Submodular Partition Functions

Omid Amini, Frédéric Mazoit, Nicolas Nisse, Stéphan Thomassé

Submodular partition functions

Omid Amini

Projet Mascotte, CNRS/INRIA/UNSA, INRIA Sophia-Antipolis, 2004 route des Lucioles BP 93, 06902 Sophia-Antipolis Cedex, France
oamini@sophia.inria.fr

Frédéric Mazoit [1]

LaBRI Université Bordeaux F-33405 Talence Cedex, France
Frederic.Mazoit@labri.fr

Nicolas Nisse

LRI, Université Paris-Sud, 91405 Orsay, France.
nisse@lri.fr

Stéphan Thomassé [1]

LIRMM-Université Montpellier II, 161 rue Ada, 34392 Montpellier Cedex, France
thomasse@lirmm.fr

Abstract

Adapting the method introduced in Graph Minors X [6], we propose a new proof of the duality between the bramble-number of a graph and its tree-width. This proof is based on a new definition of submodularity on partition functions which naturally extends the usual one on set functions. The technique simplifies the proof of bramble/tree-width duality since it does not rely on Menger's theorem. One can also derive from it all known dual notions of other classical width-parameters. Finally, it provides a dual for matroid tree-width.

1 Introduction.

In their seminal paper Graph Minors X [6], Robertson and Seymour introduced the notion of branch-width of a graph and its dual notion of tangle.

[1] research supported by the french ANR-project "Graph decompositions and algorithms (GRAAL)"

Their method is based on bias and tree-labellings. Later on, Seymour and Thomas [7] found a dual notion to tree-width, the bramble number (named after Reed [4]). The proof of the bramble-number/tree-width duality makes use of Menger's theorem to reconnect partial tree-decompositions, see for instance the textbook of Diestel [1]. Our aim in this paper is to show how the classical dual notions of width-parameters can be deduced from the original method of Graph Minors X.

In this paper, E will always denote a set with at least two elements. A *partitioning tree* on E is a tree T in which the leaves are identified with the elements of E in a one-to-one way. Therefore, every internal node v of T, if any, corresponds to the partition T_v of E which parts are the leaves of the subtrees obtained by deleting v.

An obvious way of forming a partitioning tree is simply to fix a central node which is linked to every element of E - a partitioning star. But what if we are not permitted to do so? Precisely, assume that a restricted set of partitions of E, called *admissible partitions*, is given. Is it possible to form an admissible partitioning tree? (i.e. such that every partition T_v is admissible). An obstruction to the existence of such a tree is the dual notion of bramble.

An *admissible bramble* is a nonempty set of pairwise intersecting subsets of E which contains a part of every admissible partition of E. It is routine to form an admissible bramble: just pick an element e of E, and collect, for every admissible partition, the part which contains e. Such a bramble is called *principal*. The crucial fact is that if there is a non-principal admissible bramble $\mathcal{B}$, there is no admissible partitioning tree. To see this, assume for contradiction that T is an admissible partitioning tree. For every internal node u of T, there is an element X of T_u which belongs to $\mathcal{B}$. Let v be the neighbour of u which belongs to the component of $T \setminus u$ having set of labels X. Orient the edge uv of T from u to v. Note that every internal node becomes the origin of an oriented edge. Observe also that an edge of T incident to a leaf never gets an orientation since $\mathcal{B}$ is non-principal. The contradiction follows from the fact that one edge of T carries two orientations, which is impossible since the elements of $\mathcal{B}$ are pairwise intersecting. Note that this argument fails when T has no internal vertex, i.e. E has two elements. In this case, the unique partitioning tree is by definition admissible, and every admissible bramble is principal.

Unfortunately, if no principal admissible bramble exists, there is not necessarily an admissible partitioning tree. In the first part of this paper, we prove that for some particular families of admissible partitions (e.g. generated by a submodular partition function) we have the following:

- Either there exists an admissible partitioning tree.

- Or there exists a non-principal admissible bramble.

The second part of the paper is devoted to the translation of this result into the different notions of width-parameters.

2 Submodular partition functions.

The *complement* of a subset X of E is the set $X^c := E \setminus X$. A *partition* of E is a set $\mathcal{X} = \{X_1, \ldots, X_n\}$ of subsets of E satisfying $X_1 \cup \cdots \cup X_n = E$ and $X_i \cap X_j = \emptyset$ for all $i \neq j$. The order in which the X_i's appear is irrelevant. We authorise degenerate partitions (i.e. the sets X_i can be empty). Let F be a subset of E. The partition

$$\mathcal{X}_{X_i \to F} := \{X_1 \cap F, \ldots, X_{i-1} \cap F, X_i \cup F^c, X_{i+1} \cap F, \ldots, X_n \cap F\}$$

is the partition obtained from $\mathcal{X}$ by *pushing X_i to F*.

A *partition function* is a function Φ defined from the set of partitions of E into the reals. Let $\mathcal{X}$ be a partition of E. We call $\Phi(\mathcal{X})$ the Φ-*width*, or simply *width*, of $\mathcal{X}$. Let k be an integer. A *k-partition* is a partition of width at most k. A partition function Φ is *submodular* if for every pair of partitions $\mathcal{X} = \{X_1, \ldots, X_n\}$ and $\mathcal{Y} = \{Y_1, \ldots, Y_l\}$, we have:

$$\Phi(\mathcal{X}) + \Phi(\mathcal{Y}) \geq \Phi(\mathcal{X}_{X_1 \to Y_1}) + \Phi(\mathcal{Y}_{Y_1 \to X_1})$$

To justify *a posteriori* our terminology, observe that for bipartitions, partition submodularity gives

$$\begin{aligned}\Phi(A, A^c) + \Phi(B, B^c) &= \Phi(A, A^c) + \Phi(B^c, B) \\ &\geq \Phi(A \cup (B^c)^c, A^c \cap B^c) + \Phi(B^c \cup A^c, B \cap A) \\ &\geq \Phi(A \cup B, A^c \cap B^c) + \Phi(A \cap B, A^c \cup B^c)\end{aligned}$$

which corresponds to the usual notion of submodularity when setting $\Phi(F) := \Phi(F, F^c)$, for every subset F of E.

Unfortunately, since some natural partition functions lack submodularity, we have to define a relaxed version of it. A partition function Φ is *weakly submodular* if for every pair of partitions $\mathcal{X} = \{X_1, \ldots, X_n\}$ and $\mathcal{Y} = \{Y_1, \ldots, Y_l\}$, at least one of the following holds:

(1) There exists F such that $X_1 \subseteq F \subseteq (Y_1 \setminus X_1)^c$ and $\Phi(\mathcal{X}) > \Phi(\mathcal{X}_{X_1 \to F^c})$
(2) $\Phi(\mathcal{Y}) \geq \Phi(\mathcal{Y}_{Y_1 \to X_1})$

Submodular partition functions are weakly submodular, it suffices to consider $F = Y_1^c$. Let us illustrate these notions. In what follows, $\mathcal{X} = \{X_1, \ldots, X_n\}$ is a partition of E.

- The key-example of a submodular partition function is the function *border* defined on the set of partitions of the edge set E of a graph $G = (V, E)$ by letting $\delta(\mathcal{X}) = |\Delta(\mathcal{X})|$ where

$$\Delta(\mathcal{X}) = \{x \in V \mid \exists xy \in X_i \text{ and } \exists xz \in X_j,\ i \neq j\}.$$

  We will often write, for a subset F of E, $\Delta(F)$ and $\delta(F)$ instead of $\Delta(F, F^c)$ and $\delta(F, F^c)$. The proof of submodularity is postponed to Section 5.1. As we will see, the function δ gives the tree-width of G.
- Let f be a submodular function on 2^E. We form a submodular partition function by letting $\Sigma_f(\mathcal{X}) = \sum_{i \in I} f(X_i)$. The proof of submodularity is postponed to Section 5.2. This function gives the tree-width of matroids.
- Let f be a symmetric submodular function on 2^E, i.e. satisfying moreover $f(A) = f(A^c)$ for all $A \subseteq E$. The function $\max_{i \in \{1,\ldots,n\}} f(X_i)$, which can be made weakly submodular, gives the notion of branch-width and its relatives like rank-width. It is treated in Section 5.3.
- Let Φ be a weakly submodular partition function and $p \geq 2$ be an integer. We form a weakly submodular partition function by letting $\Phi_p(\mathcal{X}) = \Phi(\mathcal{X})$ when the number of parts of $\mathcal{X}$ is at most p, and $+\infty$ otherwise (or any large constant integer).
- Let Φ be a weakly submodular partition function and $p \geq 2$ be an integer. By letting $\Phi'_p(\mathcal{X}) = \Phi(\mathcal{X})$ when the number of X_i with at least two elements is at most p, and $+\infty$ otherwise (or any large constant integer), we obtain a partition function which gives, in particular, the notion of path-width. This is a weakly submodular partition function if we only push subsets which are non-singletons.

Our choice of the partition submodularity condition is motivated by the analogy with usual submodular functions, when restricted to bipartitions. However, we never use the fact that X_1 and Y_1 may intersect, and could have defined the notion for disjoint X_1, Y_1. This less constrained definition is perfectly valid for the results presented here.

3 Search-trees.

A *bidirected tree* is a directed graph obtained from an undirected tree by replacing every edge by an oriented circuit of length two. A *search-tree* T on E is a bidirected tree on at least two nodes together with a label function l defined from the arcs of T into the subsets of E with the additional requirements:

- If u is an internal node of T, the sets $l(uv)$, for all outneighbours v of u, form a partition of E. We denote it by T_u.
- The labels of a 2-circuit do not intersect, i.e. $l(uv) \cap l(vu) = \emptyset$.

A 2-circuit uv is *exact* if $l(uv) \cup l(vu) = E$. By extension, a search-tree T is *exact* if all its 2-circuits are exact. The label of an arc with origin a leaf of T is called a *leaf-label*. Let $\mathcal{F}$ be a set of subsets of E. A search-tree T is *compatible* with $\mathcal{F}$ if every leaf-label of T contains an element of $\mathcal{F}$. Let uv be a 2-circuit of T where u is an internal node. Let F be a subset such that $l(uv) \subseteq F \subseteq l(vu)^c$. The key-fact is that replacing the partition T_u in T by $(T_u)_{l(uv)\to F^c}$ (in the obvious one-to-one way) gives a new search-tree which is still compatible with $\mathcal{F}$ since the leaf-labels are unchanged.

If Φ is a weakly submodular partition function on E, the Φ-*width* of a search-tree T with at least three nodes is the maximum of $\Phi(T_u)$, taken over the internal nodes u. If no confusion can occur, we just speak of the *width* of T. A k-*search-tree* is a search-tree with two nodes or having width at most k.

Theorem 1 *If Φ is a weakly submodular partition function and T is a k-search-tree compatible with $\mathcal{F}$, there is a relabelling of T which is an exact k-search-tree compatible with $\mathcal{F}$.*

PROOF. If T consists of a 2-circuit uv, we simply set $l(vu) := l(uv)^c$. Now, assume that amongst all relabellings of T which are k-search-trees compatible with $\mathcal{F}$, we minimise the sum of $\Phi(T_u)$, taken over all internal nodes u. Select an internal node r as the root of T. If T is not exact, we select a non exact 2-circuit uv, with u chosen closer to r than v. If v is a leaf, we simply replace $l(vu)$ by $l(uv)^c$. If v is an internal node, by the minimality of T, there is no F with $l(uv) \subseteq F \subseteq l(vu)^c$ for which $\Phi(T_u) > \Phi((T_u)_{l(uv)\to F^c})$. Since Φ is weakly submodular, we have $\Phi(T_v) \geq \Phi((T_v)_{l(vu)\to l(uv)})$. We then replace T_v by $(T_v)_{l(vu)\to l(uv)}$. Observe that both replacements strictly increase the sum of the sizes of the labels of backward arcs of T (those pointing toward the root). Thus this process stops on an exact k-search-tree which is still compatible with $\mathcal{F}$ since the leaf-labels can only increase. $\square$

In an exact search-tree T, the set of labels of the arcs entering the leaves forms a partition of E. Therefore the union of two leaf-labels is equal to E. When this partition consists of singletons and empty sets, T is a *partitioning* k-search-tree. In the full generality of partition functions, empty sets cannot be avoided, however in all the examples given below, we can *prune* partitioning trees to remove them.

4 Tree-bramble duality.

Let Φ be a weakly submodular partition function on E. A *bias* is a nonempty family $\mathcal{B}$ of subsets of E such that $\bigcap \mathcal{B} = \emptyset$. A *$k$-bramble* $\mathcal{B}$ is a nonempty family of subsets of E such that:

- For all $X, Y \in \mathcal{B}$, we have $X \cap Y \neq \emptyset$.
- For every k-partition $\mathcal{X} = \{X_1, \ldots, X_n\}$, there exists i such that $X_i \in \mathcal{B}$.

A k-bramble is *principal* if it is not a bias, i.e. $\bigcap \mathcal{B}$ is nonempty.

Theorem 2 *Let Φ be a weakly submodular partition function on a set E.*

i. Either there exists a non-principal k-bramble.
ii. Or there exists a partitioning k-search-tree.

PROOF. If there is a partitioning k-search-tree, every k-bramble is principal. The proof is given in the introduction in terms of admissible partitions. We now assume that every k-bramble is principal, and prove the existence of a partitioning k-search-tree. More generally, we show that every bias has a compatible k-search-tree. This gives our conclusion when considering the bias $\{E \setminus e \mid e \in E\}$. The proof goes by reverse induction on the inclusion order. Let $\mathcal{B}$ be a bias. We assume that the result holds for every bias $\mathcal{B}' \neq \mathcal{B}$ such that $\mathcal{B} \subseteq \mathcal{B}'$. Two cases can happen:

- For every k-partition $\mathcal{X} = \{X_1, \ldots, X_n\}$, there exists $X_i \in \mathcal{B}$. Since $\mathcal{B}$ is not a k-bramble, it contains two disjoint subsets B_i and B_j. Thus the 2-circuit labelled by B_i and B_j is a k-search-tree which is compatible with $\mathcal{B}$.
- There exists a k-partition $\mathcal{X} = \{X_1, \ldots, X_n\}$ such that $X_i \notin \mathcal{B}$, for all $i = 1, \ldots, n$. For each X_i, we choose a subset $X'_i \notin \mathcal{B}$ which contains X_i and which is maximal with respect to inclusion. We form the bias $\mathcal{B}_i := \mathcal{B} \cup \{X'_i\}$. By the induction hypothesis and Theorem 1, there exists an exact k-search-tree T_i compatible with $\mathcal{B}_i$. If T_i is also compatible with $\mathcal{B}$, we are done. If not, T_i has a leaf-label containing X'_i and no element of $\mathcal{B}$. Hence, by maximality of X'_i, this leaf-label is exactly X'_i. Observe that if T_i has two leaf-labels X'_i, since their union is E, we would have $X'_i = E$ and thus T_i would also be compatible with $\mathcal{B}$. Consequently, X'_i appears only once as a leaf-label. We form a new tree T by identifying, for every T_i, the leaf carrying the leaf-label X'_i. The tree T is not a search-tree since the labels of the arcs with origin the identified vertex are $\{X'_1, X'_2, \ldots, X'_n\}$, which is not a partition. We simply replace these labels by $X_1, X_2, \ldots, X_n$. Now T is a k-search-tree compatible with $\mathcal{B}$. $\square$

5 Examples of submodular partition functions.

5.1 The submodular partition function δ.

Let $G = (V, E)$ be a graph. Let $\mathcal{X} = \{X_1, \ldots, X_n\}$ and $\mathcal{Y} = \{Y_1, \ldots, Y_l\}$ be some partitions of E. We want to prove that:

$$\begin{aligned}\delta(\mathcal{X}) + \delta(\mathcal{Y}) &\geq \delta(\mathcal{X}_{X_1 \to Y_1}) + \delta(\mathcal{Y}_{Y_1 \to X_1}) \\ &\geq \delta(X_1 \cup Y_1^c, X_2 \cap Y_1, \ldots, X_n \cap Y_1) + \\ &\qquad \delta(Y_1 \cup X_1^c, Y_2 \cap X_1, \ldots, Y_n \cap X_1)\end{aligned}$$

Let x be a vertex of G. Two cases can happen:

- The contribution of x in the right-hand term of the previous inequality is one, say x belongs to the border of $\mathcal{X}_{X_1 \to Y_1}$. If x belongs to the border of Y_1, it contributes to $\delta(\mathcal{Y})$. If not, x belongs to the border of some X_i with $i > 1$. In both cases, its contribution to the left-hand term is at least one.
- Assume now that x both belongs to the borders of $\mathcal{X}_{X_1 \to Y_1}$ and $\mathcal{Y}_{Y_1 \to X_1}$. Since x belongs to the border of $\mathcal{X}_{X_1 \to Y_1}$ there is an edge e_x containing x in some $X_i \cap Y_1$ with $i > 1$. Similarly there is an edge f_x containing x in some $Y_j \cap X_1$ with $j > 1$. Since $e_x \in X_i$ and $f_x \in X_1$, x is in the border of $\mathcal{X}$. Similarly x is also in the border of $\mathcal{Y}$, and thus contributes also for two to the left-hand term.

5.2 The submodular partition function Σ_f.

Let f be a submodular function on 2^E.

Lemma 3 *(1) Let X and Y be two disjoint subsets of E. If $X_1 \subset X$ and $Y_1 \subset Y$, we have:*

$$f(X) + f(Y) - f(X_1) - f(Y_1) \geq f(X \cup Y) - f(X_1 \cup Y_1)$$

(2) More generally, if $X_1, \ldots, X_r$ are pairwise disjoint subsets of E, and for all $i = 1, \ldots, r$, $X_i' \subset X_i$, we have:

$$\sum_{i=1}^{r} \Big(f(X_i) - f(X_i') \Big) \geq f\Big(\bigcup_{i=1}^{r} X_i \Big) - f\Big(\bigcup_{i=1}^{r} X_i' \Big)$$

PROOF.

(1) Apply first the submodularity of f to the subsets $A = X \cup Y_1$ and $B = Y$. Since $A \cap B = Y_1$ and $A \cup B = X \cup Y$, we obtain:

$$f(X \cup Y_1) + f(Y) \geq f(X \cup Y) + f(Y_1) \tag{1}$$

Apply then the submodularity of f to the subsets $A = X_1 \cup Y_1$ and $B = X$. Since $A \cap B = X_1$ and $A \cup B = X \cup Y_1$, we obtain:

$$f(X_1 \cup Y_1) + f(X) \geq f(X \cup Y_1) + f(X_1) \tag{2}$$

The conclusion follows from (1)+(2).

(2) Follows by induction on r. □

Proposition 4 *The function Σ_f is a submodular partition function.*

PROOF. Let $\mathcal{X} = \{X_1, \ldots, X_n\}$ and $\mathcal{Y} = \{Y_1, \ldots, Y_l\}$ be two partitions of E. We want to prove that $\Sigma_f(\mathcal{X}) + \Sigma_f(\mathcal{Y}) \geq \Sigma_f(\mathcal{X}_{X_1 \to Y_1}) + \Sigma_f(\mathcal{Y}_{Y_1 \to X_1})$. We must then prove:

$$\begin{aligned} \sum_{i=1}^{n} f(X_i) + \sum_{j=1}^{l} f(Y_j) \geq f(X_1 \cup Y_1^c) + \sum_{i=2}^{n} f(Y_1 \cap X_i) \\ + f(Y_1 \cup X_1^c) + \sum_{j=2}^{l} f(X_1 \cap Y_j). \end{aligned} \tag{3}$$

By applying lemma 3 with $X_i' = Y_1 \cap X_i$ and since $X_2 \cup \cdots \cup X_n = X_1^c$, we have:

$$\sum_{i=2}^{n} f(X_i) - \sum_{i=2}^{n} f(Y_1 \cap X_i) \geq f(X_1^c) - f(Y_1 \cap X_1^c) \tag{4}$$

Similarly we obtain:

$$\sum_{j=2}^{l} f(Y_j) - \sum_{j=2}^{l} f(X_1 \cap Y_j) \geq f(Y_1^c) - f(X_1 \cap Y_1^c) \tag{5}$$

By adding (4) and (5), we obtain

$$\begin{aligned} \sum_{j=2}^{l} f(Y_j) + \sum_{i=2}^{n} f(X_i) + f(X_1 \cap Y_1^c) + f(Y_1 \cap X_1^c) \geq \\ f(Y_1^c) + f(X_1^c) + \sum_{j=2}^{l} f(X_1 \cap Y_j) + \sum_{i=2}^{n} f(Y_1 \cap X_i) \end{aligned} \tag{6}$$

By applying submodularity to X_1^c and Y_1 and to X_1 and Y_1^c, we obtain:

$$f(X_1) + f(Y_1) - f(X_1 \cap Y_1^c) - f(Y_1 \cap X_1^c) \geq f(X_1 \cup Y_1^c) + f(Y_1 \cup X_1^c) - f(Y_1^c) - f(X_1^c) \tag{7}$$

Adding (6) and (7), we obtain (3). Thus Σ_f is submodular. □

5.3 The weakly submodular partition function Max_f.

Let f be a symmetric submodular function on 2^E. Let $\mathcal{X} = \{X_1, \ldots, X_n\}$ be a partition of E. The function $\max_{i \in \{1,\ldots,n\}} f(X_i)$ is unfortunately not a weakly submodular partition function. We have to shift it a little to break ties. For some arbitrarily small $\varepsilon > 0$, we consider instead the function:

$$\mathrm{Max}_f(\mathcal{X}) = \max_{i \in \{1,\ldots,n\}} f(X_i) + \varepsilon \Sigma_f(\mathcal{X})$$

Lemma 5 *The function* Max_f *is a weakly submodular partition function.*

PROOF. Let $\mathcal{X} = \{X_1, \ldots, X_n\}$ and $\mathcal{Y} = \{Y_1, \ldots, Y_l\}$ be two partitions of E. Let F be a set such that

$$X_1 \subseteq F \subseteq (Y_1 \setminus X_1)^c \tag{8}$$

and chosen minimum with respect to f. Note that since X_1 satisfies (8), we have $f(F) \leq f(X_1)$. Assume that $f(F) < f(X_1)$. We claim that $\mathrm{Max}_f(\mathcal{X}) > \mathrm{Max}_f(\mathcal{X}_{X_1 \to F^c})$. Indeed, for every $i \geq 2$, we have by submodularity of f:

$$f(X_i) + f(F^c) \geq f(X_i \cap F^c) + f(X_i \cup F^c) \tag{9}$$

Furthermore, we have $f(F) \leq f(F \setminus X_i)$ by minimality of F, and thus by symmetry of f we get:

$$f(X_i \cup F^c) \geq f(F^c) \tag{10}$$

Adding (9) and (10), we obtain $f(X_i) \geq f(X_i \cap F^c)$. Thus the maximum of f over $\mathcal{X}$ is at least the maximum of f over $\mathcal{X}_{X_1 \to F^c}$. We now apply the submodularity of the function $\sum_f$ to the partitions $\mathcal{X}$ and $\{F^c, F\}$. We then obtain $\sum_f(\mathcal{X}) + \sum_f(F^c, F) \geq \sum_f(\mathcal{X}_{X_1 \to F^c}) + \sum_f(X_1^c, X_1)$. Since $f(X_1) > f(F)$, we have $\sum_f(F^c, F) < \sum_f(X_1^c, X_1)$, hence $\sum_f(\mathcal{X}) > \sum_f(\mathcal{X}_{X_1 \to F^c})$. Therefore, $\mathrm{Max}_f(\mathcal{X}) > \mathrm{Max}_f(\mathcal{X}_{X_1 \to F^c})$.

Assume now that $F = X_1$ is a minimum for f. By the same calculation as above, we obtain $\mathrm{Max}_f(\mathcal{Y}) \geq \mathrm{Max}_f(\mathcal{Y}_{Y_1 \to X_1})$. Thus Max_f is a weakly submodular partition function. □

6 Width parameters.

We assume in this section that the reader is somehow familiar with the usual definitions of tree-decompositions (such as tree-width, branch-width, path-width, rank-width,...). Our aim is just to associate a weakly submodular partition function to each of these parameters and show how to translate the exact partitioning k-search-tree into a tree-decomposition, and the non-principal k-bramble into the known dual notion (if any). To avoid technicalities, we assume that k is at least two and that $G = (V, E)$ is a graph with minimum degree two.

6.1 Tree-width of graphs.

The tree-width of G corresponds to the border δ defined on partitions of E.

Assume first that E has an exact partitioning k-search-tree T. Associate to every internal node u of T the bag $\Delta(T_u)$. The restriction of T to its internal nodes is a tree-decomposition of G. Indeed, for every edge xy of G, there is a leaf v of T for which $l(uv) = \{xy\}$, where uv is the arc of T with head v. Thus x and y belong to $\Delta(T_u)$, since the minimum degree in G is two. Furthermore, if a vertex of G both belongs to $\Delta(T_u)$ and $\Delta(T_v)$, it also belongs to $\Delta(T_w)$ for every node w in the (u, v)-path of T. Since every bag has size at most k, the tree-width of G is at most $k - 1$.

Now if E has a non-principal k-bramble $\mathcal{B}$, we form a bramble $\mathcal{B}'$ (in the usual sense). Let S be a subset of V with $|S| \leq k$. We associate to S the partition $\{E_1, \ldots, E_n\}$ of E where the sets E_i are the (nonempty) sets of edges minimal with respect to inclusion for the property $\Delta(E_i) \subseteq S$. Observe that this is indeed a partition since $\Delta(E_i \cap E_j) \subseteq \Delta(E_i) \cup \Delta(E_j) \subseteq S$. Since $\mathcal{B}$ is a non-principal k-bramble, one of the E_i, with at least two edges, is in $\mathcal{B}$. This means that $X_i = V(E_i) \setminus S$ is a nonempty set of vertices. In other words, E_i is the set of edges incident to at least one vertex of X_i (such a set is denoted by $E(X_i)$). We now collect, for all subsets S with $|S| \leq k$, these sets X_i to form our $\mathcal{B}'$. Observe first that, by minimality of E_i, every element X_i of $\mathcal{B}'$ induces a connected subgraph of G. We have now to prove that for every pair X_i, X_j of elements of $\mathcal{B}'$, $X_i \cup X_j$ also induces a connected subgraph of G. Indeed, let $E_i = E(X_i)$ and $E_j = E(X_j)$. Since the elements of $\mathcal{B}$ are pairwise intersecting, there is an edge xy of G in $E_i \cap E_j$. Without loss of generality, we can assume that $x \in X_i$. If we also have $x \in X_j$, the sets X_i and X_j intersect, and thus their union is connected. If $x \notin X_j$, we necessarily have $y \in X_j$, hence there is an edge of G connecting X_i and X_j. Thus $\mathcal{B}'$ is a bramble, and the minimum size covering set of $\mathcal{B}'$ has at least $k + 1$ elements. In this case

the bramble-number of G is at least $k+1$.

6.2 Branch-width of graphs.

The branch-width of G corresponds to the weakly submodular partition function $(\mathrm{Max}_\delta)_3$, which counts the maximum border of a subset in a partition of E into two or three subsets. An exact partitioning k-search tree of E is precisely a branch-decomposition of G of width k. Let us make now the correspondence between a non-principal k-bramble $\mathcal{B}$ and a tangle of G.

First of all, $\mathcal{B}$ is here a pairwise intersecting family of subsets of E such that every k-partition $\{E_1, E_2\}$ or $\{E_1, E_2, E_3\}$ contains an element of $\mathcal{B}$. The translation into a tangle of order $k+1$ is straightforward: when (G_1, G_2) is a separation of order at most k of G, we choose (G_1, G_2) in the tangle if $E(G_2) \in \mathcal{B}$, otherwise we choose (G_2, G_1). The second axiom of tangles asserts that if (A_1, B_1), (A_2, B_2) and (A_3, B_3) are in the tangle, we have $G \neq A_1 \cup A_2 \cup A_3$. It follows from the next proposition:

Proposition 6 *If E_1, E_2 and E_3 are in $\mathcal{B}$, we have $E_1 \cap E_2 \cap E_3 \neq \emptyset$.*

PROOF. Assume for contradiction that $E_1 \cap E_2 \cap E_3 = \emptyset$. Observe that

$$\delta(E_1 \cap E_2) + \delta(E_2 \cap E_3) + \delta(E_3 \cap E_1) \leq \delta(E_1) + \delta(E_2) + \delta(E_3).$$

So we can assume without loss of generality that, for instance, $\delta(E_1 \cap E_2) \leq k$. We also have that $\delta(E_1 \setminus E_2) + \delta(E_2 \setminus E_1) \leq \delta(E_1) + \delta(E_2)$. So we can assume for instance that $\delta(E_1 \setminus E_2) \leq k$. Then the partition $\{E_1^c, E_1 \setminus E_2, E_1 \cap E_2\}$ is a k-partition. But this is impossible since these three sets are respectively disjoint from E_1, E_2 and E_3, which all belong to $\mathcal{B}$. $\square$

The third axiom of tangles asserts that if (G_1, G_2) is a separation of G, we have $V(G_1) \neq V$. To see that, assume for contradiction that $V(G_1) = V$. We have $E(G_2) \in \mathcal{B}$, hence the number of vertices of G_2 is at most k. Thus every subset F of edges of G_2 is such that $\delta(F) \leq k$. Pick now F minimum with respect to inclusion such that $F \subseteq E(G_2)$ and $F \in \mathcal{B}$. Since $\mathcal{B}$ is non principal, we have $|F| \geq 2$. Let $\{F_1, F_2\}$ be a non trivial partition of F. The contradiction appears when considering the k-partition $\{F^c, F_1, F_2\}$ of E since these three sets are not in $\mathcal{B}$.

6.3 Rank-width.

The rank-width (see Oum and Seymour [3]) of G is based on the symmetric submodular function $cutrk$ defined on subsets of vertices X where $cutrk(X)$ is the rank (in $\mathbb{F}_2$) of the adjacency matrix of the bipartite graph $(X, V \setminus X)$. The submodular partition function on base set V is then $cutrk_3$. The partitioning exact k-search tree is precisely a rank-decomposition of G. A non-principal k-bramble $\mathcal{B}$ is here a pairwise intersecting family of subsets of V such that every k-partition $\{V_1, V_2\}$ or $\{V_1, V_2, V_3\}$ has an element in $\mathcal{B}$.

6.4 Path-width of graphs.

The path-width of $G = (V, E)$ corresponds to the partition function δ'_2, which is the border of partitions $\{X_1, \ldots, X_n\}$ of E with at most two parts with more than one element. The following analogue of Theorem 1 holds for partition functions Φ'_p, where Φ is a weakly submodular partition function, and $p \geq 2$ is some integer:

Theorem 7 *If T is a k-search-tree (with respect to Φ'_p) compatible with $\mathcal{F}$, there is a relabelling of a subtree of T which is an exact k-search-tree compatible with $\mathcal{F}$.*

PROOF. The proof is exactly the same as the one of Theorem 1 except in one case: One cannot always push, for u and v internal nodes of T, the part $l(uv)$ to $l(vu)$ in the partition T_u. Indeed, when $|l(uv)| \leq 1$, this could increase the number of parts of T_u with more than one element. In this case, we simply form a new tree T' by deleting the nodes of T which belong to the components of $T \setminus v$ not containing u. Now, v is a leaf of T', and we set $l(vu) = l(uv)^c$. Observe that T' is still compatible with $\mathcal{F}$. The reason for this is that $\bigcap \mathcal{F} = \emptyset$, hence one of its element is included in $l(uv)^c$. $\square$

It follows that Theorem 2 also holds for Φ'_p, and consequently for δ'_2. Now assume that T is a partitioning k-search-tree. Observe that we can assume that its internal labels have size at least 2, otherwise we just cut the branches as previously. This means that T is a caterpillar, i.e. a path with some attached leaves. We associate to every internal node x the bag $\Delta(T_x)$. This gives a path-decomposition of G in the usual sense.

The dual of path-width is the notion of *blockage*, introduced in [5]. Let us assume that $\mathcal{B}$ is a non-principal k-bramble in our sense, that is a set of pairwise intersecting subsets of edges, with overall empty intersection, and

containing a part of every partition $\mathcal{X} = \{X_1, \ldots, X_n\}$ with $\delta'_2(\mathcal{X}) \leq k$. We form a blockage as follows: A *k-cut* (V_1, V_2) is a pair of subsets of vertices with $|V_1 \cap V_2| \leq k$, $V_1 \cup V_2 = V$ and such that no edge of G joins $V_1 \setminus V_2$ to $V_2 \setminus V_1$. In a blockage $\mathcal{B}'$, either V_1 or V_2 must be chosen for every k-cut (V_1, V_2), with the additional *inclusion property* that if (V_1, V_2) and (W_1, W_2) are some k-cuts with $V_1 \subseteq W_1$, then $W_1 \in \mathcal{B}'$ implies $V_1 \in \mathcal{B}'$. The construction of $\mathcal{B}'$ is straightforward: if (V_1, V_2) is a k-cut, we let $X_1 := E(V_1) \setminus E(V_2)$, $X_2 := E(V_2) \setminus E(V_1)$, and we then list all the single edges $X_3, \ldots, X_n$ which belongs to $E(V_1 \cap V_2)$. This partition $\mathcal{X} = \{X_1, \ldots, X_n\}$ of E satisfies $\delta'_2(\mathcal{X}) \leq k$. So X_1 or X_2 belongs to $\mathcal{B}$. If $X_1 \in \mathcal{B}$, we choose V_2 in $\mathcal{B}'$, otherwise we choose V_1 in $\mathcal{B}'$. The inclusion property follows from the fact that the elements of $\mathcal{B}$ are pairwise intersecting.

6.5 Tree-width of matroids.

Let M be a matroid on ground set E with rank function r. We denote by r^c the submodular function such that $r^c(F) := r(F^c)$ for all subsets F of E and Φ the submodular partition function such that for any partition $\mathcal{X} = \{X_1, \ldots, X_l\}$,

$$\Phi(\mathcal{X}) = \Sigma_{r^c}(\mathcal{X}) - (l-1)r(E)$$

This function is submodular (Σ_{r^c} is submodular by Proposition 4) and gives the tree-width of matroids.

A *tree-decomposition* of M (see Hliněný and G. Whittle [2]) is given by a tree T whose nodes are labelled by the elements of a partition of E. A node u of T labelled F_0 thus corresponds to a partition $(F_0, \ldots, F_d)$ of E. Its *weight* is $\Sigma_{i=1}^{d} r^c(F_i) - (l-1)r(E)$. Partitioning search-tree on E with the weight function Φ thus correspond to tree-decompositions but the converse is not true. However, a tree-decomposition can be turned into a partitioning search-tree without increasing its width. Indeed, we can safely prune the empty labelled leaves and suppose u is either an internal node with a non-empty label or a leaf whose label is not a singleton. Let F_0 be its label. Attach $|F_0|$ new leaves to u and move the elements of F_0 to these leaves. The contribution of a new leaf labelled e to the weight of u is $r^c(e) - r(E) \leq 0$ and its weight is $r(e) = 1$. Since the width of a leaf labelled F is $r(F) \geq 1$, the width of this new tree-decomposition is at most the width of the previous one.

Non-principal brambles provide a dual notion to matroid tree-width.

References

[1] R. Diestel. *Graph theory*, volume 173. Springer-Verlag, 3rd edition, 2005.

[2] P. Hliněný and G. Whittle. Matroid tree-width. *European Journal of Combinatorics*, 27(7):1117–1128, 2006.

[3] S.-I. Oum and P. D. Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory Series B*, 96(4):514–528, 2006.

[4] B. A. Reed. Tree width and tangles: A new connectivity measure and some applications. *Surveys in Combinatorics*, 241:87–162, 1997.

[5] N. Robertson and P. D. Seymour. Graph Minors. I. Excluding a Forest. *Journal of Combinatorial Theory Series B*, 35(1):39–61, 1983.

[6] N. Robertson and P. D. Seymour. Graph Minors. X. Obstructions to Tree-Decomposition. *Journal of Combinatorial Theory Series B*, 52(2):153–190, 1991.

[7] P. D. 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.