebraic Discrete Methods*, (8(2)):277–284, April 1987.
- [2] Y. Bachrach, E. Elkind, R. Meir, D. Pasechnik, M. Zuckerman, J. Rothe, and J. Rosenschein. The cost of stability in coalitional games. In *Proceedings of Second Symposium on Algorithmic Game Theory (SAGT)*, pages 122–134, 2009.
- [3] H. Bodlaender. Dynamic programming on graphs with bounded treewidth. In *Proceedings of Fifteenth International Colloquium on Automata, Languages and Programming (ICALP)*, pages 105–118, 1988.
- [4] C. Chekuri and J. Chuzhoy. Polynomial bounds for the grid-minor theorem. In *Proceedings of Forty-Sixth Symposium on Theory of Computing (STOC)*, pages 60–69, 2014.
- [5] F. Edgeworth. *Mathematical psychics: an essay on the mathematics to the moral sciences*. London: Kegan Paul, 1881.
- [6] R. Gilles. *The Cooperative Game Theory of Networks and Hierarchies*. Springer, 2010.
- [7] D. Gillies. *Some Theorems on n -Person Games*. PhD thesis, Princeton, 1953.
- [8] D. Gillies. *Contributions to the Theory of Games*, volume 4, pages 47–85. Princeton University Press, 1959.
- [9] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. In *Symposium on Computational Geometry*, pages 61–71, 1986.
- [10] W. Hildenbrand and A. Kirman. *Equilibrium Analysis: Variations on Themes by Edgeworth and Walras*. Elsevier, 1988.
- [11] M. Le Breton, G. Owen, and S. Weber. Strongly balance cooperative games. *International Journal of Game Theory*, (20):419–427, 1992.
- [12] R. Meir, J. Rosenschein, and E. Malizia. Subsidies, stability, and restricted cooperation in coalitional games. In *Proceedings of Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 301–306, 2011.
- [13] R. Meir, Y. Zick, E. Elkind, and J. Rosenschein. Bounding the cost of stability in games over interaction networks. In *Proceedings of Twenty-Seventh ACM Conference on Artificial Intelligence (AAAI)*, pages 690–696, 2013.
- [14] R. Myerson. Graphs and cooperation in games. *Mathematics of Operations Research*, (2(3)):225–229, 1977.
- [15] R.B. Myerson. Conference structures and fair allocation rules. *International Journal of Game Theory*, (9(3)):169–182, 1980.
- [16] B. Reed. Personal communication, 2002.

- [17] N. Robertson and P. Seymour. Graph minors V. Excluding a planar graph,. *Journal of Combinatorial Theory, Series B*, (41):92–114, 1986.
- [18] P. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, (58(1)):22–33, 1993.
- [19] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, (16(2)):264–280, 1971.