Relaxation-Aware Heuristics for Exact Optimization in Graphical Models

Fulya Trösser [1], Simon de Givry [1], George Katsirelos [2]

[1]INRA MIA Toulouse

[2]INRA MIA Paris

5 December 2019

Journée AIGM 2019

Overview

Introduction

- Exact solvers for optimization problems on graphical models (CFN, MRF etc), typically use branch-and-bound.
- Two main factors: quality of the bound at each node and branching heuristics.
- Virtual Arc Consistency (VAC) algorithm: high quality bounds, but at a significant cost

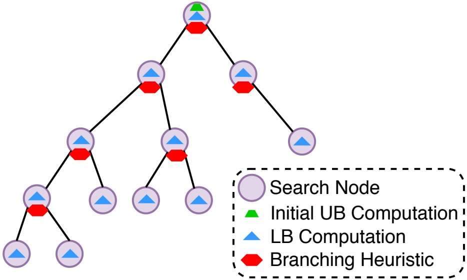

Introduction

- Branching heuristics ignore the information that VAC produces on the linear relaxation of the problem
- Branching heuristic may make decisions that are clearly ineffective
- By eliminating these ineffective decisions, we significantly reduce the size of the branch-and-bound tree.

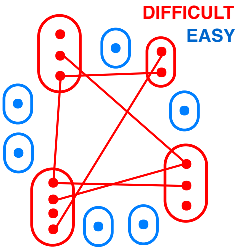

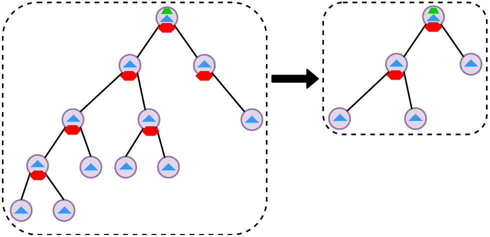

Constraint Satisfaction Problem (CSP)

A *Constraint Satisfaction Problem* [Cooper and Schiex, 2004] is a triple $\langle X, D, C \rangle$.

- X: set of n variables $X = \{1, \ldots, n\}$.
- D: set of domains $D = \{D_i : i \in X\}$.
- C: set of constraints.

A tuple t

is a *solution* iff it satisfies all the constraints in C.

Cost Function

A *cost function* is defined over the scope S of the constraint c_S to which it corresponds. It associates a cost to each tuple $t \in \ell(S)$.

- $c_\varnothing$: the nullary cost function = constant cost.
- c_i: the unary cost function on variable i.
- c_{ij}: the binary cost function on variables i and j.

x	c_x
a	0
b	2

y	c_y
a	1
b	0

x	y	c_{xy}
a	a	0
a	b	1
b	a	0
b	b	3

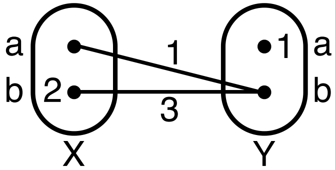

Weighted Constraint Satisfaction Problem

Weighted Constraint Satisfaction Problem (WCSP)

is a quadruple $\langle X, D, C, m \rangle$ where C is a set of cost functions and m is the *upper bound* [Cooper et al., 2010].

Find a solution such that the sum

$$c_\varnothing + \sum_{i \in X} c_i + \sum_{ij \in X^2} c_{ij}$$

- is minimized,
- is less than the upper bound m.

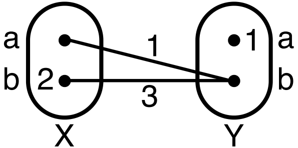

Levels of Local Consistency

Node Consistency (NC)

A WCSP is *Node Consistent* if for any variable $i \in \{1, \ldots, n\}$, $\exists a \in D_i$ such that $c_i(a) = 0$

(Soft) Arc Consistency (SAC)

A binary WCSP is *Arc Consistent* if for all $c_{xy} \in C$ we have:

$$\forall a \in D_x,\ \exists b \in D_y \text{ such that } c_{xy}(a, b) = 0.$$

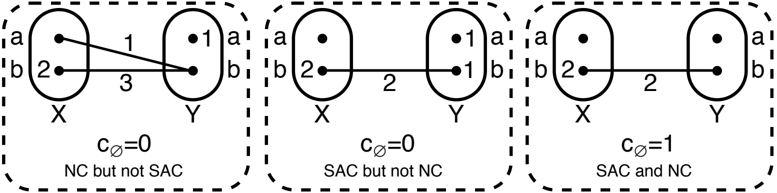

Levels of Local Consistency

Bool(P)

Values with non-zero unary cost ⇒ *REMOVED*
Assignments with non-zero binary cost ⇒ *FORBIDDEN*

Virtual Arc Consistency

A WCSP P is *virtual arc consistent* (VAC) if $Bool(P)$ is arc consistent.

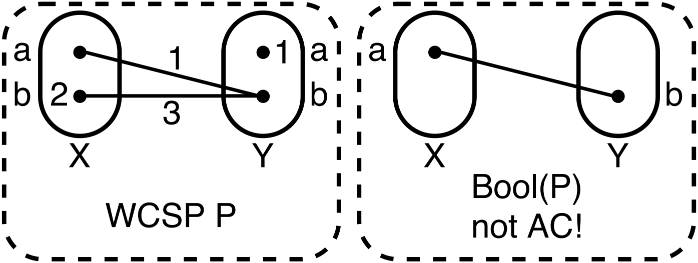

Strict Arc Consistency vs VAC-integrality

Strict AC [Savchynskyy et al., 2013]

- a **unique** value $a \in D_i$ such that $c_i(a) = 0$
- $\forall c_{ij} \in C$, a **unique** tuple $\{(i, a), (j, b)\}$ which satisfies $c_{ij}(a, b) = 0$

VAC-integrality

- a **unique** value $a \in D_i$ such that $c_i(a) = 0$
- $\forall c_{ij} \in C$, **at least one** tuple $\{(i, a), (j, b)\}$ which satisfies $c_{ij}(a, b) + c_j(b) = 0$.

Strict Arc Consistency vs VAC-integrality

Strict AC

- **unique** $a \in D_i$ s.t. $c_i(a) = 0$
- $\forall c_{ij} \in C$, **unique** $\{(i,a),(j,b)\}$ s.t. $c_{ij}(a,b) = 0$

VAC-integrality

- **unique** $a \in D_i$ s.t. $c_i(a) = 0$
- $\forall c_{ij} \in C$, **at least one** $\{(i,a),(j,b)\}$ s.t. $c_{ij}(a,b) + c_j(b) = 0$.

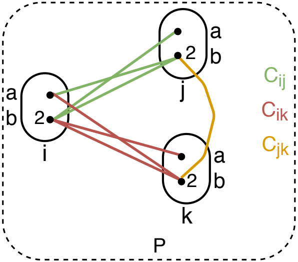

strict AC:
{ i }
strict* AC:
{ i, j, k }

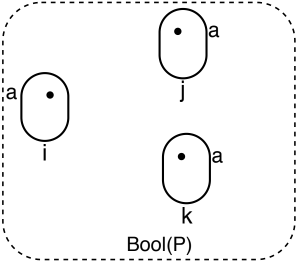

They all have a UNIQUE value in their domain in Bool(P)

Variable Ordering Heuristic

- It makes sense to AVOID branching on VAC-integral variables, since the dual bound will not be improved.

- So, we classify the variables with respect to their domain sizes in *Bool*(P): those with singleton domains, and the rest.
- Then, we give priority to the variables with more than one value.

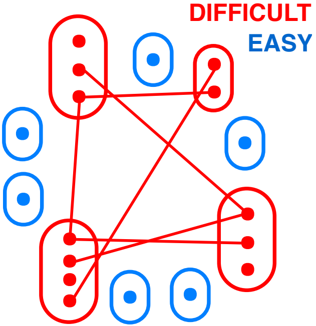

Heuristic for UB Inspired by CombiLP

- [Haller et al., 2018] develop an approach called CombiLP
- Subproblem is constructed by fixing the variables that have the same value in the incumbent and in the current relaxation.
- Following this approach, we propose a primal heuristic which runs in preprocessing.

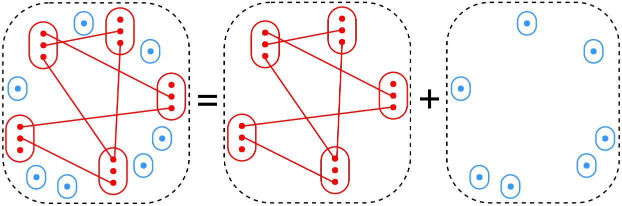

Relaxation-Aware Sub-Problem Search (RASPS)

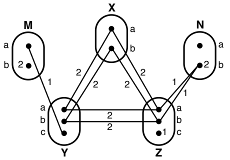

$\theta = 1$

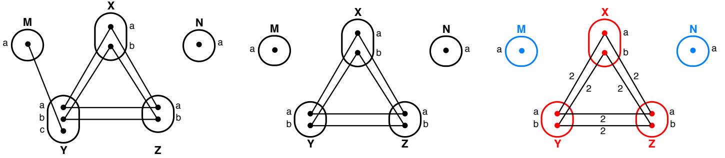

$\theta = 2$

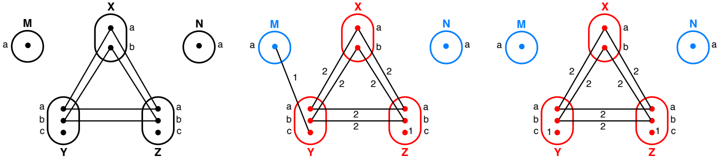

RASPS

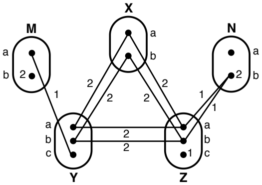

Figure: WCSP P. OPT $= \{a, a, b, c, a\}$ with total cost 1.

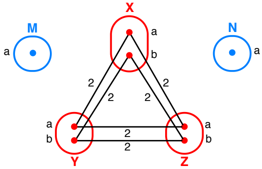

Figure: $\theta = 1$, $\{a, a, b, a, a\}$ with total cost 2.

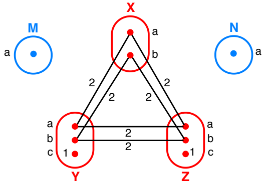

Figure: $\theta = 2$, $\{a, a, b, c, a\}$ with total cost 1.

RASPS

- θ_i: Threshold at iteration i of VAC
- r_i: Ratio of VAC-integral AC variables
- $\alpha_i = r_i/\theta_i$

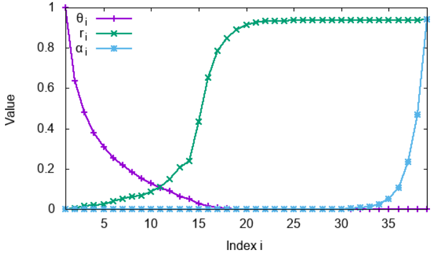

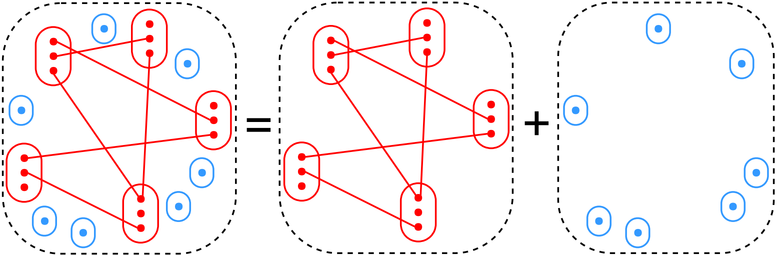

Experiments

- ToulBar2 [Cooper et al., 2010]: an open-source exact solver for cost function networks that solves various combinatorial optimization problems.
- CombiLP [Haller et al., 2018]: an open-source algorithm for energy minimization of graphical models which uses ToulBar2 as its internal combinatorial solver.
- 431 instances from various benchmarks including
  - 21 instances from CPD
  - 30 instances from worms
- Time limit: 3600 seconds (36000 for CPD)

Numerical Results for Worms

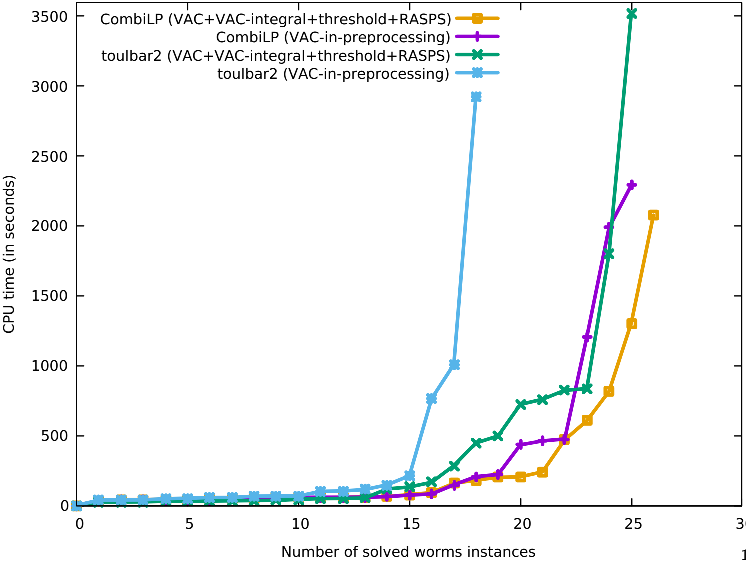

Numerical Results for CPD

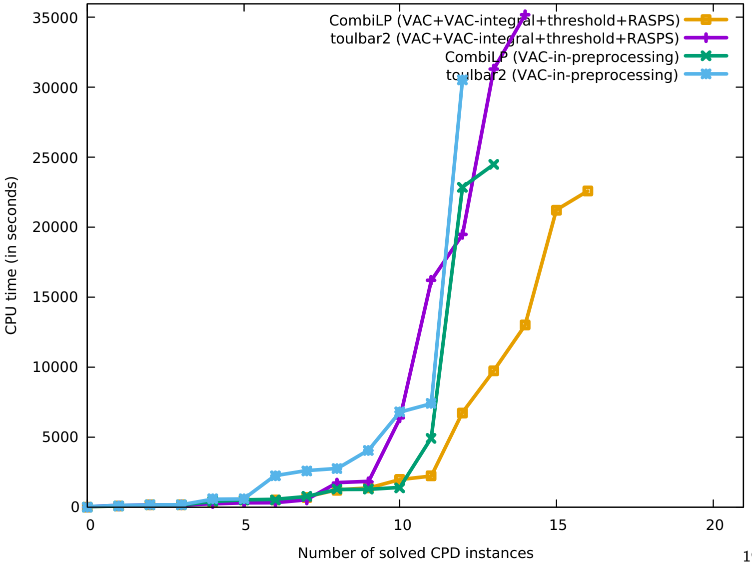

Scatter Plots

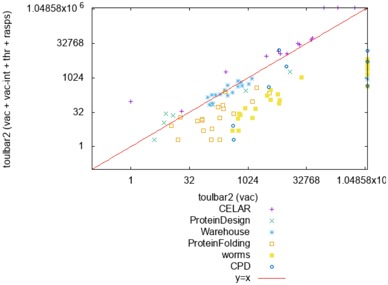

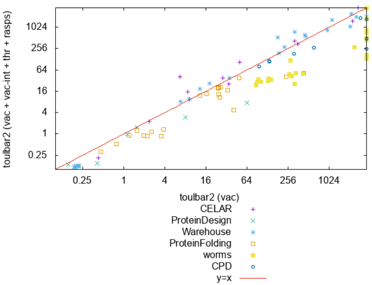

Future Work

- Further exploring the connections between WCSP and ILP
- Finding ways of making VAC more useful more often in WCSP (so it could potentially become the default option)
- Testing these heuristics with Bayesian Network Structure Learning instances

References

Cooper, Martin and Schiex, Thomas (2004)
Arc consistency for soft constraints
Artificial Intelligence 154(1-2), 199 – 227.

Cooper, Martin C and De Givry, Simon and Sánchez, Martı and Schiex, Thomas and Zytnicki, Matthias and Werner, Tomáš (2010)
Soft arc consistency revisited
Artificial Intelligence 174(7-8), 449 – 478.

Danna, Emilie and Rothberg, Edward and Le Pape, Claude (2005)
Exploring relaxation induced neighborhoods to improve MIP solutions
Mathematical Programming 102(1), 71 – 90.

Haller, Stefan and Swoboda, Paul and Savchynskyy, Bogdan (2018)
Exact MAP-Inference by Confining Combinatorial Search with LP Relaxation
AAAI.

References

Savchynskyy, Bogdan and Kappes, Jörg Hendrik and Swoboda, Paul and Schnörr, Christoph (2013)
Global MAP-optimality by shrinking the combinatorial search area with convex relaxation
Advances in Neural Information Processing Systems 1950 – 1958.

Thank You