

This is a pre-copy-editing, author-produced PDF of an article accepted for publication in “*The Computer Journal*” following peer review.

The definitive publisher-authenticated version:

Exploring Directional Path-Consistency for Solving Constraint Networks

[Shufeng Kong](#) [Sanjiang Li](#) [Michael Sioutis](#)

***The Computer Journal*, Volume 61, Issue 9, September 2018, Pages 1338–**

The article is available online at:

<https://academic.oup.com/comjnl/article-abstract/61/9/1338/4779882?redirectedFrom=fulltext>

Exploring Directional Path-Consistency for Solving Constraint Networks

Shufeng Kong¹, Sanjiang Li¹, and Michael Sioutis²

¹ QSI, FEIT, University of Technology Sydney, Sydney, Australia
shufeng.kong@student.uts.edu.au, sanjiang.li@uts.edu.au

² AASS, Örebro University, Örebro, Sweden
michael.sioutis@oru.se

Abstract. Among the local consistency techniques used for solving constraint networks, path-consistency (PC) has received a great deal of attention. However, enforcing PC is computationally expensive and sometimes even unnecessary. Directional path-consistency (DPC) is a weaker notion of PC that considers a given variable ordering and can thus be enforced more efficiently than PC. This paper shows that **DPC** (the DPC enforcing algorithm of Dechter and Pearl) decides the constraint satisfaction problem (CSP) of a constraint language Γ if it is complete and has the variable elimination property (VEP). However, we also show that no complete VEP constraint language can have a domain with more than 2 values.

We then present a simple variant of the **DPC** algorithm, called **DPC***, and show that the CSP of a constraint language can be decided by **DPC*** if it is closed under a majority operation. In fact, **DPC*** is sufficient for guaranteeing backtrack-free search for such constraint networks. Examples of majority-closed constraint classes include the classes of connected row-convex (CRC) constraints and tree-preserving constraints, which have found applications in various domains, such as scene labeling, temporal reasoning, geometric reasoning, and logical filtering. Our experimental evaluations show that **DPC*** significantly outperforms the state-of-the-art algorithms for solving majority-closed constraints.

Keywords: Path-consistency; directional path-consistency; constraint networks

1 Introduction

Many Artificial Intelligence tasks can be formulated as *constraint networks* [30], such as natural language parsing [28], temporal reasoning [14,31], and spatial reasoning [26]. A constraint network comprises a set of variables ranging over some domain of possible values, and a set of constraints that specify allowed value combinations for these variables. Solving a constraint network amounts to assigning values to its variables such that its constraints are satisfied. *Backtracking search* is the principal mechanism for solving a constraint network; it assigns values to variables in a depth-first manner, and backtracks to the previous variable assignment if there are no consistent values for the variable at hand. *Local consistency* techniques are commonly used to reduce the size of the search space before commencing search. However, searching for a complete solution for a constraint network is still hard. In fact, even deciding whether the constraint network has a solution is NP-complete in general. Therefore, given a particular

form of local consistency, a crucial task is to determine problems that can be solved by *backtrack-free* search using that local consistency [17].

This paper considers a particular form of local consistency, called *path-consistency* (PC), which is one of the most important and heavily studied local consistencies in the literature (see e.g. [27,29,32,6,10]). Recently, it was shown that PC can be used to decide the satisfiability of a problem if and only if the problem does not have the *ability to count* [2,1]; however, it remains unclear whether backtrack-free search can be used to extract a solution for such a problem after enforcing PC.

Directional path-consistency (DPC) [15] is a weaker notion of PC that considers a given variable ordering and can thus be enforced more efficiently than PC. The DPC enforcing algorithm of Dechter and Pearl [15], denoted by **DPC**, has been used to efficiently solve reasoning problems in temporal reasoning [14,31] and spatial reasoning [33]. It is then natural to ask *what binary constraint networks with finite domains can be solved by DPC*. Dechter and Pearl [15] showed that **DPC** is sufficient for enabling backtrack-free search for networks with constraint graphs of *regular width 2*. We consider the aforementioned question in the context of *constraint languages*, which is a widely adopted approach in the study of tractability of constraint satisfaction problems [8]. Specifically, we are interested in finding all binary constraint languages Γ such that the consistency of any constraint network defined over Γ can be decided by **DPC**.

To this end, we first exploit the close connection between DPC and *variable elimination* by defining constraint languages that have the (weak) variable elimination property (VEP) (which will become clear in Definition 6). We call a constraint language Γ *complete* if it contains all relations that are definable in Γ (in the sense of Definition 3). Then, we show that the constraint satisfaction problem (CSP) of a complete constraint language Γ can be decided by **DPC** if it is complete and has VEP, which is shown to be equivalent to the Helly property. However, we also show that no complete VEP constraint language can have a domain with more than 2 values.

We then present a simple variant of the algorithm **DPC**, called **DPC***, and show that the consistency of a constraint network can be decided by **DPC*** if it is defined over any majority-closed constraint language. In fact, we show that **DPC*** is sufficient for guaranteeing backtrack-free search for such constraint networks. Several important constraint classes have been found to be majority-closed. The most well-known one is the class of *connected row-convex (CRC)* constraints [16], which is further generalized to a larger class of *tree-preserving* constraints [20]. The class of CRC constraints has been successfully applied to temporal reasoning [24], logical filtering [25], and geometric reasoning [23], and the class of tree-preserving constraints can model a large subclass of the scene labeling problem [20]. We also conduct experimental evaluations to compare **DPC*** to the state-of-the-art algorithms for solving majority-closed constraints, and show that **DPC*** significantly outperforms the latter algorithms.

The remainder of this paper is organized as follows. In Section 2 we introduce basic notions and results that will be used throughout the paper. In Section 3 we present the **DPC** algorithm, and in Section 4 we discuss the connection between **DPC** and variable elimination. In Section 5 we prove that a complete constraint language Γ has weak VEP if and only if Γ is majority-closed. We then present in Section 6 our variable

elimination algorithm DPC^* , and empirically evaluate DPC^* in Section 7. Finally, Section 8 concludes the paper.

2 Preliminaries

This section recalls necessary notions and results.

Definition 1. A binary constraint network (BCN) \mathcal{N} is a triple $\langle V, \mathcal{D}, C \rangle$, where

- $V = \{v_1, \dots, v_n\}$ is a nonempty finite set of variables;
- $\mathcal{D} = \{D_1, \dots, D_n\}$, where D_i is the domain of v_i ;
- $C = \{(s_p, R_p) \mid 1 \leq p \leq m\}$ is a set of binary constraints, where $s_p = (v_i, v_j) (i \neq j)$ (called the scope of (s_p, R_p)) is a pair of variables in V and R_p (called the constraint relation of (s_p, R_p)) is a subset of $D_i \times D_j$.

Given a BCN $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ and any pair of variables (v_i, v_j) with $v_i, v_j \in V$ and $v_i \neq v_j$, we assume that there exists at most one constraint between the pair. For simplicity, we will often denote the constraint between v_i and v_j by R_{ij} . Further, we assume $R_{ij} = R_{ji}^{-1}$, which is the inverse of R_{ji} . We write $R_{ij} \in C$ throughout the paper to state that a constraint with scope (v_i, v_j) is in C . Usual operations on relations such as intersection (\cap), composition (\circ), and inverse ($^{-1}$) are also assumed.

A *partial solution* of \mathcal{N} w.r.t. a subset V' of V is an assignment of values to variables in V' such that all of the constraints R_{ij} with $v_i, v_j \in V'$ are satisfied. A partial solution w.r.t. V is called a *solution* of \mathcal{N} . We say that \mathcal{N} is *consistent* or *satisfiable* if it admits a solution, and *inconsistent* or *unsatisfiable* otherwise. A BCN $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is said to be *globally consistent* if every partial solution w.r.t. $V' \subseteq V$ can be consistently extended to a solution w.r.t. V . Further, \mathcal{N} is said to be *trivially inconsistent* if C contains an empty constraint or \mathcal{D} contains an empty domain. Two BCNs are *equivalent* if they have the same set of solutions.

Example 1. Consider a BCN $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$, where

- $V = \{v_1, v_2, v_3, v_4\}$;
- $D_i = \{a, b, c\}$ for $i = 1, 2, 3, 4$;
- $C = \{R_{12}, R_{23}, R_{34}, R_{24}\}$ where

$$R_{12} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}, R_{34} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \text{ and } R_{32} = R_{24} = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

We use Boolean matrices to represent binary relations. For example, R_{12} represents the relation $\{\langle a, a \rangle, \langle a, b \rangle, \langle a, c \rangle, \langle b, a \rangle, \langle b, b \rangle, \langle c, a \rangle\}$ between v_1 and v_2 , where the values in both D_1 and D_2 are ordered as $a < b < c$. It is easy to check that $\sigma = \langle a, a, a, a \rangle$ is a solution of \mathcal{N} .

The *constraint graph* $G_{\mathcal{N}}$ of $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is the *undirected* graph (V, E) , where $e_{ij} \in E$ iff $R_{ij} \in C$. We assume e_{ij} is always labeled with its corresponding constraint R_{ij} . Fig. 1 shows the constraint graph of the BCN in Example 1.




Fig. 1: The constraint graph $G_{\mathcal{N}}$ of the BCN in Example 1. Note that $\langle v_1, v_2, v_3, v_4 \rangle$ is a PEO in $G_{\mathcal{N}}$, whereas $\langle v_2, v_1, v_3, v_4 \rangle$ is not.




Fig. 2: Path-Consistency [6].

An undirected graph G is *triangulated* or *chordal* if every cycle of length greater than 3 has a *chord*, *i.e.*, an edge connecting two non-consecutive vertices of the cycle. The constraint graph $G_{\mathcal{N}}$ of a network $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ can be completed or triangulated by adding new edges e_{ij} labeled with the universal constraint $D_i \times D_j$.

Triangulated graphs play a key role in efficiently solving large sparse constraint networks [6,31,33]. A graph G is triangulated iff it admits a *perfect elimination ordering* (PEO) [18]. An ordering \prec of the vertices of a graph $G = (V, E)$ is a PEO if $F_v = \{w \mid (v, w) \in E, v \prec w\}$ for all of $v \in V$, *i.e.*, the set of successors of v in the ordering induces a complete subgraph of G (see Fig. 1 for an example).

A variable v_i is *arc-consistent* (AC) relative to a variable v_j (or R_{ij}) if for any $a \in D_i$ we have some $b \in D_j$ such that $\langle a, b \rangle \in R_{ij}$. Given a constraint graph $G_{\mathcal{N}} = (V, E)$, an edge e_{ij} of $G_{\mathcal{N}}$ is AC if v_i is AC relative to v_j . Let $\pi = (v_0, \dots, v_i, \dots, v_k)$ be a path in $G_{\mathcal{N}}$ with $e_{0k} \in E$. We say that π is *path-consistent* (PC) if for all $\langle c_0, c_k \rangle \in R_{0k}$ we can find values for all intermediate variables v_i ($0 < i < k$) such that all the constraints $R_{i,i+1}$ ($0 \leq i < k$) are satisfied. See Fig. 2 for an illustration. In particular, (v_0, v_2) is PC relative to a third vertex v_1 if e_{01}, e_{12} , and e_{02} are all in E and the path $\pi = (v_0, v_1, v_2)$ is PC. A constraint graph $G_{\mathcal{N}}$ is AC (resp. PC) iff all edges (resp. paths) in $G_{\mathcal{N}}$ are AC (resp. PC); $G_{\mathcal{N}}$ is *strongly PC* iff it is both AC and PC. A constraint network \mathcal{N} is AC if $G_{\mathcal{N}}$ is AC and \mathcal{N} is PC if the competition of $G_{\mathcal{N}}$ is PC [6].

Consider the BCN \mathcal{N} in Example 1. We can see that every edges in $G_{\mathcal{N}}$ is AC, but the path $\pi = (v_3, v_2, v_4)$ is not PC as R_{34} is not contained in $R_{32} \circ R_{24}$.

In this paper we are concerned with BCNs defined over a particular *constraint language* and we use *constraint languages*, *constraint classes* and *sets of relations* interchangeably.

Definition 2. [19] Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be a set of domains. An n -ary relation R over \mathcal{D} is a subset of $D_1 \times \dots \times D_n$. For any tuple $t \in R$ and any $1 \leq i \leq n$, we denote by $t[i]$ the value in the i -th coordinate position of t and write t as $\langle t[1], \dots, t[n] \rangle$.

Definition 3. [19] Given a set of binary relations Γ , we write Γ^+ for the set of relations that can be obtained from Γ using some sequence of the following operations:

- Cartesian product, i.e., for $R_1, R_2 \in \Gamma$, $R_1 \times R_2 = \{\langle t_1, t_2 \rangle \mid t_1 \in R_1, t_2 \in R_2\}$,
- equality selection, i.e., for $R \in \Gamma$, $\tau_{i=j}(R) = \{t \in R \mid t[i] = t[j]\}$, and
- projection, i.e., for $R \in \Gamma$, $\pi_{i_1, \dots, i_k}(R) = \{\langle t[i_1], \dots, t[i_k] \rangle \mid t \in R\}$.

A relation R is said to be definable in Γ if $R \in \Gamma^+$, and a set of binary relations Γ is said to be complete if every binary relation definable in Γ is also contained in Γ . The completion of Γ , written as Γ^c , is the set of all binary relations contained in Γ^+ .

The following lemma asserts that a complete set of binary relations Γ is closed under the operations that are used to achieve PC and AC.

Lemma 1. [11] Let Γ be a complete set of binary relations over a domain D . Suppose R, S are binary relations and T a unary relation, all in Γ . Then $R \cap S$, $R \circ S$, and $T' = \{a \in D \mid \langle a, b \rangle \in R, b \in T\}$ are also all in Γ .

Let Γ be a set of binary relations. A BCN $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is defined over (or, simply, over) Γ if $R \in \Gamma$ for every constraint (s, R) in C . The constraint satisfaction problem (CSP) of Γ , denoted by $\mathbf{CSP}(\Gamma)$, is the problem of deciding the consistency of BCNs defined over Γ . $\mathbf{CSP}(\Gamma^+)$ is log-space reducible to $\mathbf{CSP}(\Gamma)$ [11].

A set of binary relations Γ is weakly closed under singletons, if $\{\langle a, b \rangle\} \in \Gamma^+$ for any $R \in \Gamma$ and any $\langle a, b \rangle \in R$.

In this paper we often assume that the constraint languages are *complete* and *weakly closed under singletons*. We will see that this is not very restrictive as, for any set Γ of binary relations that is closed under a majority operation φ , the completion Γ^c of Γ is also closed under φ [19] and weakly closed under singletons (cf. Proposition 4).

3 The Strong Directional PC Algorithm

This section recalls the notions of *directional arc-consistency* (DAC) and *directional path-consistency* (DPC), and the *strong DPC* enforcing algorithm of Dechter and Pearl [13].

Definition 4. [13] Let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a BCN and $\prec = (v_1, \dots, v_n)$ an ordering of the variables in V . We say that \mathcal{N} is directionally arc-consistent (DAC) relative to \prec if v_i is arc-consistent relative to v_k for all $k > i$ with $R_{ik} \in C$. Similarly, \mathcal{N} is directionally path-consistent (DPC) relative to \prec if, for any $i \neq j$ with $R_{ij} \in C$, (v_i, v_j) is path-consistent relative to v_k for all $k > i, j$ whenever $R_{ik}, R_{jk} \in C$. Meanwhile, \mathcal{N} is strongly DPC relative to \prec if it is both DAC and DPC relative to \prec .

Algorithm 1: DPC

Input : A binary constraint network $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$;
an ordering $\prec = (v_1, \dots, v_n)$ on V .

Output: An equivalent subnetwork that is strongly DPC relative to \prec , or “Inconsistent”.

```
1 for  $k \leftarrow n$  to 1 do
2   for  $i < k$  with  $R_{ik} \in C$  do
3      $D_i \leftarrow D_i \cap R_{ki}(D_k)$ ;
4     if  $D_i = \emptyset$  then
5       return “Inconsistent”
6   for  $i, j < k$  with  $R_{ik}, R_{jk} \in C$  do
7     if  $R_{ij} \notin C$  then
8        $R_{ij} \leftarrow D_i \times D_j$ ;
9        $C \leftarrow C \cup \{R_{ij}\}$ ;
10       $R_{ij} \leftarrow R_{ij} \cap (R_{ik} \circ R_{kj})$ ;
11      if  $R_{ij} = \emptyset$  then
12        return “Inconsistent”;
13 return  $\mathcal{N}$ .
```

The strong DPC algorithm is presented as Algorithm 1. In comparison with traditional PC algorithms [10], a novelty of this single pass algorithm is its explicit reference to the constraint graph of the input constraint network. As only Line 8 may require extra working space, Algorithm 1 has a very low space complexity in practice. Further, Algorithm 1 runs in $O(w^*(\prec) \cdot e \cdot (\alpha + \beta))$ time [13], where e is the number of edges of the output constraint graph, $w^*(\prec)$ is the induced width of the ordered graph along \prec , and α, β are the runtimes of relational intersection and composition respectively. Note that $w^*(\prec) \leq n$ and α, β are bounded by $O(d)$, where d is the largest domain size.

Proposition 1. [13] *Let (\mathcal{N}, \prec) be an input to Algorithm 1, where $\prec = (v_1, \dots, v_n)$. Suppose $\mathcal{N}' = \langle V, \mathcal{D}', C' \rangle$ is the output. Then*

- (i) $G_{\mathcal{N}'}$ is triangulated and \prec^{-1} , the inverse of \prec , is a PEO of $G_{\mathcal{N}'}$;
- (ii) \mathcal{N}' is equivalent to \mathcal{N} and strongly DPC relative to \prec .

Let Γ be a set of binary relations. We say that Algorithm 1 *decides* $\mathbf{CSP}(\Gamma)$ if, for any given BCN \mathcal{N} in $\mathbf{CSP}(\Gamma)$ and any ordering \prec of variables of \mathcal{N} , Algorithm 1 returns “Inconsistent” iff \mathcal{N} is inconsistent.

The following corollary follows directly from Proposition 1.

Corollary 1. *Let Γ be a complete set of binary relations. Then the following two conditions are equivalent:*

- (i) *Algorithm 1 decides $\mathbf{CSP}(\Gamma)$.*
- (ii) *Let \mathcal{N} be a non-trivially inconsistent BCN in $\mathbf{CSP}(\Gamma)$. Suppose \mathcal{N} 's constraint graph $G_{\mathcal{N}}$ is triangulated and let $\prec^{-1} = (v_n, \dots, v_1)$ be a PEO of it. Then \mathcal{N} is consistent if \mathcal{N} is strongly DPC relative to \prec .*

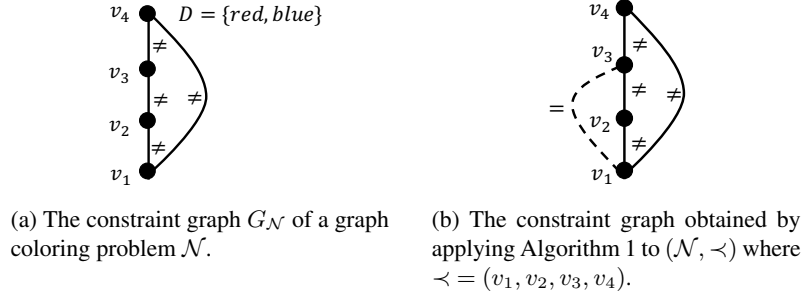


Fig. 3: A graph coloring problem with domain $D_i = \{red, blue\}$ for $i = 1, 2, 3, 4$ [13].

Example 2. The graph coloring problem \mathcal{N} with domains $\{red, blue\}$ depicted in Fig. 3 is taken from [13] and can be decided by Algorithm 1. After applying Algorithm 1 to (\mathcal{N}, \prec) , where $\prec = (v_1, v_2, v_3, v_4)$, a solution can be obtained along \prec in a backtrack-free manner (see Fig. 3b).

4 Directional PC and Variable Elimination

This section characterizes the binary constraint languages Γ such that $\text{CSP}(\Gamma)$ can be decided by **DPC**. We observe that **DPC** achieves (strong) **DPC** using the idea of *variable elimination* [13]: it iterates variables along the ordering \prec^{-1} , and propagates the constraints of a variable v_k to subsequent variables in the ordering with the update rule $R_{ij} \leftarrow R_{ij} \cap (R_{ik} \circ R_{kj})$, as if v_k is ‘eliminated’.

The following definition formalizes the process of elimination.

Definition 5. Let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a BCN with $V = \{v_1, \dots, v_n\}$ and $\mathcal{D} = \{D_1, \dots, D_n\}$. For a variable v_x in V , let $E_x = \{R_{ix} \mid R_{ix} \in C\}$. The network obtained after v_x is eliminated from \mathcal{N} , written as

$$\mathcal{N}_{-x} = \langle V \setminus \{v_x\}, \{D'_1, \dots, D'_{x-1}, D'_{x+1}, \dots, D'_n\}, C' \rangle,$$

is defined as follows:

- If $E_x = \{R_{ix}\}$, we set $C' = C \setminus E_x$ and let

$$D'_j = \begin{cases} D_i \cap R_{xi}(D_x), & \text{if } j = i \\ D_j, & \text{otherwise} \end{cases} \quad (1)$$

- If $|E_x| \neq 1$, we set $D'_j = D_j$ for all $j \neq x$, and let

$$C' = (C \setminus E_x) \cup \{R_{ix} \circ R_{xj} \cap R_{ij} \mid R_{jx}, R_{ix} \in E_x, i \neq j\}.$$

R_{ij} is assumed to be $D_i \times D_j$ if $R_{ij} \notin C$.

Fig. 4 illustrates the elimination process.




Fig. 4: Two binary constraint networks \mathcal{N} and \mathcal{N}_{-1} .

Definition 6. A BCN $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is said to have the variable elimination property (VEP), if, for any v_x in V , every solution of \mathcal{N}_{-x} can be extended to a solution of \mathcal{N} .

\mathcal{N} is said to have weak VEP, if, for any v_x in V such that v_x is AC relative to all relations in E_x , every solution of \mathcal{N}_{-x} can be extended to a solution of \mathcal{N} .

A set of binary relations Γ is said to have (weak) VEP if every BCN in $\mathbf{CSP}(\Gamma)$ has (weak) VEP. Such a set of binary relations Γ is also called a (weak) VEP class.

It is easy to see that, if a BCN (a set of binary relations) has VEP, then it also has weak VEP. The following example explains why we should take special care when eliminating variables with only one successor in Eq. (1).




Fig. 5: A constraint graph that is a chain.

Example 3. Let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a BCN defined by $V = \{v_1, v_2, v_3\}$, $D_1 = D_2 = D_3 = \{0, 1\}$, and $C = \{((v_3, v_2), R), ((v_2, v_1), R)\}$ with $R = \{(1, 0)\}$ (see Fig. 5). Suppose we do not have the operation specified in (1) and $\prec = (v_3, v_2, v_1)$ is the variable elimination ordering. Let \mathcal{N}_{-3} be the restriction of \mathcal{N} to $\{v_1, v_2\}$. Then \mathcal{N}_{-3} has a unique solution σ but it cannot be extended to a solution of \mathcal{N} .

Proposition 2. Let Γ be a complete set of binary relations that is weakly closed under singletons. Then DPC decides $\mathbf{CSP}(\Gamma)$ iff Γ has VEP.

Proof. We address the ‘if’ part first. Assume that Γ has VEP, and let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a network in $\mathbf{CSP}(\Gamma)$ that is non-trivially inconsistent and strongly DPC relative to

$\prec = (v_1, \dots, v_n)$, where $G_{\mathcal{N}}$ is triangulated and \prec^{-1} is a PEO of it. We show that \mathcal{N} is consistent. Let $V_i = \{v_1, \dots, v_i\}$ and $\mathcal{N}|_{V_i}$ be the restriction of \mathcal{N} to V_i . We claim that $\mathcal{N}|_{V_i}$ is consistent for $k = 1, \dots, n$ and prove the claim by induction on k . First, since \mathcal{N} is strongly DPC relative to \prec , D_1 is not empty and there is an $a_1 \in D_1$. Then, $\mathcal{N}|_{V_1}$ is consistent and has a solution $\sigma_1 = \langle a_1 \rangle$. Further, suppose that $\mathcal{N}|_{V_i}$ is consistent and $\sigma_i = \langle a_1, a_2, \dots, a_i \rangle$ is a solution of $\mathcal{N}|_{V_i}$. We show that σ_i can be extended to a solution $\sigma_{i+1} = \langle a_1, \dots, a_i, a_{i+1} \rangle$ of $\mathcal{N}|_{V_{i+1}}$. We consider three subcases: (i) If E_{i+1} is empty, then for any $a_{i+1} \in D_{i+1}$, $\sigma_{i+1} = \langle a_1, a_2, \dots, a_i, a_{i+1} \rangle$ is a solution of $\mathcal{N}|_{V_{i+1}}$ because there is no constraint between v_i and v_{i+1} . (ii) If $E_{i+1} = \{R_{j,i+1}\}$ is a singleton, then, since \mathcal{N} is DAC relative to \prec , we have $a_{i+1} \in D_{i+1}$ such that $\langle a_i, a_{i+1} \rangle \in R_{i,i+1}$ and $\sigma_{i+1} = \langle a_1, a_2, \dots, a_i, a_{i+1} \rangle$ is a solution of $\mathcal{N}|_{V_{i+1}}$. (iii) If E_{i+1} contains more than one variable, for every pair of distinct variables (v_x, v_y) in V_i with $R_{x,i+1}, R_{y,i+1} \in E_{i+1}$, we know $R_{xy} \in C$ because \prec^{-1} is a PEO of $G_{\mathcal{N}}$. Moreover, since (v_x, v_y) is PC relative to v_{i+1} , we have $R_{xy} \subseteq R_{x,i+1} \circ R_{i+1,y}$. Then since Γ has VEP and $\mathcal{N}|_{V_i}$ is indeed the same network as the one obtained by eliminating v_{i+1} from $\mathcal{N}|_{V_{i+1}}$, by Definition 6, σ_i can be extended to a solution σ_{i+1} of $\mathcal{N}|_{V_{i+1}}$. Thus, \mathcal{N} is consistent. By Corollary 1, **DPC** decides **CSP**(Γ).

Next, we address the ‘only if’ part. Assume that **DPC** decides **CSP**(Γ). We show that Γ has VEP. Let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a non-trivially inconsistent network in **CSP**(Γ). Given $v_x \in V$, we show that every solution of \mathcal{N}_{-x} can be extended to \mathcal{N} . Without loss of generality, we assume that $x = n$. Let $\sigma = \langle a_1, \dots, a_{n-1} \rangle$ be a solution of \mathcal{N}_{-n} , and $E_n = \{R_{in} \mid R_{in} \in C\}$. By the definition of \mathcal{N}_{-n} , for any $R_{in}, R_{jn} \in E_n (i \neq j)$, we have $\langle a_i, a_j \rangle \in R_{in} \circ R_{nj} \cap R_{ij}$. We then construct a new problem $\mathcal{N}' = \langle V, \mathcal{D}', C' \rangle$ in **CSP**(Γ), where $\mathcal{D}' = \{D'_1, \dots, D'_{n-1}, D_n\}$ with $D'_i = \{a_i\}$ for $1 \leq i < n$ and $C' = \{\{\langle a_i, a_j \rangle\} \mid 1 \leq i \neq j < n\} \cup E_n$. Clearly, σ is also a solution of \mathcal{N}'_{-n} and \mathcal{N}'_{-n} is strongly PC and, hence, strongly DPC relative to the ordering (v_1, \dots, v_{n-1}) . Further, since $\langle a_i, a_j \rangle \in R_{in} \circ R_{nj} \cap R_{ij}$ for any $R_{in}, R_{jn} \in E_n (i \neq j)$, we have that \mathcal{N}' is strong DPC relative to $\prec = (v_1, \dots, v_n)$. As $G_{\mathcal{N}'_{-n}}$ is complete, $G_{\mathcal{N}'}$ is triangulated with \prec^{-1} being a PEO of it. As **DPC** decides **CSP**(Γ) and $\mathcal{N}' \in \mathbf{CSP}(\Gamma)$, by Corollary 1, \mathcal{N}' is consistent and has a solution that extends σ and is also a solution of \mathcal{N} . This shows that \mathcal{N} is consistent and, hence, Γ has VEP. \square

Therefore, if $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is defined over a complete VEP class, then **DPC** can decide it. Note that in the above proposition we require Γ to be complete. This is important; for example, every *row-convex constraint* [3] network has VEP (cf. the proof of [34, Theorem 1]) and, hence, the class of row-convex constraints has VEP. However, **DPC** does not decide the consistency problem over the row-convex constraint class because it was shown to be NP-hard (cf. e.g. [20]).

VEP is closely related to the Helly property, defined as follows.

Definition 7. A set Γ of binary relations over $\mathcal{D} = \{D_1, \dots, D_n\}$ is said to have the Helly property if for any $k > 2$ binary relations, $R_i \subseteq D_{u_i} \times D_{u_0} (1 \leq i \leq k, 1 \leq u_i \neq u_0 \leq n)$, in Γ , and for any k values, $a_i \in D_{u_i} (1 \leq i \leq k)$, such that $R_i(a_i) = \{b \in D_{u_0} \mid \langle a_i, b \rangle \in R_i\}$ is nonempty, we have $\bigcap_{i=1}^k R_i(a_i) \neq \emptyset$ iff $R_i(a_i) \cap R_j(a_j) \neq \emptyset$ for any $1 \leq i \neq j \leq k$.




Fig. 6: An illustration of Example 4.

Example 4. Let $D_{u_0} = \{a, b, c, d\}$, $D_{u_1} = \{e\}$, $D_{u_2} = \{f\}$, $D_{u_3} = \{g\}$ and $R_1 = \{\langle e, a \rangle, \langle e, b \rangle, \langle e, c \rangle\}$, $R_2 = \{\langle f, b \rangle, \langle f, c \rangle, \langle f, d \rangle\}$, $R_3 = \{\langle g, c \rangle, \langle g, d \rangle, \langle g, a \rangle\}$. See Fig. 6 for an illustration. Then $\Gamma = \{R_1, R_2, R_3\}$ over $\mathcal{D} = \{D_{u_0}, D_{u_1}, D_{u_2}, D_{u_3}\}$ has the Helly property.

Theorem 1. *A set of binary relations Γ has VEP iff it has the Helly property.*

Proof. Suppose Γ has VEP, we show Γ has the Helly property. Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be the set of domains related to relations in Γ . Suppose $R_i \subseteq D_{u_i} \times D_{u_0}$ ($1 \leq i \leq k$, $1 \leq u_i \neq u_0 \leq n$) are $k > 2$ binary relations in Γ and $a_i \in D_{u_i}$ ($1 \leq i \leq k$) are values such that $\emptyset \neq R_i(a_i) \subseteq D_{u_0}$. Suppose $R_i(a_i) \cap R_j(a_j)$ is nonempty for any i, j with $1 \leq i \neq j \leq k$. We show that $\bigcap_{i=1}^k R_i(a_i)$ is nonempty. To this end, we construct a BCN $\mathcal{N} = \langle V, \mathcal{D}', C \rangle$ over Γ with $V = \{v_1, \dots, v_k, v_{k+1}\}$, $\mathcal{D}' = \{D_{u_1}, \dots, D_{u_k}, D_{u_0}\}$, and $C = \{R_{i,k+1} \mid 1 \leq i \leq k\}$, where $R_{i,k+1} = R_i$. Consider $\mathcal{N}_{-(k+1)}$. As $R_i(a_i) \cap R_j(a_j) \neq \emptyset$, we have $\langle a_i, a_j \rangle \in R_{i,k+1} \circ R_{k+1,j}$. This shows that $\sigma = \langle a_1, \dots, a_k \rangle$ is a solution of $\mathcal{N}_{-(k+1)}$. Since Γ and, hence, \mathcal{N} have VEP, \mathcal{N} has a solution that extends σ . Hence there exists $a \in D_{u_0}$ such that $a \in R_{i,k+1}(a_i)$ for every $1 \leq i \leq k$. Thus $\bigcap_{i=1}^k R_i(a_i) \neq \emptyset$. This proves that Γ has the Helly property.

Suppose Γ has the Helly property, we show Γ has VEP. Let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a non-trivially inconsistent BCN defined over Γ with $V = \{v_1, v_2, \dots, v_n\}$ and C is a set of binary constraints $((v_i, v_j), R)$ with $R \in \Gamma$. Let $E_n = \{R_{in} \mid R_{in} \in C\}$. Assume $\sigma = \langle a_1, a_2, \dots, a_{n-1} \rangle$ is a solution of, say, \mathcal{N}_{-n} . We show that there exists $a_n \in D_n$ such that $\langle a_1, \dots, a_{n-1}, a_n \rangle$ is a solution of \mathcal{N} . If E_n is empty, we can take any a_n from D_n which is nonempty since \mathcal{N} is non-trivially inconsistent; if E_n contains only one constraint, say, $((v_i, v_n), R_{in})$, by $a_i \in D'_i = D_i \cap R_{ni}(D_n)$, there exists $a_n \in D_n$ such that $\langle a_i, a_n \rangle \in R_{in}$; if E_n contains $k \geq 2$ constraints and let them be $((v_{u_i}, v_n), R_{u_i n})$ ($1 \leq i \leq k$), we have $\langle a_i, a_j \rangle \in R_{u_i u_j} \cap (R_{u_i n} \circ R_{n u_j})$ for $1 \leq i \neq j \leq k$. Therefore, we have $R_{u_i n}(a_i) \cap R_{u_j n}(a_j) \neq \emptyset$ for $1 \leq i \neq j \leq k$. By the Helly property of Γ , we have $\bigcap_{i=1}^k R_{u_i n}(a_i) \neq \emptyset$. So we can take any $a_n \in \bigcap_{i=1}^k R_{u_i n}(a_i)$ so that $\langle a_1, \dots, a_{n-1}, a_n \rangle$ is a solution of \mathcal{N} . Therefore, Γ has VEP. \square

The class of row-convex constraints [3] and the class of *tree-convex* constraints [35] have the Helly property and, thus, they have VEP by Theorem 1.

Proposition 3. *Suppose Γ is a set of binary relations that has the Helly property. Let $\mathcal{N} \in \text{CSP}(\Gamma)$. Suppose \mathcal{N} is non-trivially inconsistent and $G_{\mathcal{N}}$ is triangulated with*

$\prec^{-1} = (v_n, \dots, v_1)$ as a PEO of it. Then \mathcal{N} is consistent if it is strongly DPC relative to \prec .

Proof. Let $\mathcal{N} \in \mathbf{CSP}(\Gamma)$. Suppose $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is non-trivially inconsistent and $G_{\mathcal{N}}$ is triangulated with $\prec^{-1} = (v_n, \dots, v_1)$ being a PEO of it. Suppose \mathcal{N} is strongly DPC relative to \prec . We show that \mathcal{N} is consistent. Let $V_k = \{v_1, \dots, v_k\}$ and \mathcal{N}_k be the restriction of \mathcal{N} to V_k . Since \mathcal{N} is non-trivially inconsistent, we have that \mathcal{N}_1 is consistent. Suppose \mathcal{N}_k is consistent, we show \mathcal{N}_{k+1} is consistent. Let $\sigma = \langle a_1, \dots, a_k \rangle$ be a solution of \mathcal{N}_k . Let $E_{k+1} = \{R_{i,k+1} \mid R_{i,k+1} \in C, i \leq k\}$. Since $G_{\mathcal{N}}$ is triangulated and $\prec^{-1} = (v_n, \dots, v_1)$ is a PEO of it, for any two different constraints $R_{i,k+1}, R_{j,k+1} \in E_{k+1}$, we have $R_{ij} \in C$. Further, since \mathcal{N} is strongly DPC relative to \prec , we have $\langle a_i, a_j \rangle \in R_{i,k+1} \circ R_{k+1,j} \cap R_{ij}$. Thus, we have $R_{i,k+1}(a_i) \cap R_{j,k+1}(a_j) \neq \emptyset$ for any two different constraints $R_{i,k+1}, R_{j,k+1} \in E_{k+1}$. Since Γ has the Helly property, we have $\bigcap_{R_{i,k+1} \in E_{k+1}} R_{i,k+1}(a_i) \neq \emptyset$. Therefore, σ can be extended to a solution of \mathcal{N}_{k+1} and \mathcal{N}_{k+1} is consistent. By induction on k , we have that \mathcal{N} is consistent. \square

5 Majority-Closed Constraint Languages

In this section we characterize weak VEP classes. We will show that a complete set of binary relations Γ has weak VEP iff all relations in Γ are closed under a *majority* operation, which is defined as follows.

Definition 8. [7] Let $\mathcal{D} = \{D_1, \dots, D_n\}$ be a set of domains. A multi-sorted majority operation φ on \mathcal{D} is a set $\{\varphi_1, \dots, \varphi_n\}$, where φ_i is a one-sorted majority operation on D_i , i.e., a mapping from D_i^3 to D_i such that $\varphi_i(e, d, d) = \varphi_i(d, e, d) = \varphi_i(d, d, e) = d$ for all d, e in D_i .

An m -ary relation $R \subseteq D_{i_1} \times \dots \times D_{i_m}$ with $i_1, \dots, i_m \in \{1, 2, \dots, n\}$ is said to be closed under φ if $\varphi(t_1, t_2, t_3) = \langle \varphi_{i_1}(t_1[1], t_2[1], t_3[1]), \dots, \varphi_{i_m}(t_1[m], t_2[m], t_3[m]) \rangle$ is in R for any $t_1, t_2, t_3 \in R$.

A set of relations Γ is said to be closed under φ if every $R \in \Gamma$ is closed under φ .

A set of relations Γ is called a *majority-closed* language if there exists a (multi-sorted) majority operation φ such that every relation in Γ is closed under φ .

5.1 Tree-Preserving Constraints

The class of tree-preserving constraints is majority-closed.

Definition 9. [21] An undirected graph structure can often be associated to a finite domain D such that there is a bijection between the vertices in the graph and the values in D . If the graph is connected and acyclic, i.e. a tree, then we say it is a tree domain of x , denote as $T = (D, E)$ where E is a set of edges. A constraint R_{ij} over tree domains $T_i = (D_i, E_i)$ and $T_j = (D_j, E_j)$ is called tree-preserving if the image of every subtree in T_i is a subtree in T_j .

An example of tree-preserving constraint is shown in Fig. 7. A CRC constraint is a special tree-preserving constraint where related tree domains are chains [21].




Fig. 7: An example of tree-preserving constraint R_{ij} with $D_i = \{a, b\}$ and $D_j = \{e, f, g, h\}$, where a dashed arrow from a node u in T_i to a node v in T_j indicates that $\langle u, v \rangle \in R_{ij}$.

Definition 10. Let T_i be a nonempty tree domain for a variable v_i . The standard majority operation m_i on T_i is defined as:

$$(\forall a, b, c \in T_i) \quad m_i(a, b, c) = \pi_{a,b} \cap \pi_{b,c} \cap \pi_{a,c},$$

where a, b, c are not necessarily distinct and $\pi_{u,v}$ denotes the unique path from u to v in T_i .

The following result establishes the connection between tree-preserving constraints and majority-closed constraints.

Theorem 2. [21] Let T_i and T_j be two nonempty tree domains and m_i and m_j their standard majority operations. Suppose $R_{ij} \subseteq T_i \times T_j$ is a nonempty constraint such that both R_{ij} and R_{ji} are arc-consistent. Then R_{ij} is closed under $\{m_i, m_j\}$ iff both R_{ij} and R_{ji} are tree-preserving w.r.t. T_i and T_j .

5.2 Weak VEP Classes and Majority-Closed Classes

We first study a few properties of majority-closed classes.

Proposition 4. Let Γ be the set of binary relations that is closed under a multi-sorted majority operation $\varphi = \{\varphi_1, \dots, \varphi_n\}$ on $\mathcal{D} = \{D_1, \dots, D_n\}$. Then Γ is weakly closed under singletons.

Proof. Suppose R is a relation in Γ and $\langle a, b \rangle \in R \subseteq D_i \times D_j$. We show that $\{\langle a, b \rangle\}$ is closed under φ . For any $t_1, t_2, t_3 \in \{\langle a, b \rangle\}$, we have $t_1 = t_2 = t_3 = \langle a, b \rangle$, and, hence, $\varphi(t_1, t_2, t_3) = \langle \varphi_i(a, a, a), \varphi_j(b, b, b) \rangle = \langle a, b \rangle$. This shows that $\{\langle a, b \rangle\}$ is closed under φ and, hence, a relation in Γ . \square

Majority-closed relations are *decomposable*.

Definition 11. An m -ary relation R is said to be r -decomposable if, for any m -tuple t , $t \in R$ if $\pi_I(t) \in \pi_I(R)$ for all $I = (i_1, \dots, i_k)$ (a list of indices from $\{1, \dots, m\}$) with $k \leq r$, where $\pi_I(t) = \langle t[i_1], \dots, t[i_k] \rangle$ and $\pi_I(R) = \{\langle t[i_1], \dots, t[i_k] \rangle \mid t \in R\}$.

Theorem 3. [19] Let Γ be a set of binary relations over a set of finite domains $\mathcal{D} = \{D_1, \dots, D_n\}$. The following statements are equivalent:

- (1) Γ is a majority-closed constraint language.
- (2) Every $R \in \Gamma^+$ is 2-decomposable.
- (3) For every $\mathcal{N} \in \mathbf{CSP}(\Gamma)$, establishing strong PC in \mathcal{N} ensures global consistency.

Finally, we show that complete weak VEP classes are majority-closed classes.

Theorem 4. Let Γ be a complete set of binary relations over a set of finite domains $\mathcal{D} = \{D_1, \dots, D_n\}$. Then Γ has weak VEP iff it is a majority-closed class.

Proof. We first deal with the ‘only if’ part. Suppose that Γ is a complete set of binary relations that has weak VEP. By Theorem 3, we only need to show that for every BCN in $\mathbf{CSP}(\Gamma)$, establishing strong PC ensures global consistency. Let \mathcal{N}^0 be a network in $\mathbf{CSP}(\Gamma)$ and suppose $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ is the network obtained by enforcing strong PC on \mathcal{N}^0 . Since Γ is complete and thus closed under operations for achieving AC and PC by Lemma 1, \mathcal{N} is also a problem in $\mathbf{CSP}(\Gamma)$. Suppose \mathcal{N} is non-trivially inconsistent. We show that any partial solution of \mathcal{N} can be extended to a solution of \mathcal{N} .

Suppose $V' = \{v_1, \dots, v_{m-1}\} \subset V$ and $\sigma = \langle a_1, \dots, a_{m-1} \rangle$ is a solution of $\mathcal{N}|_{V'}$, which is the restriction of \mathcal{N} to V' . Assume further that $v_m \notin V'$ is a new variable and let $V'' = V' \cup \{v_m\}$. We show that σ can be consistently extended to $\mathcal{N}|_{V''}$, the restriction of \mathcal{N} to V'' . Because \mathcal{N} is strongly PC, $\mathcal{N}|_{V''}$ is strongly PC as well. In particular, v_i is AC relative to v_m for any R_{im} in C , and R_{ij} is PC relative to v_m (i.e., $R_{ij} \subseteq R_{im} \circ R_{mj}$) for any $i \neq j$ such that both R_{im} and R_{jm} are in C . By Definition 5, $\mathcal{N}|_{V'}$ is the same as $(\mathcal{N}|_{V''})_{-m}$, viz., the network obtained by eliminating v_m from $\mathcal{N}|_{V''}$. Moreover, since \mathcal{N} and, hence, $\mathcal{N}|_{V''}$ are AC, v_m is AC relative to all constraints R_{im} that are in C . By the assumption that Γ has weak VEP, σ can be consistently extended to v_m . Following this reasoning, we will find a solution of \mathcal{N} that extends σ .

Next, we consider the ‘if’ part. Suppose that Γ is a complete set of binary relations that is closed under some multi-sorted majority operation $\varphi = \{\varphi_1, \dots, \varphi_n\}$ on \mathcal{D} . Let $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$ be a problem in $\mathbf{CSP}(\Gamma)$ and v_x a variable in V . Let $E_x = \{R_{ix} \mid R_{ix} \in C\}$, and $\mathcal{N}_{-x} = \langle V \setminus \{v_x\}, \mathcal{D}, C' \rangle$, where $C' = (C \cup \{R_{ix} \circ R_{xj} \cap R_{ij} \mid R_{jx}, R_{ix} \in E_x\}) \setminus E_x$. Suppose that v_x is AC relative to all relations in E_x . We only need to show that any solution of \mathcal{N}_{-x} can be extended to a solution of \mathcal{N} . We prove this by contradiction.

Let σ be a solution of \mathcal{N}_{-x} . Assume that σ cannot be extended to a solution of \mathcal{N} . Therefore, E_x cannot be empty, otherwise σ can be trivially extended to a solution of \mathcal{N} . The case where $E_x = \{R_{ix}\}$ is a singleton is also impossible, as by (1), v_i is AC relative to R_{ix} and we could extend σ to a solution of \mathcal{N} by assigning any valid value to v_x . Suppose that E_x has $q \geq 2$ constraints and let them be $((v_1, v_x), R_1), \dots, ((v_q, v_x), R_q)$. We define a new problem $\mathcal{N}' = \langle V, \mathcal{D}, E_x \rangle$ as illustrated in Fig. 8. Since v_x is AC relative to all relations in E_x , it is easy to verify that \mathcal{N}' has a solution. For example, one can construct a solution of \mathcal{N}' by simply picking a value from D_x for v_x and then extending that valuation to v_1, \dots, v_q . Now, we construct a q -ary relation $R = \{\langle \gamma(v_1), \dots, \gamma(v_q) \rangle \mid \gamma \text{ is a solution of } \mathcal{N}'\}$. The solution set S of \mathcal{N}' can be




Fig. 8: Illustration of proof of Theorem 4.

obtained by using a sequence of the Cartesian product, equality selection, and projection operations [19]. Therefore, $S \in \Gamma^+$. Since $R = \pi_{v_1, \dots, v_q}(S)$, we have $R \in \Gamma^+$. By Theorem 3, R should be 2-decomposable; however, in the sequel we show that it is not, which is a contradiction.

Let $t = \langle \sigma(v_1), \dots, \sigma(v_q) \rangle$, where σ is a solution of \mathcal{N}_{-x} . It is clear that t is a solution of $\mathcal{N}'_{\{v_1, \dots, v_q\}}$. For any list of indices I chosen from $\{1, \dots, q\}$, with $|I| \leq 2$, we claim that $\pi_I(t) \in \pi_I(R)$. We recall that, for any two relations $R_{ix}, R_{jx} \in E_x$, the relation between v_i and v_j in \mathcal{N}_{-x} is $R_{ij} \cap (R_{ix} \circ R_{jx})$. Therefore, any partial solution $\langle \sigma(v_i), \sigma(v_j) \rangle (1 \leq i, j \leq q)$ of \mathcal{N}' can be consistently extended to v_x and, by the construction of \mathcal{N}' , further consistently extended to a solution of \mathcal{N}' . Thus, $\pi_I(t)$ is in $\pi_I(R)$ for any list of indices I chosen from $\{1, \dots, q\}$, with $|I| \leq 2$. However, $t \notin R$ because σ cannot be extended to a solution of \mathcal{N}' , which implies that R is not 2-decomposable. \square

6 The Variable Elimination Algorithm DPC*

This section presents a variant of DPC for solving BCNs defined over any weak VEP class. The new algorithm, called DPC* and presented as Algorithm 2, can solve problems that are not solvable by DPC (cf. Example 5 and Proposition 6). Compared with the variable elimination algorithm for solving CRC constraints [34], DPC* enforces a weaker AC condition instead of full AC. We first justify the correctness of Algorithm 2.

Theorem 5. *Let Γ be a complete weak VEP class. Suppose \mathcal{N} is a BCN defined over Γ and $\prec = (v_1, \dots, v_n)$ any ordering of variables of \mathcal{N} . Then, given \mathcal{N} and \prec , Algorithm 2 does not return “Inconsistent” iff \mathcal{N} is consistent.*

Proof. Suppose the input network \mathcal{N} is consistent. Since DPC* only prunes off certain infeasible domain values or relation tuples, the algorithm does not find any empty domains or relations in Lines 4, 9, and 19. Thus, it does not return “Inconsistent”.

Suppose the algorithm does not return “Inconsistent” and let $\mathcal{N}' = \langle V, \mathcal{D}', C' \rangle$ be the output network, where $\mathcal{D}' = \{D'_1, \dots, D'_n\}$. We show \mathcal{N}' is consistent.

Write $\mathcal{M}^{(0)}$ for \mathcal{N} and write $\mathcal{M}^{(i)}$ for the result of the i -th loop in calling DPC* for the input \mathcal{N} and $\prec = (v_1, v_2, \dots, v_n)$. Then $\mathcal{N}' = \mathcal{M}^{(n-1)}$ and all $\mathcal{M}^{(i)}$ ($0 \leq i < n$) are equivalent to \mathcal{N} . Let \mathcal{Q}_k be the restriction of $\mathcal{M}^{(k)}$ to $\{v_1, v_2, \dots, v_{n-k}\}$ ($0 \leq k < n$). In essence, \mathcal{Q}_k is obtained by eliminating v_{n-k+1} from \mathcal{Q}_{k-1} (Lines 2-5 or Lines 11-20), while also enforcing AC (Lines 7-10) for v_{n-k+1} relative to all its successors if

Algorithm 2: DPC*

Input : A binary constraint network $\mathcal{N} = \langle V, \mathcal{D}, C \rangle$;
an ordering $\prec = (v_1, \dots, v_n)$ on V .

Output: An equivalent subnetwork that is decomposable relative to \prec , or “Inconsistent”.

```
1 for  $k \leftarrow n$  to 1 do
2   if  $k$  has only one successor and that successor is  $i$  then
3      $D_i \leftarrow D_i \cap R_{ki}(D_k)$ 
4     if  $D_i = \emptyset$  then
5       return “Inconsistent”
6   else
7     for  $i < k$  with  $R_{ik} \in C$  do
8        $D_k \leftarrow D_k \cap R_{ik}(D_i)$ 
9       if  $D_k = \emptyset$  then
10        return “Inconsistent”
11    for  $i < k$  with  $R_{ik} \in C$  do
12       $R_{ik} \leftarrow R_{ik} \cap (D_i \times D_k)$ 
13      for  $j < i$  with  $R_{jk} \in C$  do
14         $R_{jk} \leftarrow R_{jk} \cap (D_j \times D_k)$ 
15        if  $R_{ij} \notin C$  then
16           $R_{ij} \leftarrow D_i \times D_j$ 
17           $C \leftarrow C \cup \{R_{ij}\}$ 
18         $R_{ij} \leftarrow R_{ij} \cap (R_{ik} \circ R_{kj})$ ;
19        if  $R_{ij} = \emptyset$  then
20          return “Inconsistent”
21 return  $\mathcal{N}$ .
```

it has more than one successor. Since Γ is a complete weak VEP class, every BCN defined over Γ has weak VEP. In particular, each \mathcal{Q}_{k-1} is defined over Γ and has weak VEP. This implies that every solution of \mathcal{Q}_k can be extended to a solution of \mathcal{Q}_{k-1} . Since no inconsistency is detected in the process, we have $D'_1 \neq \emptyset$ and thus \mathcal{Q}_{n-1} is consistent. By the above analysis, this implies that $\mathcal{Q}_{n-2}, \dots, \mathcal{Q}_1, \mathcal{Q}_0 = \mathcal{M}^{(0)} = \mathcal{N}$ are all consistent. \square

The preceding proof also gives a way to generate *all* solutions of a consistent input network *backtrack-free* by appropriately instantiating the variables along the input ordering \prec . Indeed, for all $1 \leq k < n$, a solution $\langle a_1, \dots, a_k \rangle$ of \mathcal{N}'_k can be extended to a solution $\langle a_1, \dots, a_{k+1} \rangle$ of \mathcal{N}'_{k+1} by choosing an element a_{k+1} from the intersection of all $R_{i,k+1}(a_i)$ with $i \leq k$ and $R_{i,k+1} \in C'$, which is always nonempty as shown in the preceding proof. As we know that if Γ is majority-closed, the completion of Γ is also majority-closed [19], and that complete majority-closed classes and complete weak VEP classes are equivalent by Theorem 4, this also proves the following result:




Fig. 9: A constraint network \mathcal{N} and its elimination \mathcal{N}_{-w} .

Proposition 5. *Suppose \mathcal{N} is a consistent BCN defined over a majority-closed class and $\prec = (v_1, \dots, v_n)$ an ordering of variables of \mathcal{N} . Then, given \mathcal{N} and \prec , Algorithm 2 returns an equivalent subnetwork \mathcal{N}' that is decomposable relative to \prec , i.e., any partial solution of \mathcal{N}' on $\{v_1, \dots, v_k\}$ for any $1 \leq k < n$ can be extended to a solution of \mathcal{N}' .*

Note that Lines 2-10 in **DPC*** do not achieve DAC of input networks. Therefore, **DPC*** does not achieve strong DPC. Since the overall runtime of Lines 2-10 is the same as enforcing DAC, this places **DPC*** in the same time complexity class as **DPC**. The following example, however, gives a BCN that can be solved by **DPC*** but not by **DPC**, which shows that the loop in Lines 7-10 is necessary.

Example 5. Let $D = \{a, b, c\}$ and φ be the majority operation on D such that for all $i, j, k \in D$, $\varphi(i, j, k) = a$ if $i \neq j$, $j \neq k$, and $i \neq k$, and $\varphi(i, j, k) = r$ otherwise, where r is the repeated value (e.g., $\varphi(b, c, b) = b$). Let $\Gamma = \{R_1, R_2, R_3, R_4, R_5, R_6\}$, where $R_1 = \{\langle a, a \rangle, \langle a, c \rangle\}$, $R_2 = \{\langle c, c \rangle, \langle c, b \rangle\}$, $R_3 = \{\langle b, b \rangle, \langle b, a \rangle\}$, $R_4 = \{\langle a, c \rangle\}$, $R_5 = \{\langle c, b \rangle\}$, and $R_6 = \{\langle a, b \rangle\}$. Every $R \in \Gamma$ is closed under the majority operation φ on D . Now, consider the constraint network $\mathcal{N} \in \mathbf{CSP}(\Gamma)$ as presented in Figure 9. Since $R_{xw} \circ R_{wz} = R_6$, $R_{xw} \circ R_{wy} = R_4$, and $R_{yw} \circ R_{wz} = R_5$, the eliminated network \mathcal{N}_{-w} is the same as the restriction of \mathcal{N} to the set of variables $\{v_x, v_y, v_z\}$. Let $\sigma(v_x) = a, \sigma(v_y) = c, \sigma(v_z) = b$. Then σ is a solution of \mathcal{N}_{-w} , but σ cannot be extended to a solution of \mathcal{N} . Thus, \mathcal{N} and hence Γ do not have VEP. By Theorem 2, **DPC** does not decide $\mathbf{CSP}(\Gamma)$.

On the other hand, since Γ is majority-closed, by Proposition 5, **DPC*** can correctly decide the consistency of \mathcal{N} . This observation is confirmed by calling the two algorithms on \mathcal{N} . Take the PEO $\prec = (w, x, y, z)$ as an example; the other PEOs are analogous. Let (\mathcal{N}, \prec) be an input to **DPC**. After processing w , we have $D_x = \{a\}$, $D_y = \{c\}$, $D_z = \{b\}$ and $R_{xy} = \{\langle a, c \rangle\}$, $R_{xz} = \{\langle a, b \rangle\}$, $R_{zy} = \{b, c\}$. We can observe that $\langle x = a, y = c, z = b \rangle$ is a solution to the eliminated subnetwork. Thus, if we keep running **DPC**, it will not detect inconsistency. On the other hand, for **DPC***, when eliminating w , **DPC*** makes w AC relative to its neighbors. Note that **DPC** does not perform this operation. After that, D_w is empty, and the algorithm will stop and output “Inconsistent”.

The following result shows that complete majority-closed classes over domains with more than two elements cannot have VEP.

Proposition 6. Let $\varphi = \{\varphi_1, \dots, \varphi_n\}$ be a majority operation on $\mathcal{D} = \{D_1, \dots, D_n\}$. If there exists a domain D_i in \mathcal{D} that contains more than two elements, then the set Γ_φ of binary relations that are closed under φ has neither the Helly property nor VEP.

Proof. Suppose a, b, c are three different values from D_i . It is easy to see that the relations $R_1 = \{\langle a, a \rangle, \langle a, b \rangle\}$, $R_2 = \{\langle a, b \rangle, \langle a, c \rangle\}$, and $R_3 = \{\langle a, a \rangle, \langle a, c \rangle\}$ are all closed under φ . Therefore, R_1 , R_2 , and R_3 are all in Γ_φ . Because any two of $R_1(a)$, $R_2(a)$, $R_3(a)$ have a common element but $R_1(a) \cap R_2(a) \cap R_3(a) = \emptyset$, this shows that Γ_φ does not have the Helly Property and, hence by Theorem 1, does not have VEP. \square

This result shows that no complete VEP class could have a domain with 3 or more values. Therefore, there are no interesting complete constraint languages except the boolean ones that can be decided by **DPC** (cf. Proposition 2), while all binary majority-closed classes (including CRC and tree-preserving constraints) can be decided by **DPC*** (cf. Proposition 5).

DPC* can also be used to solve majority-closed constraints of higher arities. This is because, by Theorem 3, every relation definable in a majority-closed language is 2-decomposable. Therefore, for each majority relation R of arity $m > 2$, if a constraint $c = ((y_1, \dots, y_m), R)$ appears in a constraint network \mathcal{N} , we could replace c with a set of binary constraints $c_{ij} = ((y_i, y_j) \mid \pi_{ij}(R))$ ($1 \leq i < j \leq m$), where $\pi_{ij}(R) = \{\langle t[y_i], t[y_j] \rangle \mid t \in R\}$.

7 Evaluations


In this section we experimentally compare algorithm **DPC*** against the state-of-the-art algorithms for solving majority-closed constraint networks. These are SAC3-SDS [4] and PC2001 [5]. SAC3-SDS is currently the best *singleton arc-consistency* (SAC) enforcing algorithm [12]. Enforcing either SAC¹ or PC correctly decides the consistency of a majority-closed constraint network [9,19].

Two different sets of data are considered for experiments, which are described as follows:


- (1) Tree-preserving constraint networks.
- (2) Random majority-closed constraint networks. These can be used to test the average performance of different algorithms. To generate such networks, we need to generate random majority-closed constraint languages as follows.
 - Randomly define a majority operation $\otimes_i : D_i^3 \rightarrow D_i$ for each domain $D_i \in \mathcal{D}$ as follows: for any $x, y, z \in D_i$,

$$\otimes_i(x, y, z) = \begin{cases} \text{any } v \in D_i, & \text{if } x, y, z \text{ are all different,} \\ \text{any repeated value of } x, y, z, & \text{otherwise.} \end{cases} \quad (2)$$


¹ *Singleton linear arc-consistency* (SLAC) is an alternative consistency notion that can be enforced to solve majority-closed constraint network [22], but no practical SLAC algorithms are available so far.




(a) Evaluations in the number n of variables. We set $\rho = 0.5$, $l = 0.3$, $d = 100$.



(b) Evaluations in the size d of domains. We set $\rho = 0.5$, $l = 0.3$, $n = 100$.



(c) Evaluations in the density ρ of networks. We set $l = 0.3$, $d = 100$, $n = 100$.



(d) Evaluations in the looseness l of constraints. We set $\rho = 0.5$, $d = 100$, $n = 100$.

Fig. 10: Performance comparisons among DPC*, SAC3-SDS and PC2001 for solving tree-preserving constraint networks.


- Randomly generate constraints $R_{ij} \subseteq D_i \times D_j$ and test whether

$$\{ \langle \otimes_i(t_x[1], t_y[1], t_z[1]), \otimes_j(t_x[2], t_y[2], t_z[2]) \rangle \mid t_x, t_y, t_z \in R_{ij} \} \subseteq R_{ij} \quad (3)$$


holds. By definition, R_{ij} is majority-closed under (\otimes_i, \otimes_j) iff (3) holds.

We used the model in [6,16] to generate random consistent constraint networks for experiments. These constraint networks were generated by varying four parameters: (1) the number of variables n , (2) the size of the domains d , (3) the density of the constraint networks ρ (i.e. the ratio of non-universal constraints to n^2) and (4) the looseness of constraints l (i.e. the ratio of the number of allowed tuples to d^2). We fix three of the four parameters and vary the remaining parameter. Experiments were carried out on a computer with an Intel Core i5-4570 processor with a 3.2 GHz frequency per CPU core, and 4 GB memory.


The graphs in Fig. 10 and Fig. 11 illustrate the experimental comparisons among algorithms DPC*, SAC3-SDS and PC2001 for solving tree-preserving and random




(a) Evaluations in the number n of variables. We set $\rho = 0.5$, $l = 0.3$, $d = 100$.



(b) Evaluations in the size d of domains. We set $\rho = 0.5$, $l = 0.3$, $n = 100$.



(c) Evaluations in the density ρ of networks. We set $l = 0.3$, $d = 100$, $n = 100$.



(d) Evaluations in the looseness l of constraints. We set $\rho = 0.5$, $d = 100$, $n = 100$.

Fig. 11: Performance comparisons among DPC*, SAC3-SDS and PC2001 for solving random majority-closed constraint networks.

majority-closed constraint networks respectively. The data points in each graph are CPU times averaged over 20 instances.

From Fig. 10 and Fig. 11, we observe that all algorithms behave similarly to one another when solving tree-preserving and random majority-closed constraint networks. Therefore, our analysis only focuses on Fig. 10 and the results are applicable to Fig. 11 as well.

We observe in Fig. 10a and Fig. 10b that all algorithms approximately show linear time behaviors with respect to n and d . On the other hand, Fig. 10c shows that PC2001 is not sensitive to the density of networks whereas DPC* and SAC3-SDS perform better when the density of networks is lower. Fig. 10d shows that the CPU time for DPC* remains almost unchanged when increasing the looseness of constraints. However, the CPU times for PC2001 and SAC3-SDS both go up and then drop down when increasing the looseness of constraints. Finally, we also observe in all the graphs in Fig. 10 that the performance differences among DPC*, PC2001 and SAC3-SDS are remarkable. DPC*

not only runs significantly faster than PC2001 and SAC3-SDS, but it also scales up to 7 times and 5 times better than PC2001 and SAC3-SDS respectively. This mainly owes to the fact that DPC^* is a single pass algorithm over the ordered input constraint networks.

8 Conclusion

This paper investigated which constraint satisfaction problems can be efficiently decided by enforcing directional path-consistency. Given a complete binary constraint language Γ , it turns out that DPC can decide $\text{CSP}(\Gamma)$ if Γ is defined over domains with less than three values. For a possibly incomplete binary constraint language Γ , we proved that Γ has the Helly property if, and only if, for any non-trivially inconsistent and triangulated binary constraint network \mathcal{N} over Γ , \mathcal{N} is consistent if it is strongly DPC relative to the inverse ordering of some perfect elimination ordering of the constraint graph of \mathcal{N} . The classes of row-convex [3] and tree-convex [35] constraints are examples of constraint classes which have the Helly property. More importantly, we presented the algorithm DPC^* , a simple variant of DPC , which can decide the CSP of any majority-closed constraint language, and is sufficient for guaranteeing backtrack-free search for majority-closed constraint networks, which have been found applications in various domains, such as scene labeling, temporal reasoning, geometric reasoning, and logical filtering. Our evaluations also show that DPC^* significantly outperforms the state-of-the-art algorithms for solving majority-closed constraint networks.

References

1. Barto, L.: The collapse of the bounded width hierarchy. *Journal of Logic and Computation* 26(3), 923–943 (2016)
2. Barto, L., Kozik, M.: Constraint satisfaction problems solvable by local consistency methods. *Journal of the ACM* 61(1), 3:1–3:19 (2014)
3. van Beek, P., Dechter, R.: On the minimality and global consistency of row-convex constraint networks. *Journal of the ACM* 42(3), 543–561 (1995)
4. Bessière, C., Cardon, S., Debruyne, R., Lecoutre, C.: Efficient algorithms for singleton arc consistency. *Constraints* 16(1), 25–53 (2011)
5. Bessière, C., Régin, J.C., Yap, R.H., Zhang, Y.: An optimal coarse-grained arc consistency algorithm. *Artificial Intelligence* 165(2), 165–185 (2005)
6. Bliet, C., Sam-Haroud, D.: Path consistency on triangulated constraint graphs. In: *IJCAI*. pp. 456–461 (1999)
7. Bulatov, A.A., Jeavons, P.: An algebraic approach to multi-sorted constraints. In: *CP*. pp. 183–198 (2003)
8. Carbonnel, C., Cooper, M.C.: Tractability in constraint satisfaction problems: A survey. *Constraints* 21(2), 115–144 (2016)
9. Chen, H., Dalmau, V., Grūbien, B.: Arc consistency and friends. *Journal of Logic and Computation* 23(1), 87–108 (2011)
10. Chmeiss, A., Jégou, P.: Efficient path-consistency propagation. *International Journal on Artificial Intelligence Tools* 7(02), 121–142 (1998)
11. Cohen, D., Jeavons, P.: The complexity of constraint languages. In: *Handbook of Constraint Programming*, pp. 169–204. Elsevier (2006)

12. Debruyne, R., Bessière, C.: Some practicable filtering techniques for the constraint satisfaction problem. In: IJCAI. pp. 412–417 (1997)
13. Dechter, R.: Constraint processing. Morgan Kaufmann (2003)
14. Dechter, R., Meiri, I., Pearl, J.: Temporal constraint networks. *Artificial Intelligence* 49(1-3), 61–95 (1991)
15. Dechter, R., Pearl, J.: Network-based heuristics for constraint-satisfaction problems. *Artificial Intelligence* 34(1), 1–38 (1987)
16. Deville, Y., Barette, O., van Hentenryck, P.: Constraint satisfaction over connected row convex constraints. *Artificial Intelligence* 109(1-2), 243–271 (1999)
17. Freuder, E.C.: A sufficient condition for backtrack-free search. *Journal of the ACM* 29(1), 24–32 (1982)
18. Fulkerson, D., Gross, O.: Incidence matrices and interval graphs. *Pacific Journal of Mathematics* 15(3), 835–855 (1965)
19. Jeavons, P., Cohen, D., Cooper, M.C.: Constraints, consistency and closure. *Artificial Intelligence* 101(1), 251–265 (1998)
20. Kong, S., Li, S., Li, Y., Long, Z.: On tree-preserving constraints. In: CP. pp. 244–261 (2015)
21. Kong, S., Li, S., Li, Y., Long, Z.: On tree-preserving constraints. *Annals of Mathematics and Artificial Intelligence* (2017)
22. Kozik, M.: Weak consistency notions for all the csps of bounded width. In: LICS. pp. 633–641 (2016)
23. Kumar, T.K.S.: On geometric CSPs with (near)-linear domains and max-distance constraints. In: Workshop on Modelling and Solving Problems with Constraints. pp. 125–138 (2004)
24. Kumar, T.K.S.: On the tractability of restricted disjunctive temporal problems. In: ICAPS. pp. 110–119 (2005)
25. Kumar, T.K.S., Russell, S.J.: On some tractable cases of logical filtering. In: ICAPS. pp. 83–92 (2006)
26. Li, S., Liu, W., Wang, S.: Qualitative constraint satisfaction problems: an extended framework with landmarks. *Artificial Intelligence* 201, 32–58 (2013)
27. Mackworth, A.K.: Consistency in networks of relations. *Artificial Intelligence* 8(1), 99–118 (1977)
28. Maruyama, H.: Structural disambiguation with constraint propagation. In: ACL. pp. 31–38 (1990)
29. Mohr, R., Henderson, T.C.: Arc and path consistency revisited. *Artificial Intelligence* 28(2), 225–233 (1986)
30. Montanari, U.: Networks of constraints: fundamental properties and applications to picture processing. *Information Science* 7, 95–132 (1974)
31. Planken, L., de Weerd, M., van der Krogt, R.: P³C: A new algorithm for the simple temporal problem. In: ICAPS. vol. 2008, pp. 256–263 (2008)
32. Singh, M.: Path consistency revisited. *International Journal on Artificial Intelligence Tools* 5(1-2), 127–142 (1996)
33. Sioutis, M., Long, Z., Li, S.: Efficiently reasoning about qualitative constraints through variable elimination. In: SETN. pp. 1:1–1:10 (2016)
34. Zhang, Y., Marisetti, S.: Solving connected row convex constraints by variable elimination. *Artificial Intelligence* 173(12), 1204–1219 (2009)
35. Zhang, Y., Yap, R.H.C.: Consistency and set intersection. In: IJCAI. pp. 263–270 (2003)