A Constraint Composite Graph-Based ILP Encoding of the Boolean Weighted CSP

Hong Xu Sven Koenig T. K. Satish Kumar
hongx@usc.edu, skoenig@usc.edu, tkskwork@gmail.com

August 29, 2017

University of Southern California
the 23rd International Conference on Principles and Practice of Constraint Programming (CP 2017)
Melbourne, Victoria, Australia

Executive Summary

- Constraint Composite Graphs (CCGs) are “lifted” representations of Weighted Constraint Satisfaction Problems (Weighted CSPs, WCSPs).
- The Integer Linear Programming (ILP) encoding based on the CCG of a WCSP allows one to find an optimal solution of the WCSP faster than the ILP encoding directly based on the WCSP itself.

Agenda

- The Weighted Constraint Satisfaction Problem (WCSP)
- The Constraint Composite Graph (CCG)
- ILP Encodings of WCSPs
- Conclusion

Agenda

- The Weighted Constraint Satisfaction Problem (WCSP)
- The Constraint Composite Graph (CCG)
- ILP Encodings of WCSPs
- Conclusion

Weighted Constraint Satisfaction Problem (WCSP): Motivation

Many real-world problems can be solved using the WCSP:

- RNA motif localization (Zytnicki et al. 2008)
- Communication through noisy channels using Error Correcting Codes in Information Theory (Yedidia et al. 2003)
- Medical and mechanical diagnostics (Milho et al. 2000; Muscettola et al. 1998)
- Energy minimization in Computer Vision (Kolmogorov 2005)
- $\cdots$

Weighted Constraint Satisfaction Problem (WCSP)

- N variables $\mathcal{X} = \{X_1, X_2, \ldots, X_N\}$.
- Each variable X_i has a discrete-valued domain $\mathcal{D}(X_i)$.
- M weighted constraints $\mathcal{C} = \{C_1, C_2, \ldots, C_M\}$.
- Each constraint C_i specifies the weight for each assignment a of values to a subset $S(C_i)$ of the variables (denoted by $E_{C_i}(a|S(C_i))$).
- Find an optimal assignment a of values to these variables so as to minimize the total weight: $\sum_{i=1}^{M} E_{C_i}(a|S(C_i))$.
- A Boolean WCSP is a WCSP in which the domain size of every variable is 2.
- Known to be NP-hard.

Boolean WCSP Example

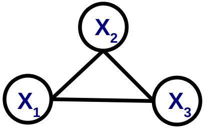

X_1	
0	0.7
1	0.2

X_2	
0	0.3
1	0.8

X_3	
0	0.1
1	1.0

X_1 \ X_2	0	1
0	0.5	0.6
1	0.7	0.3

X_2 \ X_3	0	1
0	0.6	1.3
1	1.0	1.1

X_1 \ X_3	0	1
0	0.4	0.9
1	0.7	0.8

$$E(X_1, X_2, X_3) = E_1(X_1) + E_2(X_2) + E_3(X_3) + E_{12}(X_1, X_2) + E_{13}(X_1, X_3) + E_{23}(X_2, X_3)$$

WCSP Example: Evaluate the Assignment $X_1 = 0, X_2 = 0, X_3 = 1$

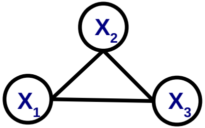

X_1	
0	0.7
1	0.2

X_2	
0	0.3
1	0.8

X_3	
0	0.1
1	1.0

X_1 \ X_2	0	1
0	0.5	0.6
1	0.7	0.3

X_2 \ X_3	0	1
0	0.6	1.3
1	1.0	1.1

X_1 \ X_3	0	1
0	0.4	0.9
1	0.7	0.8

$$E(X_1 = 0, X_2 = 0, X_3 = 1) = 0.7 + 0.3 + 1.0 + 0.5 + 1.3 + 0.9 = 4.7$$

(This is not an optimal solution.)

WCSP Example: Evaluate the Assignment $X_1 = 1, X_2 = 0, X_3 = 0$

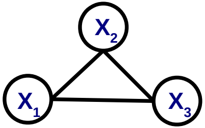

X_1	
0	0.7
1	0.2

X_2	
0	0.3
1	0.8

X_3	
0	0.1
1	1.0

X_1 \ X_2	0	1
0	0.5	0.6
1	0.7	0.3

X_2 \ X_3	0	1
0	0.6	1.3
1	1.0	1.1

X_1 \ X_3	0	1
0	0.4	0.9
1	0.7	0.8

$$E(X_1 = 1, X_2 = 0, X_3 = 0) = 0.2 + 0.3 + 0.1 + 0.7 + 0.6 + 0.7 = 2.6$$

This is an optimal solution. Using brute force, it requires exponential time to find.

Agenda

- The Weighted Constraint Satisfaction Problem (WCSP)
- The Constraint Composite Graph (CCG)
- ILP Encodings of WCSPs
- Conclusion

Two Forms of Structure in a WCSP

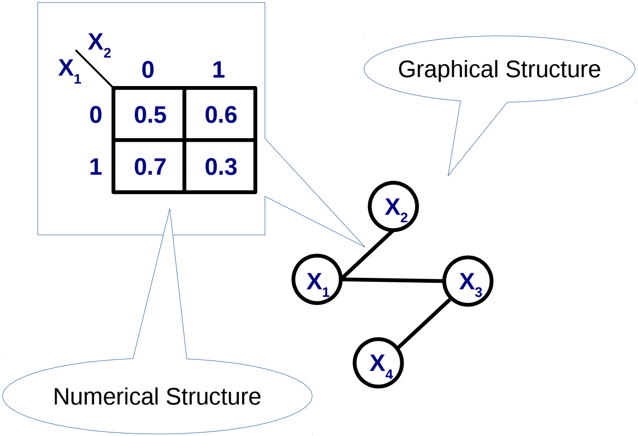

- Graphical: Which variables are in which constraints?
- Numerical: How does each constraint relate the variables in it?

How can we exploit both forms of structure computationally?

Minimum Weighted Vertex Cover (MWVC)

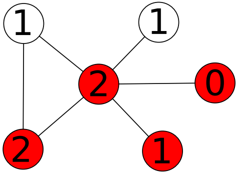

(a) ✗

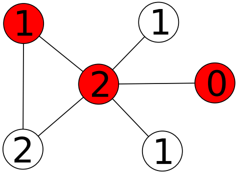

(b) ✓

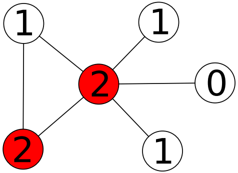

(c) ✗

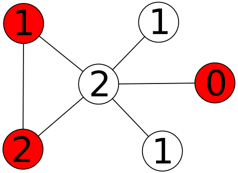

(d) ✗

Each vertex is associated with a non-negative weight. In a minimum weighted vertex cover (MWVC), the sum of the weights on the vertices in the VC is minimized.

Projection of a Minimum Weighted Vertex Cover (MWVC) onto an Independent Set

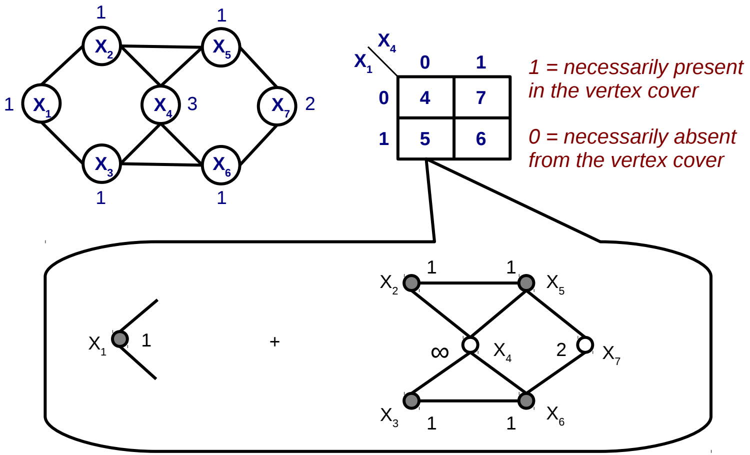

(Kumar 2008, Fig. 2)

Projection of an MWVC onto an Independent Set

Assuming Boolean variables in WCSPs

- Observation: The projection of MWVC onto an independent set looks similar to a weighted constraint.
- Question 1: Can we build the lifted graphical representation for any given WCSP? This has been answered by (Kumar 2008).
- Question 2: What is the benefit of doing so?

Lifted Representations: Example

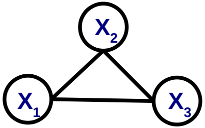

X_1	
0	0.7
1	0.2

X_2	
0	0.3
1	0.8

X_3	
0	0.1
1	1.0

X_1 \ X_2	0	1
0	0.5	0.6
1	0.7	0.3

X_2 \ X_3	0	1
0	0.6	1.3
1	1.0	1.1

X_1 \ X_3	0	1
0	0.4	0.9
1	0.7	0.8

$$E(X_1, X_2, X_3) = E_1(X_1) + E_2(X_2) + E_3(X_3) + E_{12}(X_1, X_2) + E_{13}(X_1, X_3) + E_{23}(X_2, X_3)$$

Lifted Representations: Example

X_1	
0	0.7
1	0.2

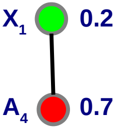

X_1 \ X_2	0	1
0	0.5	0.6
1	0.7	0.3

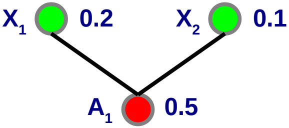

X_2	
0	0.3
1	0.8

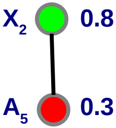

X_2 \ X_3	0	1
0	0.6	1.3
1	1.0	1.1

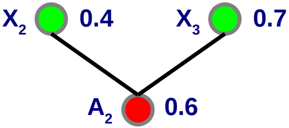

X_3	
0	0.1
1	1.0

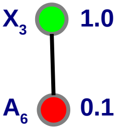

X_1 \ X_3	0	1
0	0.4	0.9
1	0.7	0.8

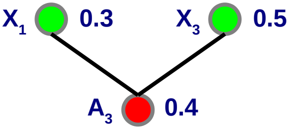

Constraint Composite Graph (CCG)

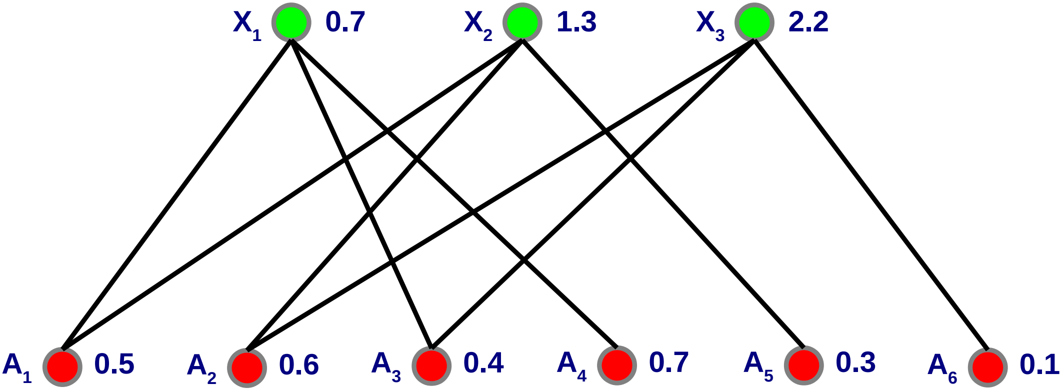

MWVC on the Constraint Composite Graph (CCG)

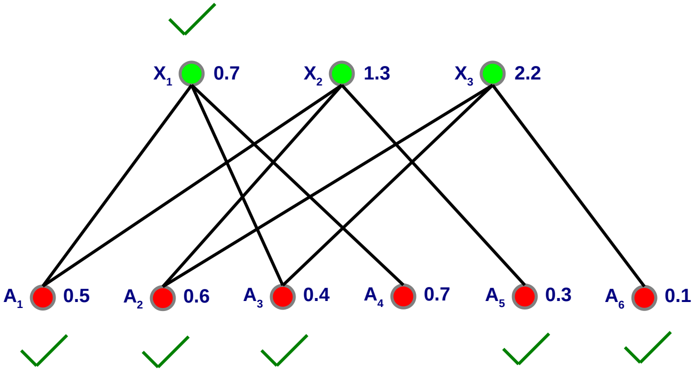

An MWVC of the CCG encodes an optimal solution of the original WCSP (Kumar 2008)!

$X_i \in \text{MWVC} \implies X_i = 1$; $X_i \notin \text{MWVC} \implies X_i = 0$.

Agenda

- The Weighted Constraint Satisfaction Problem (WCSP)
- The Constraint Composite Graph (CCG)
- ILP Encodings of WCSPs
- Conclusion

Direct ILP Encoding

Consider the WCSP $\langle \mathcal{X}, \mathcal{D}, \mathcal{C} \rangle$.

$$
\begin{aligned}
\underset{q_a^C : q_a^C \in \boldsymbol{q}}{\text{minimize}} \quad & \sum_{C \in \mathcal{C}} \sum_{a \in A(S(C))} w_a^C q_a^C \\
\text{s.t.} \quad q_a^C \quad & \in \quad \{0, 1\} \quad \forall q_a^C \in \boldsymbol{q} \\
\sum_{a \in A(S(C))} q_a^C \quad & = \quad 1 \quad \forall C \in \mathcal{C} \\
\sum_{a \in A(S(C)) : a|S(C) \cap S(C') = s} q_a^C \quad & = \sum_{a' \in A(S(C')) : a'|S(C) \cap S(C') = s} q_{a'}^{C'} \quad \forall C, C' \in \mathcal{C} \text{ and} \\
& \qquad s \in A(S(C) \cap S(C')),
\end{aligned}
$$

where $q_a^C = 1$ iff the assignment a to the variables in C is part of the to-be-determined optimal solution (Koller et al. 2009, Section 13.5).

CCG-Based ILP Encoding

Denoting its CCG by $G = \langle V, E, w \rangle$.

$$
\begin{aligned}
\underset{x_i : v_i \in V}{\text{minimize}} \quad & \sum_{i=1}^{|V|} w_i x_i \\
\text{s.t.} \quad & x_i && \in \{0,1\} && \forall\ v_i \in V \\
& x_i + x_j && \geq 1 && \forall\ (v_i, v_j) \in E,
\end{aligned}
$$

where x_i represents the presence of v_i in the MWVC.

Comparison

Encoding	Direct	CCG-Based
Number of ILP Variables	$\mathcal{O}\left(\lvert\mathcal{C}\rvert 2^{\hat{C}}\right)$	$\mathcal{O}\left(\lvert\mathcal{C}\rvert 2^{\hat{C}}\hat{C}\right)$
Number of ILP Constraints	$\mathcal{O}\left(\lvert\mathcal{C}\rvert^2 2^{\hat{C}}\right)$	$\mathcal{O}\left(\lvert\mathcal{C}\rvert 2^{\hat{C}}\hat{C}\right)$
Number of ILP Variables per ILP Constraint	$\mathcal{O}\left(2^{\hat{C}}\right)$	≤ 2

- $\lvert\mathcal{C}\rvert$: Number of WCSP constraints
- $\hat{C}$: Maximum number of WCSP variables in a WCSP constraint

The CCG-based ILP encoding is more advantageous if $\hat{C}$ is bounded!

Experimental Evaluation: Instances and Setup

- The UAI 2014 Inference Competition: PR and MMAP benchmark instances (with ten thousands variables and constraints in some cases)
  - Converted to WCSP instances by taking negative logarithms and normalizing.
- Only instances in which variables have only binary domains are used.
- Experiments were performed on a GNU/Linux workstation with an Intel Xeon processor E3-1240 v3 (8MB Cache, 3.4GHz) and 16GB RAM.
- Each benchmark instance is encoded into ILPs using both encoding methods.
- Each benchmark instance has a running time limit of 2 minutes.
- All ILPs were solved using the Gurobi Optimizer (Gurobi Optimization, Inc. 2017).

Experimental Evaluation: Running Time Comparison

Termination Status	Total	CCG-Based Only	Direct Only	Neither	Both
Number of Benchmark Instances	160	23	5	14	118

The number of benchmark instances on which the direct and CCG-based algorithms terminated within a running time limit of 120 seconds.

Experimental Evaluation: Running Time Comparison

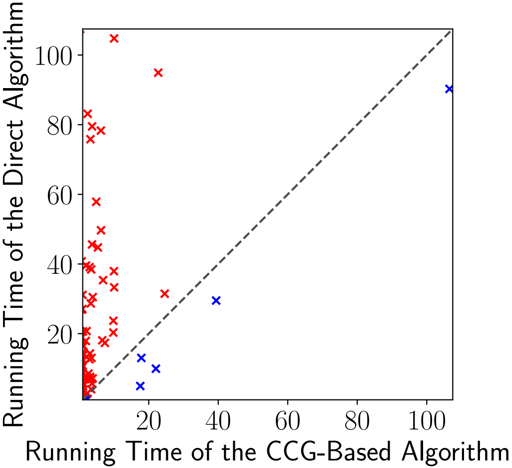

Projection of an MWVC onto an Independent Set

Assuming Boolean variables in WCSPs

- Observation: The projection of MWVC onto an independent set looks similar to a weighted constraint.
- Question 1: Can we build the lifted graphical representation for any given WCSP? This has been answered by (Kumar 2008).
- Question 2: What is the benefit of doing so? A more efficient ILP encoding

Agenda

- The Weighted Constraint Satisfaction Problem (WCSP)
- The Constraint Composite Graph (CCG)
- ILP Encodings of WCSPs
- Conclusion

Conclusion

- We developed a new ILP encoding of (Boolean) WCSPs based on the CCG.
- On Boolean WCSPs,
  - In theory, the CCG-based ILP encoding scales better in the numbers of ILP variables and constraints than the direct ILP encoding.
  - In practice, the time to solve the ILPs produced by the CCG-based ILP encoding is in general much shorter than those produced by the direct ILP encoding.

References I

- Gurobi Optimization, Inc. *Gurobi Optimizer Reference Manual*. 2017. URL: http://www.gurobi.com.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009. ISBN: 978-0262258357.
- Vladimir Kolmogorov. *Primal-dual Algorithm for Convex Markov Random Fields*. Tech. rep. MSR-TR-2005-117. Microsoft Research, 2005.
- T. K. Satish Kumar. "A Framework for Hybrid Tractability Results in Boolean Weighted Constraint Satisfaction Problems". In: *the Proceedings of the International Conference on Principles and Practice of Constraint Programming*. 2008, pp. 282–297.
- Isabel Milho, Ana Fred, Jorge Albano, Nuno Baptista, and Paulo Sena. "An Auxiliary System for Medical Diagnosis Based on Bayesian Belief Networks". In: *Portuguese Conference on Pattern Recognition*. 2000.
- Nicola Muscettola, P. Pandurang Nayak, Barney Pell, and Brian C. Williams. "Remote Agent: To Boldly Go Where No AI System Has Gone Before". In: *Artificial Intelligence* 103.1–2 (1998), pp. 5–47.

References II

Jonathan S Yedidia, William T Freeman, and Yair Weiss. "Understanding belief propagation and its generalizations". In: *Exploring Artificial Intelligence in the New Millennium* 8 (2003), pp. 236–239.

Matthias Zytnicki, Christine Gaspin, and Thomas Schiex. "DARN! A Weighted Constraint Solver for RNA Motif Localization". In: *Constraints* 13.1 (2008), pp. 91–109.