A Unified Theory of Structural Tractability for Constraint Satisfaction Problems

David Cohen[1] Peter Jeavons[2] Marc Gyssens

[1]Department of Computer Science, Royal Holloway, University of London, UK
d.cohen@rhul.ac.uk

[2]Computing Laboratory, University of Oxford, UK
peter.jeavons@comlab.ox.ac.uk

[3]Department WNI, Hasselt University, Agoralaan, B-3590 Diepenbeek, Belgium
Marc.Gyssens@uhasselt.be

Dagstuhl Meeting on the Complexity of CSP, October 2006

Outline

Guarded Blocks.

Definition

A guarded block of a hypergraph *H* is a pair $\langle \lambda, \chi \rangle$ where the guard, λ, is a subset of the hyperedges of *H*, and the block, χ, is a subset of the vertices of the guard.

Guarded Decompositions

Definition

For any CSP instance P, and any guarded block b of $\sigma(P)$, the constraint generated by P on b is the projection onto the block of b join of all the constraints of P whose scopes are elements of the guard of b.

Definition

A set of guarded blocks Ξ of a hypergraph H is a guarded decomposition if for every CSP instance P with structure H, the instance generated by P on Ξ, is solution-equivalent to P.

Example Guarded Decompositions.

Example

For any hypergraph $H = \langle V, E \rangle$ there are two trivial guarded decompositions.

The set $\{\langle \{e\}, e \rangle \mid e \in E\}$ is a guarded decomposition.

The set $\{\langle E, V \rangle\}$ is also a guarded decomposition.

Definition

- A guarded block $\langle \lambda, \chi \rangle$ of H covers a hyperedge e of H if $e \subseteq \chi$.
- A set of guarded blocks Ξ of H is called a guarded cover if each hyperedge of H is covered by some guarded block of Ξ.
- A set of guarded blocks Ξ of a hypergraph H is called a complete guarded cover for H if each hyperedge e of H occurs in the guard of some guarded block of Ξ which covers e.

Theorem

A set of guarded blocks Ξ of a hypergraph H is a guarded decomposition of H if and only if it is a complete guarded cover for H.

Structural Decompositions

Definition

A class of CSP instances is called structural if it is defined by a set of hypergraphs.

- A guarded decomposition for a hypergraph H associates any CSP instance whose structure is H with a solution-equivalent instance having a different structure, which may be easier to solve.
- Structural classes of CSP instances are the class of all instances whose structure has a guarded decomposition of a certain kind.

Tractable Structural Classes

The most important structural classes are those which are *tractable*, in the following sense.

Definition

A class of CSP instances is called tractable if there exists

- a polynomial-time algorithm to decide membership; and
- a polynomial-time algorithm to solve all instances.

Finding Tractable Structural Classes

Tractable discovery For any given hypergraph it must be possible to decide in polynomial time whether there exists a guarded decomposition of the type we are considering, and, in the process, obtain such a decomposition if it exists.

Tractable construction Given such a guarded decomposition, it must be possible to generate each of the new constraints in the corresponding solution-equivalent instance in polynomial time.

Tractable solution Given such a solution-equivalent instance, it must be possible to solve the resulting instance in polynomial time.

Tractable Construction

Definition

The width of a set of guarded blocks is the maximum number of hyperedges in any of its guards.

For any fixed value of k, the class of guarded decompositions of width at most k has the tractable-construction property.

Tractable Solution

One way to achieve the tractable solution property is to ensure that the structure of decomposed instances is *acyclic*.

Definition

A join tree of a set of guarded blocks Ξ of H is a tree, T, whose nodes are the elements of Ξ, where the set of nodes of T for which any vertex x occurs in the block induces a (connected) subtree of T.
A set of guarded blocks is acyclic if it has a join tree.

Any class of *acyclic* guarded decompositions has the tractable-solution property.

Completing Acyclic Guarded Covers

Theorem

If the set of guarded blocks Ξ is an acyclic guarded cover for H then H has a complete acyclic guarded cover of the same width.

It is because of these results, codified here, that the search for structural decompositions has really been the search for acyclic guarded covers.

Tractable Discovery

- For any fixed choice of k, any CSP instance whose structure has an acyclic guarded cover with width at most k can be solved in polynomial time.
- To define a tractable structural class we still need some way to determine in polynomial time whether a given instance has such a guarded cover or not.
- Luckily there are simple extra conditions on guarded covers which are sufficient to enable us to do this.

Components

Definition

Let $H = \langle V, E \rangle$ be a hypergraph and $\chi \subseteq V$ be any subset of vertices.
A pair of vertices x, y is χ-connected if there is a sequence of hyperedges $e_0, \ldots, e_m$ such that $x \in e_0 - \chi$, $y \in e_m - \chi$ and $e_i \cap e_{i+1} \not\subseteq \chi$, for $i = 0, \ldots, m - 1$.
A χ-component of H is a maximal non-empty set of vertices C such that each pair of vertices in C is χ-connected.

Properties of Acyclic Guarded Covers

Definition

Let T be a tree of guarded blocks of H. For any pair of adjacent nodes n and n', define the n'-branch of T with respect to n, $\mathrm{br}_n(n')$, to be the nodes whose path to n includes n'.
The vertices of $\mathrm{br}_n(n')$, denoted $\chi(\mathrm{br}_n(n'))$, are the vertices in the blocks of the elements of $\mathrm{br}_n(n')$ which are not in $\chi(n)$.

Lemma

Let H be a hypergraph. Any join tree T of a guarded cover satisfies:

- JT1 *For every arc $\langle n, n' \rangle$ of T, and every hyperedge e, if $e \cap \chi(\mathrm{br}_n(n')) \neq \emptyset$, then e is covered by some node of $\mathrm{br}_n(n')$;*
- JT2 *For every arc $\langle n, n' \rangle$ of T, $\chi(\mathrm{br}_n(n'))$ is a union of $\chi(n)$-components.*

Decomposition Trees

Definition

A decomposition tree, T, of a hypergraph H is a rooted join tree of a set of guarded blocks of H satisfying the following conditions:

- **DT1** For every arc $\langle n, n' \rangle$ of T, and every edge e of H, if $e \cap \chi(\mathrm{br}_n(n')) \neq \emptyset$, then e is covered by some node of $\mathrm{br}_n(n')$;
- **DT2** For every arc $\langle n, n' \rangle$ of T, there exists a single $\chi(n)$-component, $C_{\langle n,n' \rangle}$, of H such that $\chi(\mathrm{br}_n(n')) = C_{\langle n,n' \rangle}$.

Theorem

Let $H = \langle V, E \rangle$ be a hypergraph, and let β be a binary predicate on pairs of guarded blocks of H. In $O(|\beta||E||V|^2)$ time it is possible to decide whether H has an acyclic guarded cover Ξ which has a decomposition tree where every arc satisfies β, and to construct such a cover if it exists.

Unbroken Components

Definition

A guarded block b of a hypergraph H has unbroken components if each $\chi(b)$-component of H meets (has non-empty intersection with) at most one $(\cup\lambda(b))$-component of H.

Respecting Labels

Definition

Let λ be any set of hyperedges of a hypergraph $H = \langle V, E \rangle$.
We define the label, $L_\lambda(v)$, of any vertex $v \in (\cup\lambda)$ to be a pair,

- the first component is the set of $(\cup\lambda)$-components which meet a hyperedge containing v,
- the second component is the the set of hyperedges of λ which contain v.

Definition

We say that a guarded block $\langle \lambda, \chi \rangle$ respects labels if the vertices outside of the block have different labels from those inside the block.

Spread Cut Decompositions

Lemma

For any fixed k and any hypergraph $H = \langle V, E \rangle$, the set of guarded blocks of H with width k which have unbroken components and respect labels can be enumerated in polynomial time in the size of H.

Definition

A spread-cut decomposition of H is an acyclic guarded cover Ξ with a decomposition tree, where every guarded block in Ξ has unbroken components and respects labels.

Corollary

For any fixed k, the class of CSP instances whose structure has a spread cut decomposition of width at most k is a tractable structural class.

The Decomposition Hierarchy

Spread Cut decompositions strongly generalise:

- Hinge-tree decompositions
- Biconnected Component Decompositions
- Cycle-hypercutset decompositions
- Cycle cutset decompositions.

Spreadcuts and Hypertrees

- Hypertrees generalise pure query decompositions. Spread cuts do not.
- Spread cuts can have arbitrarily smaller width than pure query decompositions on some hypergraphs. This is not known for hypertrees.
- Hypertrees are within a factor of three of the (optimal) acyclic cover width.
- Spread Cuts can have significantly smaller width than hypertrees.

Where's Your Guard Going!

- The descendant condition (special condition) on hypertrees means that the guard of any block can only meet the blocks of one branch of a hypertree.
- There is no such condition on Spread cuts.
- The requirement for a decomposition tree in the spread cut definition is very strong and prevents spread cuts from generalising pure query decompositions.
- Hypertrees do not require a decomposition tree. They have a normalisation result which reduces arbitrary hypertrees to those having a decomposition tree. Hence they generalise pure query decompositions.
- We need a more general form of spread cut which does not require a decomposition tree but instead has a normalisation result similar to that for hypertrees.

Smaller width than hypertrees

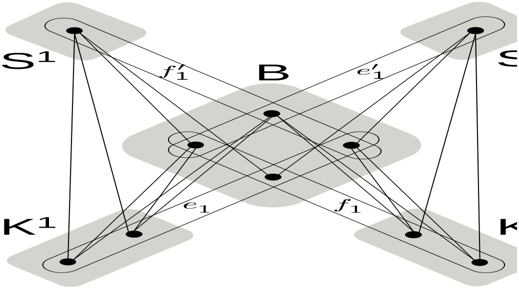

This hypergraph is the first of a family that have

- Optimal acyclic cover width $2n + 1$
- Spread cut width $2n + 1$
- Hypertree width $3n$
- Pure query width $3n$

A Spreadcut Decomposition

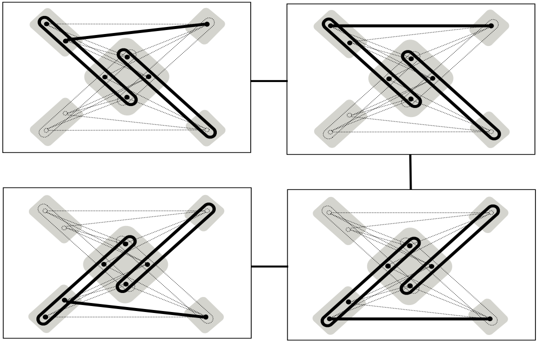

Summary

- All decompositions (except fractional edge covers) known are acyclic guarded covers. These can be motivated by a few simplifying assumptions.
- There is a very concise algorithm to find acyclic guarded covers which only relies on the number of allowable guarded blocks available for a hypergraph.
- This restriction still allows us to define spread cuts which are incomparable with hypertrees.

Open Questions

- Can we relax the definition of a spread cut to make normalisation work and to make spread cuts++ generalise pure query decompositions.
- Can hypertrees have arbitrarily smaller width than pure query decompositions.
- Do (normalised) hypertrees have a polynomial bound on the number of permitted guarded blocks for any hypertree (and are these then effectively enumerable).