Tangles are decided by weighted vertex sets

Jakob Kneip
Universität Hamburg

joint work with Christian Elbracht and Maximilian Teegen

November 16, 2020

$G = (V, E)$ a graph. A *separation* is a pair (A, B) of vertex sets with $A \cup B = V$ such that no edge of G has one endpoint in $A \setminus B$ and one in $B \setminus A$.

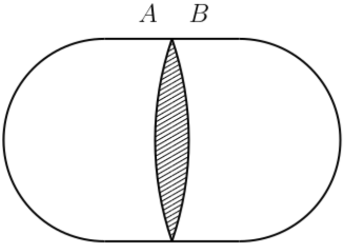

$G = (V, E)$ a graph. A *separation* is a pair (A, B) of vertex sets with $A \cup B = V$ such that no edge of G has one endpoint in $A \setminus B$ and one in $B \setminus A$.

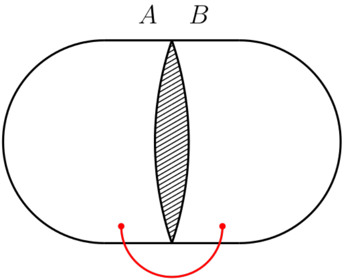

$G = (V, E)$ a graph. A *separation* is a pair (A, B) of vertex sets with $A \cup B = V$ such that no edge of G has one endpoint in $A \setminus B$ and one in $B \setminus A$.

The *order* of a separation is the size of its separator $A \cap B$.

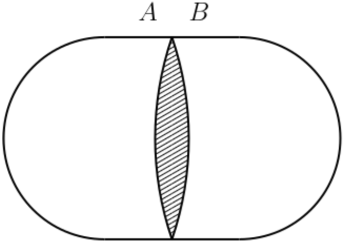

A k-*tangle* is a set τ of separations of order $< k$ such that:

A *k-tangle* is a set τ of separations of order $< k$ such that:

- τ contains exactly one of (A, B) and (B, A) for all (A, B) of order $< k$

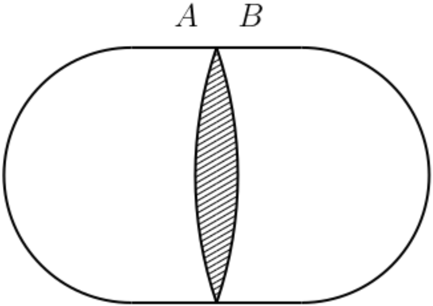

A k-*tangle* is a set τ of separations of order $< k$ such that:

- τ contains exactly one of (A, B) and (B, A) for all (A, B) of order $< k$
- if τ contains $(A_1, B_1), (A_2, B_2)$ and (A_3, B_3), then

$$G[A_1] \cup G[A_2] \cup G[A_3] \neq G.$$

A k-*tangle* is a set τ of separations of order $< k$ such that:

- τ contains exactly one of (A, B) and (B, A) for all (A, B) of order $< k$
- if τ contains $(A_1, B_1), (A_2, B_2)$ and (A_3, B_3), then

$$G[A_1] \cup G[A_2] \cup G[A_3] \neq G.$$

Motivation

A large cluster in a graph *orients* all the low order separations of a graph,

A k-*tangle* is a set τ of separations of order $< k$ such that:

- τ contains exactly one of (A, B) and (B, A) for all (A, B) of order $< k$
- if τ contains $(A_1, B_1), (A_2, B_2)$ and (A_3, B_3), then

$$G[A_1] \cup G[A_2] \cup G[A_3] \neq G.$$

Motivation

A large cluster in a graph *orients* all the low order separations of a graph, the second condition (*tangle property*) ensures that we point to something substantial.

Examples

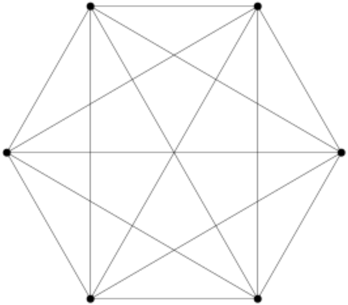

cliques

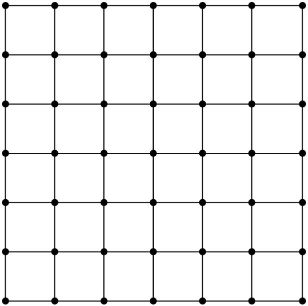

grids

Examples

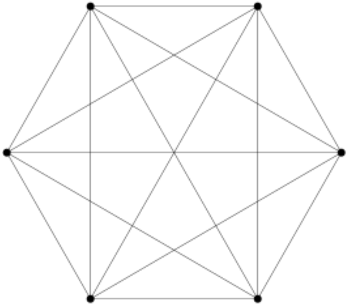

cliques

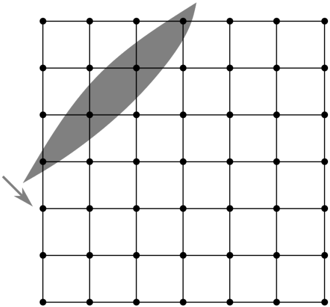

grids

Examples

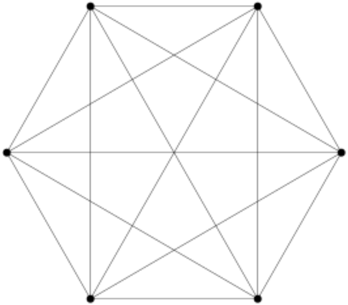

cliques

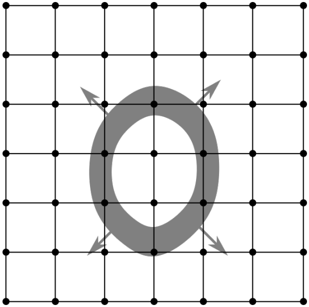

grids

The tangles in these examples have a set of vertices which determines their orientation by majority vote:

The tangles in these examples have a set of vertices which determines their orientation by majority vote:

there is a *decider set* $X \subseteq V$ such that for every separation (A, B) of order $< k$ we have

$$(A, B) \in \tau \text{ if and only if } |B \cap X| > |A \cap X|.$$

The tangles in these examples have a set of vertices which determines their orientation by majority vote:

there is a *decider set* $X \subseteq V$ such that for every separation (A, B) of order $< k$ we have

$$(A, B) \in \tau \text{ if and only if } |B \cap X| > |A \cap X|.$$

Do we always have such a decider set?

The tangles in these examples have a set of vertices which determines their orientation by majority vote:

there is a *decider set* $X \subseteq V$ such that for every separation (A, B) of order $< k$ we have

$$(A, B) \in \tau \text{ if and only if } |B \cap X| > |A \cap X|.$$

Do we always have such a decider set? *Maybe.*

We can show that *weighted deciders* exist:

Theorem (Elbracht, K, Teegen, 2020)

*Let $G = (V, E)$ be a finite graph and τ a k-tangle in G.
Then there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$,
where $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.*

We can show that *weighted deciders* exist:

Theorem (Elbracht, K, Teegen, 2020)

Let $G = (V, E)$ be a finite graph and τ a k-tangle in G.
Then there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$,
where $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.

How did we prove this?

First observation

The separations come with a natural partial order:

$$(A, B) \leqslant (C, D) :\Longleftrightarrow A \subseteq C \text{ and } B \supseteq D.$$

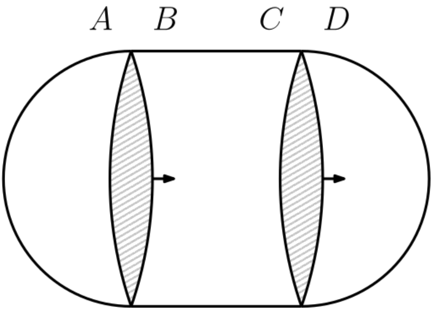

First observation

The separations come with a natural partial order:

$$(A, B) \leqslant (C, D) :\Longleftrightarrow A \subseteq C \text{ and } B \supseteq D.$$

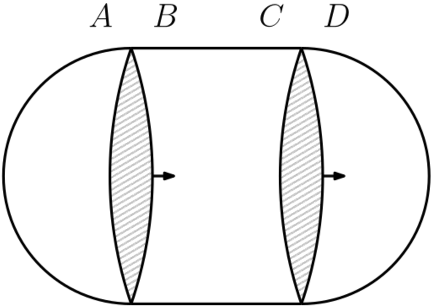

It suffices to find a weighted decider for the *maximal separations* of a k-tangle w.r.t. this partial order.

Second observation

Let (A, B) and (C, D) be two distinct maximal separations in a k-tangle. Consider the *separators* $A \cap B$ and $C \cap D$.

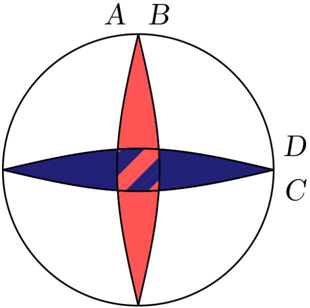

Second observation

Let (A, B) and (C, D) be two distinct maximal separations in a k-tangle.
Consider the *separators* $A \cap B$ and $C \cap D$.

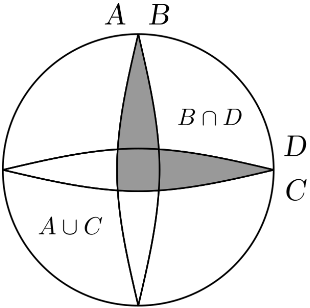

The separation $(A \cup C, B \cap D)$ cannot be of order $< k$.

Second observation

Let (A, B) and (C, D) be two distinct maximal separations in a k-tangle.
Consider the *separators* $A \cap B$ and $C \cap D$.
The separation $(A \cup C, B \cap D)$ cannot be of order $< k$.

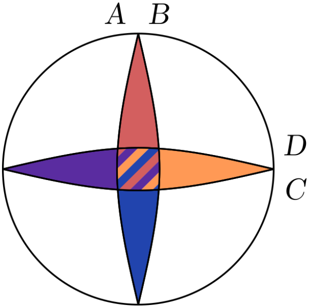

Taken together, the separator vertices are 'more often right than wrong':

$$(|B \cap (C \cap D)| + |D \cap (A \cap B)|) - (|A \cap (C \cap D)| + |C \cap (A \cap B)|) > 0$$

Idea: Find a weighting of the separators, and derive the weighting of the vertices from it.

Idea: Find a weighting of the separators, and derive the weighting of the vertices from it.

We can phrase this problem as a linear program:

- Enumerate the maximal separations of τ as $(A_1, B_1), \ldots, (A_m, B_m)$.

Idea: Find a weighting of the separators, and derive the weighting of the vertices from it.

We can phrase this problem as a linear program:

- Enumerate the maximal separations of τ as $(A_1, B_1), \ldots, (A_m, B_m)$.
- Define an $m \times m$ matrix M where m_{ij} specifies how well $A_j \cap B_j$ decides (A_i, B_i), by setting

$$m_{ij} := \underline{|B_i \cap (A_j \cap B_j)|} - \underline{|A_i \cap (A_j \cap B_j)|}\,.$$

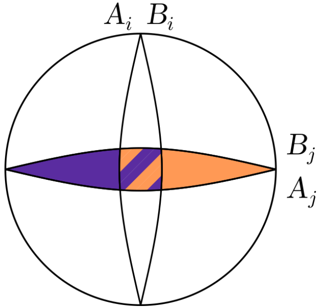

We can phrase this problem as a linear program:

- Enumerate the maximal separations of τ as $(A_1, B_1), \dots, (A_m, B_m)$.
- Define an $m \times m$ matrix M where m_{ij} specifies how well $A_j \cap B_j$ decides (A_i, B_i), by setting

$$m_{ij} := \underline{|B_i \cap (A_j \cap B_j)|} - \underline{|A_i \cap (A_j \cap B_j)|}\,.$$

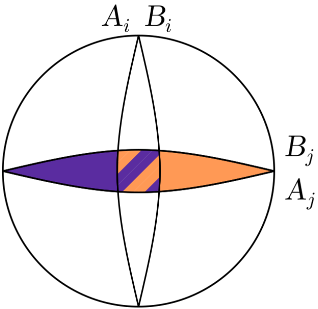

Now if x is a vector of weights for the separators, then

$$w(B_i) - w(A_i) = \sum_j m_{ij} \cdot x_j\,,$$

We can phrase this problem as a linear program:

- Enumerate the maximal separations of τ as $(A_1, B_1), \ldots, (A_m, B_m)$.
- Define an $m \times m$ matrix M where m_{ij} specifies how well $A_j \cap B_j$ decides (A_i, B_i), by setting

$$m_{ij} := \underline{|B_i \cap (A_j \cap B_j)|} - \underline{|A_i \cap (A_j \cap B_j)|}\,.$$

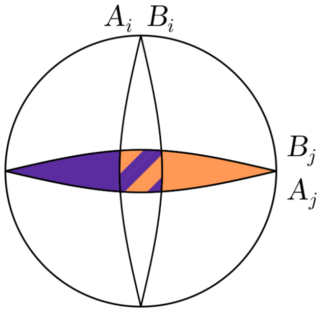

Now if x is a vector of weights for the separators, then

$$w(B_i) - w(A_i) = \sum_j m_{ij} \cdot x_j\,,$$

so Mx is the vector of the 'net scores' of the (A_i, B_i) in the weighting x.

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.
By the second observation we have $m_{ij} + m_{ji} > 0$, so $M + M^T$ has positive off-diagonal entries.

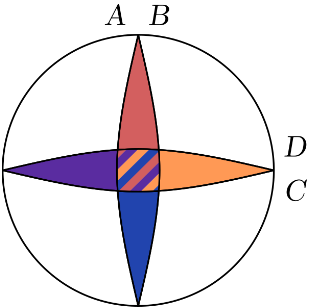

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.
By the second observation we have $m_{ij} + m_{ji} > 0$, so $M + M^T$ has positive off-diagonal entries.

Lemma (Farkas' Lemma)

Given A and b, either $Ax \geqslant b$ has a solution with $x \geqslant 0$, or there exists a $y \geqslant 0$ such that $A^T y \leqslant 0$ and $b^T y > 0$.

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.
By the second observation we have $m_{ij} + m_{ji} > 0$, so $M + M^T$ has positive off-diagonal entries.

Lemma (Farkas' Lemma)

Given A and b, either $Ax \geqslant b$ has a solution with $x \geqslant 0$, or there exists a $y \geqslant 0$ such that $A^T y \leqslant 0$ and $b^T y > 0$.

Apply Farkas with $A = M$ and $b = (1, \ldots, 1)^T$. Two possible outcomes:

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.
By the second observation we have $m_{ij} + m_{ji} > 0$, so $M + M^T$ has positive off-diagonal entries.

Lemma (Farkas' Lemma)

Given A and b, either $Ax \geqslant b$ has a solution with $x \geqslant 0$, or there exists a $y \geqslant 0$ such that $A^T y \leqslant 0$ and $b^T y > 0$.

Apply Farkas with $A = M$ and $b = (1, \ldots, 1)^T$. Two possible outcomes:

- there is $x \geqslant 0$ with $Mx \geqslant (1, \ldots, 1)^T > 0$.

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.
By the second observation we have $m_{ij} + m_{ji} > 0$, so $M + M^T$ has positive off-diagonal entries.

Lemma (Farkas' Lemma)

Given A and b, either $Ax \geqslant b$ has a solution with $x \geqslant 0$, or there exists a $y \geqslant 0$ such that $A^T y \leqslant 0$ and $b^T y > 0$.

Apply Farkas with $A = M$ and $b = (1, \ldots, 1)^T$. Two possible outcomes:

- there is $x \geqslant 0$ with $Mx \geqslant (1, \ldots, 1)^T > 0$.
- there is $y \geqslant 0$ with $M^T y \leqslant 0$ and $y \neq 0$.

We need to find a weight vector $x \geqslant 0$ with $Mx > 0$.
By the second observation we have $m_{ij} + m_{ji} > 0$, so $M + M^T$ has positive off-diagonal entries.

Lemma (Farkas' Lemma)

Given A and b, either $Ax \geqslant b$ has a solution with $x \geqslant 0$, or there exists a $y \geqslant 0$ such that $A^T y \leqslant 0$ and $b^T y > 0$.

Apply Farkas with $A = M$ and $b = (1, \ldots, 1)^T$. Two possible outcomes:

- there is $x \geqslant 0$ with $Mx \geqslant (1, \ldots, 1)^T > 0$.
- there is $y \geqslant 0$ with $M^T y \leqslant 0$ and $y \neq 0$. But then

$$0 \leqslant (M + M^T)y \leqslant My\,.$$

Theorem (Elbracht, K, Teegen, 2020)

Let $G = (V, E)$ be a finite graph and τ a k-tangle in G.
Then there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$,
where $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.

Theorem (Elbracht, K, Teegen, 2020)

*Let $G = (V, E)$ be a finite graph and τ a k-tangle in G.
Then there exists a function $w\colon V \to \mathbb{N}$ such that a separation (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$,
where $w(U) := \sum_{u \in U} w(u)$ for $U \subseteq V$.*

We never made use of the graph's edges,
so the same result holds for k-tangles of hypergraphs.

Similar to tangles, we can define k-*edge-tangles* of a graph $G = (V, E)$:

Similar to tangles, we can define k-*edge-tangles* of a graph $G = (V, E)$: For a *cut* (A, B), a bipartition of V, we define as *order* the number of cross-edges $|E(A, B)|$.

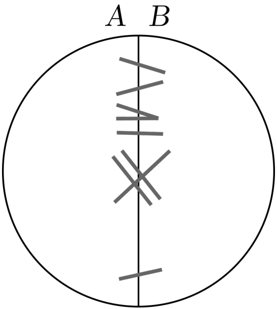

Similar to tangles, we can define k*-edge-tangles* of a graph $G = (V, E)$: For a *cut* (A, B), a bipartition of V, we define as *order* the number of cross-edges $|E(A, B)|$.

A k*-edge-tangle* is a set τ of cuts of order $< k$ such that:

- For every cut (A, B) of order $< k$ either (A, B) or (B, A) is in τ.

Similar to tangles, we can define k*-edge-tangles* of a graph $G = (V, E)$: For a *cut* (A, B), a bipartition of V, we define as *order* the number of cross-edges $|E(A, B)|$.

A k*-edge-tangle* is a set τ of cuts of order $< k$ such that:

- For every cut (A, B) of order $< k$ either (A, B) or (B, A) is in τ.
- If τ contains (A_1, B_1),(A_2, B_2) and (A_3, B_3) then

$$A_1 \cup A_2 \cup A_3 \neq V.$$

Similar to tangles, we can define *k-edge-tangles* of a graph $G = (V, E)$: For a *cut* (A, B), a bipartition of V, we define as *order* the number of cross-edges $|E(A, B)|$.

A *k-edge-tangle* is a set τ of cuts of order $< k$ such that:

- For every cut (A, B) of order $< k$ either (A, B) or (B, A) is in τ.
- If τ contains (A_1, B_1),(A_2, B_2) and (A_3, B_3) then

$$A_1 \cup A_2 \cup A_3 \neq V.$$

- For every (A, B) in τ there are at least k edges incident with vertices in B.

Theorem (Elbracht, K, Teegen, 2020)

Let $G = (V, E)$ be a finite (multi-)graph and τ a k-edge-tangle in G. Then there exists a function $w \colon V \to \mathbb{N}$ such that a cut (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$.

Theorem (Elbracht, K, Teegen, 2020)

Let $G = (V, E)$ be a finite (multi-)graph and τ a k-edge-tangle in G. Then there exists a function $w\colon V \to \mathbb{N}$ such that a cut (A, B) of G of order $< k$ lies in τ if and only if $w(A) < w(B)$.

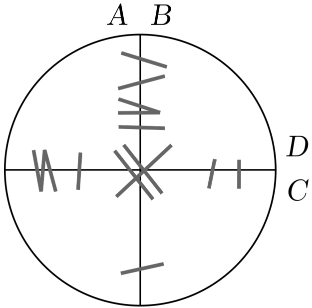

The edge-tangle result does not extend to hypergraphs.

The edge-tangle result does not extend to hypergraphs.

Consider the hypergraph $H = (V, E)$ where

- V is the set of 3-element subsets of $\{1, \ldots, 7\}$;

The edge-tangle result does not extend to hypergraphs.

Consider the hypergraph $H = (V, E)$ where

- V is the set of 3-element subsets of $\{1, \ldots, 7\}$;
- for each $i \in \{1, \ldots, 7\}$ we have a hyperedge $e_i = \{v \in V \mid i \in v\}$.

The edge-tangle result does not extend to hypergraphs.

Consider the hypergraph $H = (V, E)$ where

- V is the set of 3-element subsets of $\{1, \ldots, 7\}$;
- for each $i \in \{1, \ldots, 7\}$ we have a hyperedge $e_i = \{v \in V \mid i \in v\}$.

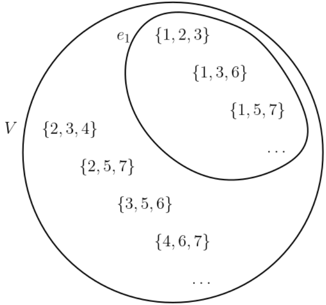

Consider the hypergraph $H = (V, E)$ where

- V is the set of 3-element subsets of $\{1, \ldots, 7\}$;
- for each $i \in \{1, \ldots, 7\}$ we have a hyperedge $e_i = \{v \in V \mid i \in v\}$.

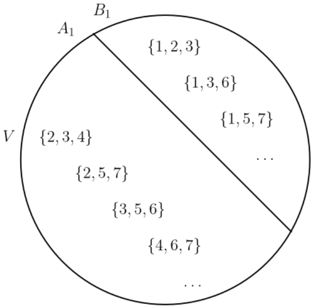

In this hypergraph we have an edge-tangle

$$\tau := \{(A, B) \text{ of order } < 7 \mid \; B \supseteq e_i \text{ for some } e_i\} \; .$$

Consider the hypergraph $H = (V, E)$ where

- V is the set of 3-element subsets of $\{1, \ldots, 7\}$;
- for each $i \in \{1, \ldots, 7\}$ we have a hyperedge $e_i = \{v \in V \mid i \in v\}$.

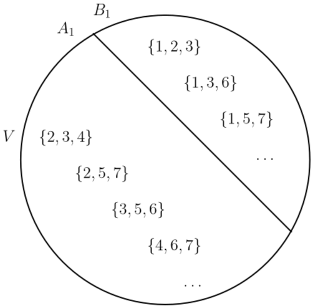

In this hypergraph we have an edge-tangle

$$\tau := \{(A, B) \text{ of order} < 7 \mid \ B \supseteq e_i \text{ for some } e_i\} \ .$$

This edge tangle has no weighted decider: consider

$$\sum_{1 \leqslant i \leqslant 7} w(B_i) - w(A_i) \, .$$

End

Thank you!