



The Bergman Complex of a Matroid and Phylogenetic Trees

Federico Ardila and Caroline J. Klivans

ABSTRACT. We study the Bergman complex $\mathcal{B}(M)$ of a matroid M : a polyhedral complex which arises in algebraic geometry, but which we describe purely combinatorially. We prove that a natural subdivision of the Bergman complex of M is a geometric realization of the order complex of its lattice of flats. In addition, we show that the Bergman fan $\tilde{\mathcal{B}}(K_n)$ of the graphical matroid of the complete graph K_n is homeomorphic to the space of phylogenetic trees \mathcal{T}_n .

1. Introduction

In [?], Bergman defined the *logarithmic limit-set* of an algebraic variety in order to study its exponential behavior at infinity. We follow [?] in calling this set the *Bergman complex* of the variety. Bergman conjectured that this set is a finite, pure polyhedral complex. He also posed the question of studying the geometric structure of this set; *e.g.*, its connectedness, homotopy, homology and cohomology. Bieri and Groves first proved the conjecture in [?] using valuation theory.

Recently, Bergman complexes have received considerable attention in several areas, such as tropical algebraic geometry and dynamical systems. They are the *non-archimedean amoebas* of [?] and the *tropical varieties* of [?, ?]. In particular, Sturmfels [?] gave a new description of the Bergman complex and an alternative proof of Bergman's conjecture in the context of Gröbner basis theory. Moreover, when the variety is a linear space, so the defining ideal I is generated by linear forms, he showed that the Bergman complex can be described solely in terms of the matroid associated to the linear ideal.

Sturmfels used this description to define the Bergman complex of an arbitrary matroid, and suggested studying its combinatorial, geometric and topological properties [?]. The goal of the paper is to undertake this study.

In Section ?? we study the collection of bases of minimum weight of a matroid. We show that this collection is itself the collection of bases of a matroid, and we give several descriptions of it.

In Section ?? we prove the main result of the paper. We show that, appropriately subdivided, the Bergman complex of a matroid M is the order complex of the proper part of the lattice of flats L_M of the matroid. These order complexes are well-understood objects [?], and an immediate corollary of our result is an answer to Bergman's questions about the geometry of $\mathcal{B}(M)$ in this special case. The Bergman complex of an arbitrary matroid M is a finite, pure polyhedral complex. In fact, it is homotopy equivalent to a wedge of $(r - 2)$ -dimensional spheres, where r is the rank of M .

In Section ??, we take a closer look at the Bergman complex of the graphical matroid of the complete graph K_n . We show that the Bergman fan $\tilde{\mathcal{B}}(K_n)$ is exactly the space of ultrametrics on $[n]$, which is

1991 *Mathematics Subject Classification*. Primary 05B35, 52B05, 92B10.

Key words and phrases. matroids, phylogenetics trees, discrete geometry.

homeomorphic to the space of phylogenetic trees as in [?]. As a consequence, we show that the order complex of the partition lattice Π_n is a subdivision of the link of the origin of this space. This provides a new explanation for the known result that these two simplicial complexes are homotopy equivalent [?, ?, ?, ?, ?].

Finally, in the appendix, we review some matroid theory which we will use throughout the paper. For a more thorough introduction, we refer the reader to [?].

2. The bases of minimum weight of a matroid

Let M be a matroid of rank r on the ground set $[n] = \{1, 2, \dots, n\}$, and let $\omega \in \mathbb{R}^n$. Regard ω as a weight function on M , so that the weight of a basis $B = \{b_1, \dots, b_r\}$ of M is given by $\omega_B = \omega_{b_1} + \omega_{b_2} + \dots + \omega_{b_r}$.

Let M_ω be the collection of bases of M having minimum ω -weight. This is one of the central objects of our study, and we wish to understand it from three different points of view: geometric, algorithmic and matroid theoretic.

Geometrically, we can understand M_ω in terms of the matroid polytope. We will use the following characterization of matroid polytopes, due to Gelfand and Serganova:

THEOREM 2.1. [?, Theorem 1.11.1] *Let S be a collection of r -subsets of $[n]$. Let P_S be the polytope in \mathbb{R}^n with vertex set $\{e_{b_1} + \dots + e_{b_r} \mid \{b_1, \dots, b_r\} \in S\}$, where e_i is the i -th unit vector. Then S is the collection of bases of a matroid if and only if every edge of P_S is a translate of the vector $e_i - e_j$ for some $i, j \in [n]$.*

Let P_M be the matroid polytope of M . We can now think of ω as a linear functional in \mathbb{R}^n . The bases in M_ω correspond to the vertices of P_M which minimize the linear functional ω . Their convex hull is P_{M_ω} , the face of P_M where ω is minimized. It follows that the edges of P_{M_ω} , being edges of P_M also, are parallel to vectors of the form $e_i - e_j$. Therefore M_ω is the collection of bases of a matroid.

Algorithmically, matroids have the property that their ω -minimum bases are precisely the possible outputs of the greedy algorithm: Start with $B = \emptyset$. At each stage, look for an ω -minimum element of $[n]$ which can be added to B without making it dependent, and add it. After r steps, output the basis B . [?, Theorem 1.8.5]

DEFINITION 2.2. Given $\omega \in \mathbb{R}^n$, let $\mathcal{F}(\omega)$ denote the unique flag of subsets

$$\emptyset =: F_0 \subset F_1 \subset \dots \subset F_k \subset F_{k+1} := E$$

for which ω is constant on each set $F_i - F_{i-1}$ and has $\omega|_{F_i - F_{i-1}} < \omega|_{F_{i+1} - F_i}$. The *weight class* of a flag \mathcal{F} is the set of vectors ω such that $\mathcal{F}(\omega) = \mathcal{F}$.

We can describe weight classes by their defining equalities and inequalities. For example, one of the weight classes in \mathbb{R}^5 is the set of vectors ω such that $\omega_1 = \omega_4 < \omega_2 < \omega_3 = \omega_5$. It corresponds to the flag $\{\emptyset \subset \{1, 4\} \subset \{1, 2, 4\} \subset \{1, 2, 3, 4, 5\}\}$.

PROPOSITION 2.3. *If ω is in the weight class of $\mathcal{F} = \{\emptyset =: F_0 \subset \dots \subset F_{k+1} := E\}$, then the ω -minimum bases of M are exactly those containing $r(F_i) - r(F_{i-1})$ elements of $F_i - F_{i-1}$, for each i . Consequently, M_ω depends only on \mathcal{F} , and we call it $M_{\mathcal{F}}$.*

PROOF. The greedy algorithm picks $r(F_1)$ elements of the lowest weight, until it reaches a basis of F_1 ; then it picks $r(F_2) - r(F_1)$ elements of the second lowest weight, until it reaches a basis of F_2 , and so on. Therefore, the possible outputs of the algorithm are precisely the ones described. □

Matroid theoretically, M_ω can be constructed as a direct sum of minors of M , and its lattice of flats L_{M_ω} can be constructed from intervals of L_M , as follows:

PROPOSITION 2.4. *If $\mathcal{F} = \{\emptyset =: F_0 \subset \dots \subset F_{k+1} := E\}$, then*

$$M_{\mathcal{F}} = \bigoplus_{i=1}^{k+1} (M|F_i)/F_{i-1} \quad \text{and} \quad L_{M_{\mathcal{F}}} \cong \prod_{i=1}^{k+1} [F_{i-1}, F_i].$$

PROOF. After $r(F_{i-1})$ steps, the greedy algorithm has chosen a basis of F_{i-1} . In the following $r(F_i) - r(F_{i-1})$ steps, it needs to choose elements which, when added to F_{i-1} , give a basis of F_i . The possible choices are, precisely, the bases of $(M|F_i)/F_{i-1}$. The first equality follows, and the second one follows from it. \square

3. The Bergman complex

We now define the two main objects of study of this paper.

DEFINITION 3.1. The *Bergman fan* of a matroid M with ground set $[n]$ is the set

$$\tilde{\mathcal{B}}(M) := \{\omega \in \mathbb{R}^n : M_{\omega} \text{ has no loops}\}.$$

The *Bergman complex* of M is

$$\mathcal{B}(M) := \{\omega \in S^{n-2} : M_{\omega} \text{ has no loops}\},$$

where S^{n-2} is the sphere $\{\omega \in \mathbb{R}^n : \omega_1 + \dots + \omega_n = 0, \omega_1^2 + \dots + \omega_n^2 = 1\}$.

For the moment, we are slightly abusing notation by calling these two objects a *fan* and a *complex*. We will very soon see that they are a polyhedral fan and a spherical polyhedral complex, respectively; this justifies their name. We will concentrate on the Bergman complex, but the same arguments apply to the Bergman fan.

DEFINITION 3.2. The weight class of a flag \mathcal{F} is *valid* for M if $M_{\mathcal{F}}$ has no loops.

Since the matroid M_{ω} only depends on the weight class that ω is in, the Bergman complex of M is a disjoint union of the following weight classes:

DEFINITION 3.3. The weight class of a flag \mathcal{F} is *valid* for M if $M_{\mathcal{F}}$ has no loops.

We will study two polyhedral subdivisions of $\mathcal{B}(M)$, one of which is clearly finer than the other.

DEFINITION 3.4. The *fine subdivision* of $\mathcal{B}(M)$ is the subdivision of $\mathcal{B}(M)$ into valid weight classes: two vectors u and v of $\mathcal{B}(M)$ are in the same class if and only if $\mathcal{F}(u) = \mathcal{F}(v)$.

The *coarse subdivision* of $\mathcal{B}(M)$ is the subdivision of $\mathcal{B}(M)$ into M_{ω} -equivalence classes: two vectors u and v of $\mathcal{B}(M)$ are in the same class if and only if $M_u = M_v$.

THEOREM 3.5. *The weight class of a flag \mathcal{F} is valid for M if and only if \mathcal{F} is a flag of flats of M . Therefore, the fine subdivision of the Bergman complex $\mathcal{B}(M)$ is a geometric realization of $\Delta(L_M - \{\hat{0}, \hat{1}\})$, the order complex of the proper part of the lattice of flats of M .*

PROOF. Assume F_i in \mathcal{F} is not a flat of M , so there exists some $e \in \overline{F_i} - F_i$. By Proposition ??, any basis B in $M_{\mathcal{F}}$ contains $r(F_i)$ elements of F_i ; since e is dependent on them, it cannot be in B . Hence e is a loop in $M_{\mathcal{F}}$, so the weight class of \mathcal{F} is not valid.

Conversely, assume every F_i in \mathcal{F} is a flat of M . Consider any $e \in E$, and find the value of i such that $e \in F_i - F_{i-1}$. After $r(F_{i-1})$ steps the greedy algorithm produces a basis of F_{i-1} . Since F_{i-1} is a flat, e is not dependent on it, and in the next step of the algorithm we can choose e . In the end, this produces an ω -minimum basis of M containing e . Therefore the weight class of \mathcal{F} is valid. \square

The order complex $\Delta(L_M - \{\hat{0}, \hat{1}\})$ is a well understood object [?]. As an immediate consequence of Theorem ??, we get the following result.

COROLLARY 3.6. *The Bergman complex $\mathcal{B}(M)$ is homotopy equivalent to a wedge of $\hat{\mu}(L_M)$ $(r - 2)$ -dimensional spheres. Its subdivision into weight classes is a pure, shellable simplicial complex.*

Here $\widehat{\mu}(L_M) = (-1)^{r(M)} \mu_{L_M}(\hat{0}, \hat{1})$ is an evaluation of the *Möbius function* μ_{L_M} of the lattice L_M . The Möbius function is an extremely useful combinatorial invariant of a poset; for more information, see [?, Chapter 3].

Example: Let M be the graphical matroid of the complete graph on four nodes. The bases of this matroid are given by spanning trees. The flats are complete subgraphs and vertex disjoint unions of complete subgraphs (see Figure ??). Note that in this case, the flats are in correspondence with the partitions of the set $\{A, B, C, D\}$. In general, the flats of the graphical matroid of K_n are in bijection with partitions of the set $[n]$. Furthermore, the lattice of flats is the partition lattice Π_n , which orders partitions by refinement.




FIGURE 1. The graph K_4 and the lattice of flats of the corresponding matroid.

The Bergman complex $\mathcal{B}(K_4)$ is shown in Figure ??. It is a wedge of six 1-spheres. More generally, $\mathcal{B}(K_n)$ is a wedge of $\widehat{\mu}(\Pi_n) = (n - 1)!$ spheres of dimension $n - 3$. The vertices of $\mathcal{B}(K_4)$ are labeled with the corresponding flats, and a few of the corresponding weight classes are shown. Notice that the ground set of a matroid is always a flat, which corresponds to the weight class in which all weights are equal. We removed this weight class when normalizing the Bergman complex to the sphere.




FIGURE 2. The Bergman complex of the graphical matroid of K_4 .

The fine subdivision of the Bergman complex is almost the Petersen graph. The only difference is the presence of the three extra vertices, 13, 24 and 56. In the coarse subdivision into M_ω -equivalence classes, these three vertices do not appear. For example, the weight class $\omega_1 < \omega_3 < \omega_2 = \omega_4 = \omega_5 = \omega_6$ induces the same matroid M_ω as $\omega_1 = \omega_3 < \omega_2 = \omega_4 = \omega_5 = \omega_6$ and $\omega_3 < \omega_1 < \omega_2 = \omega_4 = \omega_5 = \omega_6$. Next we describe the relationship between these two subdivisions in general.

The coarse decomposition of $\mathcal{B}(M)$ into cells which induce the same M_ω is also a pure, polyhedral complex: it is a subcomplex of the spherical polar to the matroid polytope of M . To describe this decomposition, it is enough to describe its full-dimensional cells.

Therefore, we only need to determine when two full-dimensional weight classes give the same matroid M_ω . It is clearly enough to answer this question when the two weight classes are *adjacent*; *i.e.*, the intersection of their closures is a facet of both. This happens when the two corresponding flags, which have one flat in each rank, are equal in all but one rank.

Let \diamond be the *diamond poset*; *i.e.*, the rank 2 poset consisting of a minimum element, a maximum element, and two rank 1 elements.

THEOREM 3.7. *Suppose that the weight classes of two maximal flags \mathcal{F} and \mathcal{F}' are adjacent. Say \mathcal{F} and \mathcal{F}' only differ in rank i ; that is, $\mathcal{F} - F_i = \mathcal{F}' - F'_i$. Then the following conditions are equivalent:*

- (i) $M_{\mathcal{F}} = M_{\mathcal{F}'}$.
- (ii) $M_{\mathcal{F}} = M_{\mathcal{F} - F_i}$.
- (iii) $F_i \cup F'_i = F_{i+1}$.
- (iv) *The interval $[F_{i-1}, F_{i+1}]$ of L_M is a diamond poset.*

PROOF. Let $M_j = (M|F_j)/F_{j-1}$, $M'_j = (M|F'_j)/F'_{j-1}$, $N_i = (M|F_{i+1})/F_{i-1}$, and $N = M_1 \oplus \cdots \oplus M_{i-1} \oplus M_{i+2} \oplus \cdots \oplus M_{k+1}$. By Proposition ??,

$$M_{\mathcal{F}} = N \oplus M_i \oplus M_{i+1}, \quad M_{\mathcal{F}'} = N \oplus M'_i \oplus M'_{i+1}, \quad M_{\mathcal{F} - F_i} = N \oplus N_i.$$

Since M_i, M_{i+1}, M'_i and M'_{i+1} have rank 1 and N_i has rank 2,

$$\begin{aligned} L_{M_i \oplus M_{i+1}} &= \{\emptyset, F_i - F_{i-1}, F_{i+1} - F_i, F_{i+1} - F_{i-1}\} \cong \diamond, \\ L_{M'_i \oplus M'_{i+1}} &= \{\emptyset, F'_i - F_{i-1}, F_{i+1} - F'_i, F_{i+1} - F_{i-1}\} \cong \diamond, \\ L_{N_i} &= \{F - F_{i-1} : F \in [F_{i-1}, F_{i+1}]\} \cong [F_{i-1}, F_{i+1}]. \end{aligned}$$

If (iv) does not hold, then we know immediately that $L_{N_i} \neq L_{M_i \oplus M_{i+1}}$. Also $F_i \cup F'_i \neq F_{i+1}$, and therefore $L_{M_i \oplus M_{i+1}} \neq L_{M'_i \oplus M'_{i+1}}$.

If (iv) holds, then F_i and F'_i are the only rank i flats of M in $[F_{i-1}, F_{i+1}]$. Since N_i has no loops, (iii) holds; and therefore $L_{M_i \oplus M_{i+1}} = L_{M'_i \oplus M'_{i+1}} = L_{N_i}$. \square

4. The space of phylogenetic trees

In this section, we show that the Bergman fan $\tilde{\mathcal{B}}(K_n)$ of the matroid of the complete graph K_n is homeomorphic to the space of phylogenetic trees \mathcal{T}_n , as defined in [?]. To do so, we start by reviewing the connection between phylogenetic trees and ultrametrics.

DEFINITION 4.1. A *dissimilarity map* on $[n]$ is a map $\delta : [n] \times [n] \rightarrow \mathbb{R}$ such that $\delta(i, i) = 0$ for all $i \in [n]$, and $\delta(i, j) = \delta(j, i)$ for all $i, j \in [n]$.

DEFINITION 4.2. A dissimilarity map is an *ultrametric* if, for all $i, j, k \in [n]$, two of the values $\delta(i, j), \delta(j, k)$ and $\delta(i, k)$ are equal and not less than the third.

Let T be a rooted metric n -tree; that is, a tree with n leaves labeled $1, 2, \dots, n$, together with a length assigned to each one of its edges. For each pair of leaves u, v of the tree, we define the *distance* $d_T(u, v)$ to be the length of the unique path joining leaves u and v in T . This gives us a distance function $d_T : [n] \times [n] \rightarrow \mathbb{R}$.

We will only consider *equidistant n -trees*. These are the rooted metric n -trees such that the distance between the root and any leaf is equal to 1, and the lengths of the interior edges are positive. (For technical reasons, the edges incident to a leaf are allowed to have negative lengths.)

We can think of equidistant trees as a model for the evolutionary relationships between a certain set of species. The various species, represented by the leaves, descend from a single root. The descent from the root to a leaf tells us the history of how a particular species branched off from the others. For more information on the applications of this and other similar models, see for example [?] and [?].

The connection between equidistant trees and ultrametrics is given by the following theorem.

THEOREM 4.3. [?, Theorem 7.2.5] *A map $\delta : [n] \times [n] \rightarrow \mathbb{R}$ is an ultrametric if and only if it is the distance function of an equidistant n -tree.*

We can think of a dissimilarity map $\delta : [n] \times [n] \rightarrow \mathbb{R}$ as a weight function ω_δ on the edges of the complete graph K_n . This leads us to the following result, which connects these ideas to the Bergman fan.

THEOREM 4.4. *A dissimilarity map $\delta : [n] \times [n] \rightarrow \mathbb{R}$ is an ultrametric if and only if ω_δ is in the Bergman fan $\tilde{\mathcal{B}}(K_n)$.*

PROOF. We claim that the following three statements about a weight function on the edges of K_n are equivalent.

- (i) In any triangle, the largest weight is achieved (at least) twice.
- (ii) In any cycle, the largest weight is achieved (at least) twice.
- (iii) Every edge is in a spanning tree of minimum weight.

The theorem will follow from this claim, because ultrametrics are characterized by (i) and weight functions in the Bergman complex are characterized by (iii).

The implication (ii) \Rightarrow (i) is trivial. Conversely, assume that (i) holds and (ii) does not. Without loss of generality, assume that the cycle $v_1 v_2 \dots v_k$ has $v_1 v_2$ as its unique edge of largest weight. The largest weight in triangle $v_1 v_2 v_3$ must be achieved at $\omega(v_1 v_2) = \omega(v_1 v_3)$. The largest weight in triangle $v_1 v_3 v_4$ must then be achieved at $\omega(v_1 v_3) = \omega(v_1 v_4)$. Continuing in this way we get that $\omega(v_1 v_2) = \omega(v_1 v_3) = \dots = \omega(v_1 v_k)$, and (ii) follows.

Now we prove (ii) \Rightarrow (iii). Consider an arbitrary edge f . Let T be a spanning tree of minimum weight. If $f \in T$ we are done; otherwise, $T \cup f$ has a unique cycle. There is at least one edge e in this cycle with $\omega(e) \geq \omega(f)$. Therefore, the weight of the spanning tree $T \setminus e \cup f$ is not larger than the weight of T . This is then a spanning tree of minimum weight containing f .

Finally, assume that (iii) holds and (i) does not. Assume that the triangle with edges e, f, g has $\omega(e) > \omega(f), \omega(g)$, and consider a spanning tree T of minimum weight which contains edge e . If f is in T , then g cannot be in T , and replacing e with g will give a spanning tree of smaller weight. If neither f nor g is in T , we can still replace e with one of them to obtain a spanning tree of smaller weight. If we could not, that would imply that both f and g form a cycle when added to $T \setminus e$. Call these cycles C_f and C_g . But then $(C_f \setminus f) \cup (C_g \setminus g) \cup e$ would contain a cycle in T , a contradiction. \square

The previous two theorems give us a one-to-one correspondence between the vectors in the Bergman fan $\tilde{\mathcal{B}}(K_n)$ and the equidistant n -trees: $\tilde{\mathcal{B}}(K_n)$ parameterizes equidistant n -trees by the distances between their leaves. This leads us to consider the space of trees \mathcal{T}_n of [?]. This space parameterizes equidistant n -trees in a different way: it keeps track of their combinatorial type, and the lengths of their internal edges. We recall the construction of the space \mathcal{T}_n . Each maximal cell corresponds to a combinatorial type of rooted binary tree on n labeled leaves; *i.e.*, a rooted tree where each internal vertex has two descendants. Such trees have $n - 2$ internal edges, and are parameterized by vectors in $\mathbb{R}_{>0}^{n-2}$ recording these edge lengths. Moving to a lower dimensional face of a maximal cell corresponds to setting some of these edge lengths to 0, which gives non-binary degenerate cases of the original tree. Maximal cells are glued along these lower-dimensional cells when two trees specialize to the same degenerate tree.

Given a fixed combinatorial type of tree and the vector of internal edge lengths, we can recover the pairwise distances of leaves as linear functions on the internal edge lengths. For example, consider the tree type of Figure ???. We obtain $(\delta(A, B), \delta(A, C), \delta(A, D), \delta(B, C), \delta(B, D), \delta(C, D)) \in \mathcal{B}(K_4)$ from (x, y) by the map $f : (x, y) \mapsto (2(1 - x - y), 2(1 - y), 2, 2(1 - y), 2, 2)$. The converse is also true; given the pairwise distances of leaves we can recover the internal edge lengths via linear relations on these distances [?].




FIGURE 3. Combinatorial type of tree with 4 leaves.

In general, doing this for each type of tree, we get a map $f : \mathcal{T}_n \rightarrow \tilde{\mathcal{B}}(K_n)$. It follows from the previous two theorems that f is a one-to-one correspondence between \mathcal{T}_n and $\tilde{\mathcal{B}}(K_n)$. We will now see that, in fact, \mathcal{T}_n and $\tilde{\mathcal{B}}(K_n)$ have the same combinatorial structure.

PROPOSITION 4.5. *The map $f : \mathcal{T}_n \rightarrow \tilde{\mathcal{B}}(K_n)$ is a piecewise linear homeomorphism. It identifies the decomposition of the space of trees \mathcal{T}_n into combinatorial tree types with the coarse decomposition of the Bergman fan $\tilde{\mathcal{B}}(K_n)$.*

PROOF. Restricting to a maximal cell of \mathcal{T}_n , corresponding to a fixed tree type, f is a linear map from the lengths of internal edges (in the space of trees) to the pairwise distances of the leaves (in the space of ultrametrics). Also, it is clear that when two maximal cells of \mathcal{T}_n intersect, the linear restrictions of f to these two cells agree on their intersection. The first claim follows.

Suppose we are given a combinatorial type of equidistant n -tree. From the branching order of each triple of leaves (*i.e.*, which, if any, of the three branched off first), we can recover which edges of each triangle of K_n are maximum in the corresponding weight vector. In turn, this allows us to recover which edges of any cycle are maximum: one can check that an edge is maximum in a cycle C if and only if it is maximum in each triangle that it forms with a vertex of C . Knowing the maximum edges of each cycle of the graph, we can determine M_ω using the following version of the greedy algorithm. Start with the complete graph K_n and break its cycles successively: at each step pick an existing cycle, and remove one of its maximum edges. The trees which can result by applying this procedure are precisely the ω -minimum spanning trees [?]. Therefore f maps a fixed tree type class of \mathcal{T}_n to a fixed M_ω -equivalence class; *i.e.*, a fixed cell in the coarse subdivision of $\tilde{\mathcal{B}}(K_n)$.

Conversely, suppose we are given M_ω (which has no loops) and we want to determine the combinatorial tree type of $f^{-1}(\omega)$. Consider the edges $\{e, f, g\}$ of any triangle in K_n ; we can find out whether e is maximum in this triangle as follows. Take a minimum spanning tree T containing e . Either $T \setminus e \cup f$ or $T \setminus e \cup g$ is a spanning tree; assume it is the first. If $T \setminus e \cup f$ is a minimum spanning tree, then $\omega(e) = \omega(f)$, and e is maximum in the triangle. Otherwise $\omega(e) < \omega(f)$ and e is not maximum in the triangle. Determining this information for each triangle tells us, for each triple of leaves, which one (if any) branched off first in the corresponding tree. It is easy to reconstruct the combinatorial type of the tree from this data, in the same way that one recovers an equidistant tree from its corresponding ultrametric [?, Theorem 7.2.5]. \square

The link of the origin in the coarse subdivision of \mathcal{T}_n , which we call T_n , is a simplicial complex which has appeared in many different contexts. It was first considered by Boardman [?], and also studied by Readdy [?], Robinson and Whitehouse [?], Sundaram [?], Trappmann and Ziegler [?], Vogtmann [?], and Wachs [?],

among others. By Theorem ??, the link of the origin in the fine subdivision of $\tilde{\mathcal{B}}(K_n)$ is the order complex of the partition lattice Π_n . We conclude the following result.

COROLLARY 4.6. *The order complex of the partition lattice Π_n is a subdivision of the complex T_n .*

This provides a new explanation of the known result [?, ?, ?, ?, ?, ?] that these two simplicial complexes are homotopy equivalent; namely, they have the homotopy type of a wedge of $(n - 1)!$ $(n - 3)$ -dimensional spheres.




FIGURE 4. The fine subdivision of $\mathcal{B}(K_4)$.

Let us now revisit the example of the last section. In Figure ?? we show the Bergman complex $\mathcal{B}(K_4)$, with some of the corresponding trees. We now know that this is a subdivision of T_4 , the link of the origin in the space of phylogenetic trees with 4 leaves, which is the Petersen graph. The three extra vertices in the fine subdivision are 13, 24 and 56. The tree corresponding to vertex 24 of the fine subdivision has the property that the vertex joining the leaves C and D is at the same height as the vertex joining the leaves A and B . This information is not captured by the combinatorial type of the tree; *i.e.*, by the coarse subdivision.

In Figure ??, we show a representative piece of the fine subdivision of the space of trees with 5 leaves, with K_5 labeled as shown.

5. Acknowledgments

We would like to thank Bernd Sturmfels for suggesting to us the project of studying the Bergman complex of a matroid, and for many helpful conversations on this topic. We would also like to thank Lou Billera, Günter Ziegler, and an anonymous referee, for reading preliminary versions of this manuscript and helping us improve the presentation.




FIGURE 5. A piece of the fine subdivision of $\mathcal{B}(K_5)$.

6. Appendix: Matroids

DEFINITION 6.1. A *matroid* M on a finite ground set E is a collection of subsets \mathcal{I} such that:

- (1) $\emptyset \in \mathcal{I}$
- (2) If $I_1 \in \mathcal{I}$ and $I_2 \subseteq I_1$ then $I_2 \in \mathcal{I}$.
- (3) If $I_1, I_2 \in \mathcal{I}$ and $|I_1| < |I_2|$ then there exists an element $x \in I_2 - I_1$ such that $I_1 \cup x \in \mathcal{I}$.

Elements of \mathcal{I} are referred to as the *independent sets* of M

For example, let V be a finite set of vectors in \mathbb{R}^n . Let \mathcal{I} be the collection of all linearly independent subsets of V . Then \mathcal{I} is the collection of independent sets of a matroid on the ground set V . This fundamental example motivates many concepts and results in matroid theory.

First we review some standard matroid concepts. A *basis* is a maximal independent set. A *circuit* is a minimal dependent set and we call a one element dependent set a *loop*.

DEFINITION 6.2. The *rank* of a subset $X \subset E$, $r(X)$, is the size of the largest independent set in X . A set X is a *flat* if for all elements $e \in E - X$, $r(X \cup e) > r(X)$. The *closure* \overline{X} of a subset X of E is the minimal flat that contains it. The poset of flats ordered by containment is a lattice, which we call the *lattice of flats*, L_M .

DEFINITION 6.3. The *order complex* of a finite poset P is the simplicial complex $\Delta(P) = \{C \subset P \mid C \text{ is a chain of } P\}$.

DEFINITION 6.4. The *matroid polytope* of M is the polytope P_M in \mathbb{R}^E with vertex set $\{e_{b_1} + \dots + e_{b_r} \mid \{b_1, \dots, b_r\} \text{ is a basis of } M\}$, where e_i is the i -th unit vector.

DEFINITION 6.5. For a subset $T \subset E$, the *restriction of M to T* , or *deletion of $E - T$ from M* , is the matroid on the ground set T , whose rank function is $r_{M|T}(X) = r_M(X)$ for $X \subseteq T$. This matroid is denoted $M|T$ or $M \setminus (E - T)$.

DEFINITION 6.6. For a subset $T \subset E$, the *contraction of T from M* is the matroid on the ground set $E - T$, whose rank function is $r_{M/T}(X) = r_M(X \cup T) - r_M(T)$. This matroid is denoted M/T .

DEFINITION 6.7. Given two matroids M_1 and M_2 on disjoint sets E_1 and E_2 , there is a matroid $M_1 \oplus M_2$ on the ground set $E_1 \cup E_2$, called the *direct sum* of M_1 and M_2 . Its bases are the sets of the form $B_1 \cup B_2$, where B_1 and B_2 are bases of M_1 and M_2 respectively. Its flats are the sets of the form $F_1 \cup F_2$, where F_1 and F_2 are flats of M_1 and M_2 respectively.

References

- [1] G. Bergman. *The logarithmic limit-set of an algebraic variety*. Transactions of the American Mathematical Society, vol. 157, 459-469, 1971.
- [2] R. Bieri and J. Groves. *The geometry of the set of characters induced by valuations*. J. Reine Angew. Math., vol. 347, 168-195, 1984.
- [3] L. Billera, S. Holmes, and K. Vogtmann. *Geometry of the space of phylogenetic trees*. Advances in Applied Mathematics, vol. 27, 733-767, 2001.
- [4] A. Björner. *The homology and shellability of matroids and geometric lattices*. Matroid Applications. Encyclopedia of Mathematics and Its Applications, vol. 40. Cambridge University Press, Cambridge, 1992.
- [5] J. M. Boardman. *Homotopy structures and the language of trees*. Proc. Sympos. Pure Math., vol. 22, 37-58, 1971.
- [6] A. Borovik, I. Gelfand, and N. White. *Coxeter matroids*. Birkhauser, Boston, 2003.
- [7] M. Einsiedler, M. Kapranov, D. Lind, and T. Ward. *Non-archimedean amoebas*. Preprint, 2003.
- [8] D. Kozen. *The design and analysis of algorithms*. Lecture Notes in Computer Science, Springer-Verlag, 1991.
- [9] J. G. Oxley. *Matroid theory*. Oxford University Press, New York, 1992.
- [10] M. A. Readdy. *The pre-WDVV ring of physics and its topology*, preprint, 2002.
- [11] A. Robinson and S. Whitehouse. *The tree representation of σ_{n+1}* . J. Pure Appl. Algebra, vol. 111, 245-253, 1996.
- [12] C. Semple and M. Steel. *Phylogenetics*. Oxford Lecture Series in Mathematics and Its Applications, Oxford University Press, 2003.
- [13] D. Speyer and B. Sturmfels. *The tropical Grassmannian*. Preprint, arXiv:math.AG/0304218.
- [14] R. P. Stanley. *Enumerative combinatorics, vol. 1*. Cambridge University Press, New York, 1986.
- [15] B. Sturmfels. *Solving systems of polynomial equations*. CBMS Regional Conference Series in Mathematics, vol. 97. American Mathematical Society, Providence, RI, 2002.
- [16] S. Sundaram. *Homotopy of the non-modular partitions and the Whitehouse module*. J. Algebraic Combinatorics, vol. 9, 251-269, 1999.
- [17] H. Trappmann and G. M. Ziegler. *Shellability of complexes of trees*. Journal of Combinatorial Theory, Series A, vol. 82, 168-178, 1998.
- [18] K. Vogtmann. *Local structure of some $OUT(F_n)$ -complexes*. Proc. Edinburgh Math. Soc. (2), vol. 33, 367-379, 1990.
- [19] M. Wachs, personal communication, 2003.

MATHEMATICAL SCIENCES RESEARCH INSTITUTE, BERKELEY, CA, 94709

E-mail address: federico@msri.org

URL: <http://www.msri.org/~federico>

MATHEMATICAL SCIENCES RESEARCH INSTITUTE, BERKELEY, CA, 94709

Current address: 4114 Upson Hall, Cornell University, Ithaca, NY, 14853

E-mail address: cjk11@cornell.edu

URL: <http://www.cs.cornell.edu/~cjk>