

## View-based Propagator Derivation

Christian Schulte · Guido Tack

**Abstract** When implementing a propagator for a constraint, one must decide about variants: When implementing  $\min$ , should one also implement  $\max$ ? Should one implement linear constraints both with unit and non-unit coefficients? Constraint variants are ubiquitous: implementing them requires considerable (if not prohibitive) effort and decreases maintainability, but will deliver better performance than resorting to constraint decomposition.

This paper shows how to use *views* to *derive* propagator variants, combining the efficiency of dedicated propagator implementations with the simplicity and effortlessness of decomposition.

A model for views and derived propagators is introduced. Derived propagators are proved to be *perfect* in that they inherit essential properties such as correctness and domain and bounds consistency. Techniques for systematically deriving propagators such as transformation, generalization, specialization, and type conversion are developed. The paper introduces an implementation architecture for views that is independent of the underlying constraint programming system. A detailed evaluation of views implemented in Gecode shows that derived propagators are efficient and that views often incur no overhead. Views have proven essential for implementing Gecode, substantially reducing the amount of code that needs to be written and maintained.

### 1 Introduction

When implementing a propagator for a constraint, one typically must also decide whether to implement some of its variants. When implementing a propagator for the constraint  $\max\{x_1, \dots, x_n\} = y$ , should one also implement  $\min\{x_1, \dots, x_n\} = y$ ? The latter can be implemented using the former as  $\max\{-x_1, \dots, -x_n\} = -y$ . When implementing a propagator for the linear equation  $\sum_{i=1}^n a_i x_i = k$  for integer variables  $x_i$  and integers  $a_i$  and  $k$ ,

---

Christian Schulte  
ICT, KTH - Royal Institute of Technology, Sweden  
E-mail: cschulte@kth.se

Guido Tack  
National ICT Australia (NICTA) and Faculty of Information Technology, Monash University, Caulfield, Victoria, Australia  
E-mail: guido.tack@monash.edu

should one also implement the special case  $\sum_{i=1}^n x_i = k$  for better performance? When implementing a propagator for the reified linear equation  $(\sum_{i=1}^n x_i = c) \leftrightarrow b$ , should one also implement  $(\sum_{i=1}^n x_i \neq c) \leftrightarrow b$ ? These two constraints only differ by the sign of  $b$ , as the latter is equivalent to  $(\sum_{i=1}^n x_i = c) \leftrightarrow \neg b$ .

The two straightforward approaches for implementing constraint variants are to either implement dedicated propagators for the variants, or to decompose. In the last example, for instance, the reified constraint could be decomposed into two propagators, one for  $(\sum_{i=1}^n x_i = c) \leftrightarrow b'$ , and one for  $b \leftrightarrow \neg b'$ , introducing an additional variable  $b'$ .

Implementing the variants inflates code and documentation and is error prone. Given the potential code explosion, one may be able to only implement some variants (say, min and max). Other variants important for performance (say, ternary min and max) may be infeasible due to excessive programming and maintenance effort. Decomposing, on the other hand, massively increases memory consumption and runtime.

A third approach, and the one we follow in this paper, is to *derive* propagators, combining the efficiency of dedicated propagator implementations with the simplicity and effortlessness of decomposition.

Techniques for deriving propagators have been used in constraint programming systems for a long time. Some techniques have been the topic of extensive previous research, like the *indexicals* approach, which derives propagators from range expressions [5,43]. Other similar techniques are part of the constraint programming folklore and exist in implementations such as integer *expressions* in IBM ILOG CP Optimizer [24] and Google or-tools [19], which can be regarded as generalized bi-directional indexicals, or *literals* in SAT solvers [12], which combine Boolean variables with negation.

In this paper, we consolidate this long line of work by presenting a model for propagator derivation that allows us to analyze the properties of derived propagators, and by discussing and evaluating implementation strategies.

We call the basic building block for propagator derivation a *view*. A view can be regarded as a restricted form of a bi-directional indexical or expression, where the restrictions have been chosen carefully such that the resulting derived propagator satisfies important properties concerning correctness and effectiveness, and such that the implementation does not incur any overhead.

The following example shows how to derive a propagator using views.

**Example 1 (Deriving a minimum propagator)** Consider a propagator for the constraint  $\max(x, y) = z$ . Given three additional propagators for  $x' = -x$ ,  $y' = -y$ , and  $z' = -z$ , we could propagate the constraint  $\min(x', y') = z'$  using the propagator for  $\max(x, y) = z$ . Instead of these three additional propagators, we will derive a propagator for max from the propagator for min using views that perform the simple negation transformations.

Views transform input and output of a propagator. For example, a minus view on a variable  $x$  transforms the variable domain of  $x$  by negating each element, passes the transformed domain to the propagator, and performs the inverse transformation on the domain returned by the propagator. With views, the implementation of the maximum propagator can be reused: a propagator for the minimum constraint can be derived from a propagator for the maximum constraint and a minus view for each variable. \*

While the underlying technique is not new, this paper presents the first thorough academic discussion of deriving propagators using views. We systematically explore the *scope* of this technique, establish the *properties* of derived propagators, discuss the *limitations* given the desired properties, and present and evaluate different *implementation techniques*

including *compile-time* resolution of views. The evaluation shows that views are widely applicable, drastically reduce programming effort, and yield efficient propagators in practice.

More specifically, we identify the properties of views that are essential for deriving *perfect* propagators. The paper establishes a formal model that defines a view as a function and a derived propagator as functional composition of views (mapping values to values) with a propagator (mapping domains to domains). This model enables us to reason about derived propagators: they are indeed propagators; they faithfully implement the intended constraints; domain consistency carries over to derived propagators; different forms of bounds consistency over integer variables carry over provided that the views satisfy additional yet natural properties.

We introduce techniques for deriving propagators that use views for transformation, generalization, specialization, and type conversion of propagators. We show how to apply these techniques for different variable domains using various views and how views can be used for the derivation of dual scheduling propagators.

We present and evaluate different implementation approaches for views and derived propagators. An implementation using parametric polymorphism (such as templates in C++) is shown to incur no or very low overhead. The architecture is orthogonal to the used constraint programming system and has been fully implemented in Gecode [14]. Views have proven essential for implementing Gecode – the technique is applicable to most propagators and results in a substantial reduction of the amount of code that needs to be written and maintained.

**Plan of the paper.** Section 2 introduces constraints and propagators. Section 3 establishes views and propagator derivation. Section 4 presents techniques for propagator derivation. Section 5 discusses related approaches. Section 6 describes an implementation architecture based on parametric propagators and range iterators. Section 7 discusses limitations of views. The implementation is evaluated in Section 8, and Section 9 concludes the paper.

## 2 Preliminaries

This section introduces the notation used in the rest of the paper for constraints, propagators, and propagation strength.

**Variables, constraints, and domains.** Constraint satisfaction problems use a *finite set of variables*  $X$  and a *finite set of values*  $V$ . We typically write variables as  $x, y, z \in X$  and values as  $v, w \in V$ .

A solution of a constraint satisfaction problem assigns a single value to each variable. A constraint restricts which assignments of values to variables are allowed.

**Definition 1 (Assignments and constraints)** An *assignment*  $a$  is a function mapping variables to values. The set of all assignments is  $\text{Asn} = X \rightarrow V$ . A *constraint*  $c$  is a set of assignments,  $c \in \text{Con} = \mathcal{P}(\text{Asn}) = \mathcal{P}(X \rightarrow V)$  (we write  $\mathcal{P}(S)$  for the power set of  $S$ ). Any assignment  $a \in c$  is a *solution* of  $c$ . \*

Constraints are defined on assignments as total functions on all variables. For a typical constraint  $c$ , only a subset  $\text{vars}(c)$  of the variables is *significant*; the constraint is the full relation for all  $x \notin \text{vars}(c)$ . Constraints are either written as sets of assignments (for example,  $\{a \in \text{Asn} \mid a(x) < a(y)\}$ ) or as expressions with the usual meaning, using the notation  $\llbracket \cdot \rrbracket$  (for example,  $\llbracket x < y \rrbracket$ ).

**Example 2 (Sum constraint)** Let  $X = \{x, y, z\}$  and  $V = \{1, 2, 3, 4\}$ . The sum constraint  $\llbracket x = y + z \rrbracket$  corresponds to the following set of assignments:

$$\begin{aligned} \llbracket x = y + z \rrbracket &= \{(x \mapsto a, y \mapsto b, z \mapsto c) \mid a, b, c \in V \wedge a = b + c\} \\ &= \{(x \mapsto 2, y \mapsto 1, z \mapsto 1), (x \mapsto 3, y \mapsto 1, z \mapsto 2), \\ &\quad (x \mapsto 3, y \mapsto 2, z \mapsto 1), (x \mapsto 4, y \mapsto 2, z \mapsto 2), \\ &\quad (x \mapsto 4, y \mapsto 1, z \mapsto 3), (x \mapsto 4, y \mapsto 3, z \mapsto 1)\} \end{aligned} \quad *$$

**Definition 2 (Domains)** A *domain*  $d$  is a function mapping variables to sets of values, such that  $d(x) \subseteq V$ . The set of all domains is  $\text{Dom} = X \rightarrow \mathcal{P}(V)$ . The set of values in  $d$  for a particular variable  $x$ ,  $d(x)$ , is called the *variable domain* of  $x$ . The following definition expresses that any domain  $d$  can be regarded as a set of assignments, and thus a constraint:

$$\text{con}(d) = \{a \in \text{Asn} \mid \forall x \in X : a(x) \in d(x)\}$$

An assignment  $a \in \text{con}(d)$  is *licensed* by  $d$ . \*

Domains thus represent *Cartesian* sets of assignments. In this sense, any domain is also a constraint. For a more uniform representation, we take the liberty to use domains as constraints. In particular,  $a \in d$  (instead of  $a \in \text{con}(d)$ ) denotes an assignment  $a$  licensed by  $d$ , and  $c \cap d$  denotes  $c \cap \text{con}(d)$ .

A domain  $d$  that maps some variable to the empty value set is *failed*, written  $d = \emptyset$ , as it represents no valid assignments ( $\text{con}(d) = \emptyset$ ). A domain  $d$  representing a single assignment,  $\text{con}(d) = \{a\}$ , is *assigned*, and is written as  $d = \{a\}$ .

**Definition 3 (Constraint satisfaction problems)** A *constraint satisfaction problem* (CSP) is a pair  $\langle d, C \rangle$  of a domain  $d$  and a set of constraints  $C$ . The *solutions* of a CSP  $\langle d, C \rangle$  are the assignments licensed by  $d$  that satisfy all constraints in  $C$ , defined as  $\text{sol}(\langle d, C \rangle) = \{a \in \text{con}(d) \mid \forall c \in C : a \in c\}$ . \*

**Propagators.** A propagation-based constraint solver employs *propagators* to implement constraints. A propagator for a constraint  $c$  takes a domain  $d$  as input and removes values from the variable domains in  $d$  that are in conflict with  $c$ .

A domain  $d$  is *stronger* than a domain  $d'$ , written  $d \subseteq d'$ , if and only if  $d(x) \subseteq d'(x)$  for all  $x \in X$ . A domain  $d$  is *strictly stronger* than a domain  $d'$ , written  $d \subset d'$ , if and only if  $d$  is stronger than  $d'$  and  $d(x) \subset d'(x)$  for some variable  $x$ . The goal of constraint propagation is to prune values from variable domains, thus inferring stronger domains, without removing solutions of the constraints.

A propagator is a function  $p$  that takes a domain as its argument and returns a stronger domain, it may only *prune* assignments. If the original domain is an assigned domain  $\{a\}$ , the propagator either accepts ( $p(\{a\}) = \{a\}$ ) or rejects ( $p(\{a\}) = \emptyset$ ) it. We can therefore say that a propagator realizes a *decision procedure* for the unique constraint defined by all accepted assignments. The pruning and the decision procedure must be consistent: if the decision procedure accepts an assignment, the pruning procedure must never remove this assignment from any domain. This property is enforced by requiring propagators to be monotonic. The following definition captures these properties.

**Definition 4 (Propagators)** A *propagator* is a function  $p \in \text{Dom} \rightarrow \text{Dom}$  that is

- *contracting*:  $p(d) \subseteq d$  for any domain  $d$ ;

- *monotonic*:  $p(d') \subseteq p(d)$  for any domains  $d' \subseteq d$ .

The set of all propagators is  $\text{Prop}$ . If a propagator  $p$  returns a *strictly* stronger domain ( $p(d) \subset d$ ), we say that  $p$  *prunes the domain*  $d$ . The propagator  $p$  *induces* the unique constraint  $c_p$  defined by the set of assignments accepted by  $p$ :

$$c_p = \{a \in \text{Asn} \mid p(\{a\}) = \{a\}\} \quad *$$

Propagators can also be *idempotent* ( $p(p(d)) = p(d)$  for any domain  $d$ ). Idempotency is not required to make propagation sound or complete, but it can make propagation more efficient [40]. Like idempotency, monotonicity as defined here is not necessary for soundness or completeness of a solver [41]. Most definitions and theorems in this paper are independent of whether propagators are monotonic or not. Non-monotonicity will thus only be discussed where it is relevant.

The propagator defined as  $p_U(d) = d$  induces the universal constraint (i.e., it does not prune anything), while  $p_\emptyset(d) = \emptyset$  induces the empty constraint (which does not permit any solution). Note that, contrary to other popular definitions,  $p_U$  is not a “valid” propagator for arbitrary constraints, it merely induces a constraint that is weaker than any other constraint.

**Propagation strength.** Each propagator induces a single constraint, but different propagators inducing the same constraint can differ in *strength*. Typical examples are propagators for the *all-different* constraint that perform naive pruning when variables are assigned, or establish bounds consistency [34] or domain consistency [37].

In the literature, propagation strength is usually defined as a property of a domain in relation to a constraint. For example, a domain  $d$  is *domain-consistent* (also known as generalized arc-consistent) with respect to a constraint  $c$  if  $d(x)$  only contains values that appear in at least one solution of  $c$  for each variable  $x$ . As this paper is concerned with propagators, propagation strength is defined with respect to a propagator.

A propagator  $p$  is *domain-complete* if any domain it returns is domain-consistent with respect to  $c_p$ . For any constraint  $c$ , there is exactly one domain-complete propagator for  $c$  (as domains form a lattice). It is defined as  $\hat{p}_c(d) = \text{dom}(c_p \cap d)$ , where  $\text{dom}(c)$  is the *domain relaxation* of  $c$ , the strongest domain that contains all assignments of  $c$ ,  $\text{dom}(c) = \min\{d \mid c \subseteq d\}$ .

For constraints over integer variables ( $V \subseteq \mathbb{Z}$ ), several weaker notions of propagation strength are known. The most well-known is *bounds consistency*, which in fact can mean one of four special cases: range, bounds( $D$ ), bounds( $\mathbb{Z}$ ), and bounds( $\mathbb{R}$ ) consistency (as discussed in [7, 35]).

The first three differ in whether holes are ignored in the original domain, in the resulting domain, or in both, in that order. Holes in a domain are ignored by the function  $\text{hull}(d)(x) = [\min(d(x)) .. \max(d(x))]$ , which returns the convex hull of a variable domain  $d(x)$  in  $\mathbb{Z}$ . Bounds( $\mathbb{R}$ ) consistency only requires solutions to be found in the real-valued relaxation of the constraint (written  $c_{\mathbb{R}}$ ), and is defined using the real-valued convex hull and domain relaxation (written  $\text{hull}_{\mathbb{R}}$  and  $\text{dom}_{\mathbb{R}}$ ). The different notions of bounds consistency give rise to the respective definitions of bounds completeness.

**Definition 5 (Bounds completeness)** A propagator  $p$  is

- range-complete if and only if  $p(d) \subseteq \text{dom}(c_p \cap \text{hull}(d))$ ,
- bounds( $D$ )-complete if and only if  $p(d) \subseteq \text{hull}(\text{dom}(c_p \cap d))$ ,
- bounds( $\mathbb{Z}$ )-complete if and only if  $p(d) \subseteq \text{hull}(\text{dom}(c_p \cap \text{hull}(d)))$ , and
- bounds( $\mathbb{R}$ )-complete if and only if  $p(d) \subseteq \text{hull}_{\mathbb{R}}(\text{dom}_{\mathbb{R}}(c_{p_{\mathbb{R}}} \cap \text{hull}_{\mathbb{R}}(d)))$

for any domain  $d$ . \*

### 3 Views

This section defines the formal model for views and derived propagators. Based on this model, the section proves essential properties of propagators derived using views, such as correctness and completeness. Derived propagators are thus indeed *perfect* with respect to the formal model.

#### 3.1 Views and Derived Propagators

Given a propagator  $p$ , a view is represented by two functions,  $\varphi$  and  $\varphi^-$ , that can be composed with  $p$  such that  $\varphi^- \circ p \circ \varphi$  is the desired derived propagator. The function  $\varphi$  transforms the input domain, and  $\varphi^-$  applies the inverse transformation to the propagator's output domain.

**Definition 6 (Variable views and views)** A *variable view*  $\varphi_x \in V \rightarrow V'$  for a variable  $x$  is an injective function mapping values to values. The set  $V'$  may be different from  $V$ , and the corresponding sets of assignments, domains, constraints, and propagators are called  $\text{Asn}'$ ,  $\text{Dom}'$ ,  $\text{Con}'$ , and  $\text{Prop}'$ , respectively.

Given a family of variable views  $\varphi_x$  for all  $x \in X$ , we lift them point-wise to assignments:  $\varphi_{\text{Asn}}(a)(x) = \varphi_x(a(x))$ . A *view*  $\varphi \in \text{Con} \rightarrow \text{Con}'$  is a family of variable views, lifted to constraints:  $\varphi(c) = \{\varphi_{\text{Asn}}(a) \mid a \in c\}$ . The *inverse* of a view is defined as  $\varphi^-(c) = \{a \in \text{Asn} \mid \varphi_{\text{Asn}}(a) \in c\}$ . \*

**Definition 7 (Derived propagators and constraints)** Given a propagator  $p \in \text{Prop}'$  and a view  $\varphi$ , the *derived propagator*  $\widehat{\varphi}(p) \in \text{Prop}$  is defined as  $\widehat{\varphi}(p) = \varphi^- \circ p \circ \varphi$ . Similarly, a *derived constraint* is defined to be  $\varphi^-(c) \in \text{Con}$  for a given  $c \in \text{Con}'$ . \*

**Example 3 (Scale views)** Given a propagator  $p$  for the constraint  $c = \llbracket x = y \rrbracket$ , we want to derive a propagator for  $c' = \llbracket x = 2y \rrbracket$  using a view  $\varphi$  such that  $\varphi^-(c) = c'$ .

Intuitively, the function  $\varphi$  leaves  $x$  as it is and scales  $y$  by 2, while  $\varphi^-$  does the inverse transformation. We thus define  $\varphi_x(v) = v$  and  $\varphi_y(v) = 2v$ . That clarifies the need for different sets  $V$  and  $V'$ , as  $V'$  must contain all elements of  $V$  multiplied by 2.

The derived propagator is  $\widehat{\varphi}(p) = \varphi^- \circ p \circ \varphi$ . We say that  $\widehat{\varphi}(p)$  “uses a scale view on”  $y$ , meaning that  $\varphi_y$  is the function defined as  $\varphi_y(v) = 2v$ . Similarly, using an identity view on  $x$  amounts to  $\varphi_x$  being the identity function on  $V$ .

Given the assignment  $a = (x \mapsto 2, y \mapsto 1)$ , we first apply  $\varphi$  and get  $\varphi(\{a\}) = \{(x \mapsto 2, y \mapsto 2)\}$ . This is accepted by  $p$  and returned unchanged, so  $\varphi^-$  transforms it back to  $\{a\}$ . Another assignment  $a' = (x \mapsto 1, y \mapsto 2)$  is transformed to  $\varphi(\{a'\}) = \{(x \mapsto 1, y \mapsto 4)\}$ , rejected ( $p(\varphi(\{a'\})) = \emptyset$ ), and the empty domain is mapped to the empty domain by  $\varphi^-$ . The propagator  $\widehat{\varphi}(p)$  induces  $\varphi^-(c)$ . \*

#### 3.2 Correctness of Derived Propagators

Derived propagators are well-defined and correct: a derived propagator  $\widehat{\varphi}(p)$  is in fact a propagator, and it induces the desired constraint ( $c_{\widehat{\varphi}(p)} = \varphi^-(c_p)$ ). The proofs of these statements employ the following direct consequences of the definitions of views:

- P1.  $\varphi$  and  $\varphi^-$  are monotonic by construction (as  $\varphi$  and  $\varphi^-$  are defined point-wise).
- P2.  $\varphi^- \circ \varphi = \text{id}$  (the identity function, as  $\varphi$  is injective).

P3.  $|\varphi(\{a\})| = 1$ ,  $\varphi(\emptyset) = \emptyset$ .

P4. For any view  $\varphi$  and domain  $d \in \text{Dom}$ , we have  $\varphi(d) \in \text{Dom}'$ , and for any  $d' \in \text{Dom}'$ , we have  $\varphi^-(d') \in \text{Dom}$  (as views are defined point-wise).

**Proposition 1 (Correctness)** For a propagator  $p$  and view  $\varphi$ ,  $\widehat{\varphi}(p)$  is a propagator. \*

*Proof.* The derived propagator is well-defined because both  $\varphi(d)$  and  $\varphi^-(d)$  are domains (see P4 above). We have to show that  $\varphi^- \circ p \circ \varphi$  is contracting and monotonic.

For contraction, we have  $p(\varphi(d)) \subseteq \varphi(d)$  as  $p$  is contracting. From monotonicity of  $\varphi^-$  (with P1), it follows that  $\varphi^-(p(\varphi(d))) \subseteq \varphi^-(\varphi(d))$ . As  $\varphi^- \circ \varphi = \text{id}$  (with P2), we have  $\varphi^-(p(\varphi(d))) \subseteq d$ , which proves that  $\widehat{\varphi}(p)$  is contracting.

Monotonicity is shown as follows for all domains  $d', d$  with  $d' \subseteq d$ :

$$\begin{aligned} & \varphi(d') \subseteq \varphi(d) && (\varphi \text{ monotonic, P1}) \\ \implies & p(\varphi(d')) \subseteq p(\varphi(d)) && (p \text{ monotonic}) \\ \implies & \varphi^-(p(\varphi(d'))) \subseteq \varphi^-(p(\varphi(d))) && (\varphi^- \text{ monotonic, P1}) \end{aligned}$$

In summary, for any propagator  $p$ ,  $\widehat{\varphi}(p) = \varphi^- \circ p \circ \varphi$  is a propagator. ■

Non-monotonic propagators as defined in [41] must at least be *weakly* monotonic, which means that  $p(\{a\}) \subseteq p(d)$  for all domains  $d$  and assignments  $a \in d$ . The above proof can be easily adjusted to weakly monotonic propagators by replacing  $d'$  with  $\{a\}$  and using P3 in the second line of the proof.

**Proposition 2 (Induced constraints)** Let  $p$  be a propagator, and let  $\varphi$  be a view. Then  $\widehat{\varphi}(p)$  induces the constraint  $\varphi^-(c_p)$ . \*

*Proof.* As  $p$  induces  $c_p$ , we know  $p(\{a\}) = c_p \cap \{a\}$  for all assignments  $a$ . With  $|\varphi(\{a\})| = 1$  (P3), we have  $p(\varphi(\{a\})) = c_p \cap \varphi(\{a\})$ . Furthermore, we know that  $c_p \cap \varphi(\{a\})$  is either  $\emptyset$  or  $\varphi(\{a\})$ .

- *Case  $\emptyset$ :* We have  $\varphi^-(p(\varphi(\{a\}))) = \emptyset = \{a' \in \text{Asn} \mid a = a' \wedge \varphi_{\text{Asn}}(a) \in c_p\} = \varphi^-(c_p) \cap \{a\}$ .
- *Case  $\varphi(\{a\})$ :* With P2, we have  $\varphi^- \circ \varphi = \text{id}$  and hence  $\varphi^-(p(\varphi(\{a\}))) = \{a\}$ . Furthermore,  $\varphi^-(c_p) \cap \{a\} = \{a' \in \text{Asn} \mid a = a' \wedge \varphi_{\text{Asn}}(a) \in c_p\} = \{a\}$ .

Together, this shows that  $\varphi^- \circ p \circ \varphi(\{a\}) = \{a\} \cap \varphi^-(c_p)$ . ■

Another important property is that views preserve contraction: if a propagator  $p$  prunes a domain, the pruning will not be lost after transformation by  $\varphi^-$ .

**Proposition 3 (Views preserve contraction)** Let  $p$  be a propagator, let  $\varphi$  be a view, and let  $d$  be a domain such that  $p(\varphi(d)) \subset \varphi(d)$ . Then  $\widehat{\varphi}(p)(d) \subset d$ . \*

*Proof.* The definition of  $\varphi^-(c)$  is  $\{a \in \text{Asn} \mid \varphi_{\text{Asn}}(a) \in c\}$ . Hence,  $|\varphi^-(c)| \leq |c|$ . Similarly, we know that  $|\varphi(c)| = |c|$ . From  $p(\varphi(d)) \subset \varphi(d)$ , it follows that  $|p(\varphi(d))| < |\varphi(d)|$ . Together, this yields  $|\widehat{\varphi}(p)(d)| < |\varphi(d)| = |d|$ . We have seen in Proposition 1 that  $\widehat{\varphi}(p)(d) \subseteq d$ , so we can conclude that  $\widehat{\varphi}(p)(d) \subset d$ . ■

### 3.3 Completeness of Derived Propagators

Ideally, a propagator derived from a domain- or bounds-complete propagator should inherit its completeness. Unfortunately this turns out not to be true in general for all notions of completeness and all views. We now first show how  $\text{bounds}(\mathbb{Z})$  completeness is inherited, and then generalize this result to the other notions.

The key insight is that  $\text{bounds}(\mathbb{Z})$  completeness of propagators derived using a view  $\varphi$  depends on whether  $\varphi$  and  $\varphi^-$  commute with the hull operator, as defined below.

**Definition 8** A constraint  $c$  is a  $\varphi$ -constraint for a view  $\varphi$  if and only if for all  $a \in c$ , there is a  $b \in \text{Asn}$  such that  $a = \varphi_{\text{Asn}}(b)$ . A view  $\varphi$  is *hull-injective* if and only if  $\varphi^-(\text{hull}(\text{dom}(c))) = \text{hull}(\text{dom}(\varphi^-(c)))$  for all  $\varphi$ -constraints  $c$ . It is *hull-surjective* if and only if  $\varphi(\text{hull}(d)) = \text{hull}(\varphi(d))$  for all domains  $d$ . It is *hull-bijective* if and only if it is hull-injective and hull-surjective. \*

The proofs rely on the additional fact that views commute with set intersection.

**Lemma 1** For any view  $\varphi$ , the equation  $\varphi^-(c_1 \cap c_2) = \varphi^-(c_1) \cap \varphi^-(c_2)$  holds. \*

*Proof.* By definition of  $\varphi^-$ , we have

$$\varphi^-(c_1 \cap c_2) = \{a \in \text{Asn} \mid \varphi_{\text{Asn}}(a) \in c_1 \wedge \varphi_{\text{Asn}}(a) \in c_2\}$$

As  $\varphi_{\text{Asn}}$  is a function, this is equal to

$$\{a \in \text{Asn} \mid \varphi_{\text{Asn}}(a) \in c_1\} \cap \{a \in \text{Asn} \mid \varphi_{\text{Asn}}(a) \in c_2\} = \varphi^-(c_1) \cap \varphi^-(c_2) \quad \blacksquare$$

**Theorem 1** (**Bounds( $\mathbb{Z}$ ) completeness**) Let  $p$  be a  $\text{bounds}(\mathbb{Z})$ -complete propagator. For any hull-bijective view  $\varphi$ , the propagator  $\widehat{\varphi}(p)$  is  $\text{bounds}(\mathbb{Z})$ -complete. \*

*Proof.* From [Proposition 2](#), we know that  $\widehat{\varphi}(p)$  induces the constraint  $\varphi^-(c_p)$ . By monotonicity of  $\varphi^-$  (with [P1](#)) and  $\text{bounds}(\mathbb{Z})$  completeness of  $p$ , we know that

$$\varphi^- \circ p \circ \varphi(d) \subseteq \varphi^-(\text{hull}(\text{dom}(c_p \cap \text{hull}(\varphi(d))))))$$

We now use the fact that both  $\varphi^-$  and  $\varphi$  commute with  $\text{hull}(\cdot)$  and set intersection:

$$\begin{aligned} & \varphi^-(\text{hull}(\text{dom}(c_p \cap \text{hull}(\varphi(d)))))) \\ &= \varphi^-(\text{hull}(\text{dom}(c_p \cap \varphi(\text{hull}(d)))))) && \text{(hull-surjective)} \\ &= \text{hull}(\text{dom}(\varphi^-(c_p \cap \varphi(\text{hull}(d)))))) && \text{(hull-injective)} \\ &= \text{hull}(\text{dom}(\varphi^-(c_p) \cap \varphi^-(\varphi(\text{hull}(d)))))) && \text{(commute with } \cap \text{)} \\ &= \text{hull}(\text{dom}(\varphi^-(c_p) \cap \text{hull}(d))) && \text{(P2)} \end{aligned}$$

The second step uses hull injectivity, so it requires  $c_p \cap \varphi(\text{hull}(d))$  to be a  $\varphi$ -constraint. All assignments in a  $\varphi$ -constraint have to be the image of some assignment under  $\varphi_{\text{Asn}}$ . This is the case here, as the intersection with  $\varphi(\text{hull}(d))$  can only contain such assignments. So in summary, we get

$$\varphi^- \circ p \circ \varphi(d) \subseteq \text{hull}(\text{dom}(\varphi^-(c_p) \cap \text{hull}(d)))$$

which is the definition of  $\widehat{\varphi}(p)$  being  $\text{bounds}(\mathbb{Z})$ -complete.  $\blacksquare$



**Stronger notions of completeness.** Similar theorems hold for domain completeness, range and bounds( $\mathbb{Z}$ ) completeness. The theorems directly follow from the fact that any view  $\varphi$  is *domain-injective*, meaning that  $\varphi^-(\text{dom}(c)) = \text{dom}(\varphi^-(c))$  for all constraints  $c$ . We split this statement into the following two lemmas.

**Lemma 2** Given a constraint  $c$ , let  $d = \text{dom}(c)$ . Then for all  $x \in X$ , we have  $v \in d(x) \Leftrightarrow \exists a \in c : a(x) = v$ . \*

*Proof.* We prove both directions of the equivalence:

- $\Rightarrow$  There must be such an assignment  $a$  because otherwise one can construct a strictly stronger  $d' \subset d$  with  $v \notin d'(x)$  such that still  $c \subseteq d'$ .
- $\Leftarrow$  Each domain  $d'$  in the intersection  $\bigcap \{d' \in \text{Dom} \mid c \subseteq \text{con}(d')\}$  must contain the value  $v \in d'(x)$  as  $c \subseteq d'$ . So for the result of the intersection  $d$ ,  $v \in d(x)$ . ■

**Lemma 3** Any view  $\varphi$  is domain-injective. \*

*Proof.* We have to show that  $\varphi^-(\text{dom}(c)) = \text{dom}(\varphi^-(c))$  holds for any constraint  $c$  and any view  $\varphi$ . For clarity, we write the equation including the implicit  $\text{con}(\cdot)$  operations:  $\varphi^-(\text{con}(\text{dom}(c))) = \text{con}(\text{dom}(\varphi^-(c)))$ . By definition of  $\varphi^-$  and  $\text{con}(\cdot)$ , we have

$$\begin{aligned} \varphi^-(\text{con}(\text{dom}(c))) &= \{a \in \text{Asn} \mid \forall x \in X : \varphi_{\text{Asn}}(a)(x) \in \text{dom}(c)(x)\} \\ &= \{a \in \text{Asn} \mid \forall x \in X \exists b \in c : \varphi_{\text{Asn}}(a)(x) = b(x)\} \quad (\text{Lemma 2}) \end{aligned}$$

As  $\varphi_{\text{Asn}}$  is an injective function, we can find such a  $b$  that is in the range of  $\varphi_{\text{Asn}}$ , if and only if there is also a  $b' \in \varphi^-(c)$  such that  $\varphi_{\text{Asn}}(b') = b$ . Therefore, we get

$$\begin{aligned} &\{a \in \text{Asn} \mid \forall x \in X \exists b' \in \varphi^-(c) : a(x) = b'(x)\} \\ &= \{a \in \text{Asn} \mid \forall x \in X : a(x) \in \text{dom}(\varphi^-(c))(x)\} \\ &= \text{con}(\text{dom}(\varphi^-(c))) \quad \blacksquare \end{aligned}$$

The following three theorems express under which conditions the different notions of completeness are preserved when deriving propagators. The proofs for these theorems are analogous to the proof of [Theorem 1](#), using [Lemma 3](#).

**Theorem 2 (Bounds( $D$ ) completeness)** Let  $p$  be a bounds( $D$ )-complete propagator. For any hull-injective view  $\varphi$ , the propagator  $\widehat{\varphi}(p)$  is bounds( $D$ )-complete. \*

**Theorem 3 (Range completeness)** Let  $p$  be a range-complete propagator. For any hull-surjective view  $\varphi$ , the propagator  $\widehat{\varphi}(p)$  is range-complete. \*

**Theorem 4 (Domain completeness)** Let  $p$  be a domain-complete propagator, and let  $\varphi$  be a view. Then  $\widehat{\varphi}(p)$  is domain-complete. \*

This last theorem is particularly interesting. It states that deriving a propagator from a domain-complete propagator yields a domain-complete propagator, no matter what view is used.

An additional case not discussed above is a propagator derived from a bounds( $\mathbb{Z}$ )-complete propagator and a hull-injective but not hull-surjective view. This is a common case, as for example scale views and propagators for linear constraints fall in this category. The resulting propagator is, in general, only bounds( $\mathbb{R}$ )-complete. This is exactly what we would expect from a propagator for linear equations, as the next example demonstrates.

**Example 4 (Linear constraints)** A propagator for a linear constraint  $c_\Sigma = \llbracket \sum_{i=1}^n x_i = c \rrbracket$  and  $n$  scale views (see Example 3) yield a propagator for a linear constraint with coefficients  $c_{\Sigma a} = \llbracket \sum_{i=1}^n a_i x_i = c \rrbracket$ .

The usual propagator for a linear constraint with coefficients achieves bounds( $\mathbb{R}$ ) consistency in linear time  $O(n)$  [21]. However, it is bounds( $\mathbb{Z}$ )-complete for unit coefficients. Hence, the above-mentioned property applies: The propagator for  $c_\Sigma$  is bounds( $\mathbb{Z}$ )-complete, scale views are only hull-injective, so the derived propagator for  $c_{\Sigma a}$  is bounds( $\mathbb{R}$ )-complete. Implementing the simpler propagator without coefficients and deriving the variant with coefficients yields propagators with exactly the same runtime complexity and propagation strength as manually implemented propagators. \*

### 3.4 Additional Properties of Derived Propagators

The following discussion shows how views can be composed, and how derived propagators behave with respect to idempotency and subsumption.

**View composition.** A derived propagator permits further derivation. Consider a propagator  $p$  and two views  $\varphi, \varphi'$ . Then  $\widehat{\varphi}'(\widehat{\varphi}(p))$  is a perfectly acceptable derived propagator, and properties like correctness and completeness carry over transitively. For instance, we can derive a propagator for  $\llbracket x - y = c \rrbracket$  from a propagator for  $\llbracket x + y = 0 \rrbracket$ , combining an *offset view* ( $\varphi_y(v) = v + c$ ) and a *minus view* ( $\varphi'_y(v) = -v$ ) on  $y$ . This yields a propagator for  $\llbracket x + (-(y + c)) = 0 \rrbracket = \llbracket x - y = c \rrbracket$ .

**Fixed points.** Schulte and Stuckey [40] show how to optimize the scheduling of propagators that are known to be at a fixed point. Views preserve fixed points of propagators, so the same optimizations apply to derived propagators.

**Proposition 4** Let  $p$  be a propagator, let  $\varphi$  be a view, and let  $d$  be a domain. If  $\varphi(d)$  is a fixed point of  $p$ , then  $d$  is a fixed point of  $\widehat{\varphi}(p)$ . \*

*Proof.* Assume  $p(p(\varphi(d))) = p(\varphi(d))$ . We have to show  $\widehat{\varphi}(p)(d) = \widehat{\varphi}(p)(\widehat{\varphi}(p)(d))$ . With the assumption, we can write  $\widehat{\varphi}(p)(d) = (\varphi^- \circ p \circ \varphi)(d)$ . We know that  $\varphi \circ \varphi^-(c) = c$  if  $|\varphi^-(c)| = |c|$ . As we first apply  $\varphi$ , this is the case here, so we can add  $\varphi \circ \varphi^-$  in the middle, yielding  $(\varphi^- \circ p \circ (\varphi \circ \varphi^-) \circ p \circ \varphi)(d)$ . With function composition being associative, this is equal to  $\widehat{\varphi}(p)(\widehat{\varphi}(p)(d))$ . ■

**Subsumption.** A propagator is *subsumed* (also known as entailed) by a domain  $d$  if and only if for all stronger domains  $d' \subseteq d$ ,  $p(d') = d'$ . Subsumed propagators cannot do any pruning in the remaining subtree of the search, and can therefore be removed. Deciding subsumption is coNP-complete in general, but for many practically relevant propagators an approximation can be decided easily (such as when a domain becomes assigned). The following theorem states that the approximation is also valid for the derived propagator.

**Proposition 5** Let  $p$  be a propagator and let  $\varphi$  be a view. The propagator  $\widehat{\varphi}(p)$  is subsumed by a domain  $d$  if and only if  $p$  is subsumed by  $\varphi(d)$ . \*

*Proof.* With P2 we get that  $\forall d' \subseteq d : \varphi^-(p(\varphi(d'))) = d'$  is equivalent to

$$\forall d' \subseteq d : \varphi^-(p(\varphi(d'))) = \varphi^-(\varphi(d'))$$

As  $\varphi^-$  is a function, and because it preserves contraction (see [Proposition 3](#)), this is equivalent to  $\forall d' \subseteq d : p(\varphi(d')) = \varphi(d')$ . This can be rewritten to  $\forall d'' \subseteq \varphi(d) : p(d'') = d''$  because all  $\varphi(d')$  are subsets of  $\varphi(d)$ . ■

## 4 Propagator Derivation Techniques

This section introduces systematic techniques for deriving propagators using views. The techniques capture the transformation, generalization, specialization, and type conversion of propagators and are shown to be widely applicable across different variable domains and application areas.

### 4.1 Transformation

**Boolean connectives.** For Boolean variables, where  $V = \{0, 1\}$ , the only view apart from identity for Boolean variables captures negation. A *negation view* on  $x$  defines  $\varphi_x(v) = 1 - v$  for  $x \in X$  and  $v \in V$ . Deriving propagators using Boolean views thus means to propagate using *literals* rather than variables.

The obvious application of negation views is to derive propagators for all Boolean connectives from just three propagators. A negation view for  $x$  in  $x = y$  yields a propagator for  $\neg x = y$ . From disjunction  $x \vee y = z$  one can derive conjunction  $x \wedge y = z$  with negation views on  $x, y, z$ , and implication  $x \rightarrow y = z$  with a negation view on  $x$ . From equivalence  $x \leftrightarrow y = z$  one can derive exclusive or  $x \oplus y = z$  with a negation view on  $z$ .

As Boolean constraints are widespread, it pays off to optimize frequently occurring cases of propagators for Boolean connectives. One important propagator is for  $\bigvee_{i=1}^n x_i = y$  with arbitrarily many variables. Again, conjunction can be derived with negation views on the  $x_i$  and on  $y$ . Another important propagator implements the constraint  $\bigvee_{i=1}^n x_i = 1$ . A dedicated propagator for this constraint is essential as the constraint occurs frequently and can be implemented efficiently using watched literals, see for example [17]. With views all implementation work is readily reused for conjunction. This shows a general advantage of views: effort put into optimizing a single propagator directly pays off for all other propagators derived from it.

**Boolean cardinality.** Like the constraint  $\bigvee_{i=1}^n x_i = 1$ , the Boolean cardinality constraint  $\sum_{i=1}^n x_i \geq c$  occurs frequently and can be implemented efficiently using watched literals (requiring  $c + 1$  watched literals, Boolean disjunction corresponds to the case where  $c = 1$ ). But also a propagator for  $\sum_{i=1}^n x_i \leq c$  can be derived using negation views on the  $x_i$  with the following transformation:

$$\begin{aligned} \sum_{i=1}^n x_i \leq c &\iff -\sum_{i=1}^n x_i \geq -c &\iff n - \sum_{i=1}^n x_i \geq n - c \\ &\iff \sum_{i=1}^n 1 - x_i \geq n - c &\iff \sum_{i=1}^n \neg x_i \geq n - c \end{aligned}$$

**Reification.** Many reified constraints (such as  $(\sum_{x=1}^n x_i = c) \leftrightarrow b$ ) also exist in a negated version (such as  $(\sum_{x=1}^n x_i \neq c) \leftrightarrow b$ ). Deriving the negated version is trivial by using a negation view on the Boolean control variable  $b$ . This contrasts nicely with the effort without views: either the entire code must be duplicated or the parts that perform checking whether the constraint or its negation is subsumed must be factored out and combined differently for the two variants.




Fig. 1 Task  $t$  and its dual task  $t'$  using a minus view

**Transformation using set views.** Set constraints deal with variables whose values are finite sets. Using *complement views* (analogous to Boolean negation, as sets with their usual operations also form a Boolean algebra) on  $x, y, z$  with a propagator for  $x \cap y = z$  yields a propagator for  $x \cup y = z$ . A complement view on  $y$  yields  $x \setminus y = z$ .

**Transformation using integer views.** The obvious integer equivalent to negation views for Boolean variables are *minus views*: a minus view on an integer variable  $x$  is defined as  $\varphi_x(v) = -v$ . Minus views help to derive propagators following simple transformations: for example,  $\min(x, y) = z$  can be derived from  $\max(x, y) = z$  by using minus views for  $x, y$ , and  $z$ .

Transformations through minus views can improve performance in subtle ways. Consider a  $\text{bounds}(\mathbb{Z})$ -complete propagator for multiplication  $x \times y = z$  (for example, [1, Section 6.5] or [39]). Propagation depends on whether zero is still included in the domains of  $x, y$ , or  $z$ . Testing for inclusion of zero each time the propagator is executed may be inefficient, and it leads to a convoluted implementation. Instead, one would like to rewrite the propagator to special variants where  $x, y$ , and  $z$  are either strictly positive or negative. These variants can propagate more efficiently, in particular because propagation can easily be made idempotent. Instead of implementing three different propagators ( $x, y, z$  strictly positive; only  $x$  or  $y$  strictly positive; only  $z$  strictly positive), a single propagator assuming that all views are strictly positive is sufficient. The other propagators can be derived using minus views.

Again, with views it becomes realistic to optimize a single implementation of a propagator and derive other, equally optimized, implementations. The effort to implement all required specialized versions without views is typically unrealistic.

**Scheduling propagators.** An important application area is constraint-based scheduling, see for example [2]. Many propagation algorithms for constraint-based scheduling are based on tasks, where a task  $t$  is characterized by its start time, processing time (how long does the task take to be executed on a resource), and end time. Scheduling algorithms are typically expressed in terms of earliest start time ( $\text{est}(t)$ ), latest start time ( $\text{lst}(t)$ ), earliest completion time ( $\text{ect}(t)$ ), and latest completion time ( $\text{lct}(t)$ ).

Another particular aspect of scheduling algorithms is that they are often required in two, mutually dual, variants. Let us consider not-first/not-last propagation as an example. Assume a set of tasks  $T$  and a task  $t \notin T$  to be scheduled on the same resource. Then  $t$  cannot be scheduled before the tasks in  $T$  ( $t$  is not-first in  $T \cup \{t\}$ ), if  $\text{ect}(t) > \text{lst}(T)$  (where  $\text{lst}(T)$  is a conservative estimate of the latest start time of all tasks in  $T$ ). Hence,  $\text{est}(t)$  can be adjusted to leave some room for at least one task from  $T$ . The dual variant is that  $t$  is not-last: if  $\text{ect}(T) > \text{lst}(t)$  (again,  $\text{ect}(T)$  estimates the earliest completion time of  $T$ ), then  $\text{lct}(t)$  can be adjusted.

Running the dual variant of a scheduling algorithm on tasks  $t \in T$  is the same as running the original algorithm on the *dual tasks*  $t' \in T'$ , which are simply mirrored at the 0-origin of

the time scale (see [Figure 1](#)):

$$\text{est}(t') = -\text{lct}(t) \quad \text{ect}(t') = -\text{lst}(t) \quad \text{lst}(t') = -\text{ect}(t) \quad \text{lct}(t') = -\text{est}(t)$$

The dual variant of a scheduling propagator can be automatically derived using a minus view that transforms the time values. In our example, only a propagator for not-first needs to be implemented and the propagator for not-last can be derived (or vice versa). This is in particular beneficial if the algorithms use sophisticated data structures such as  $\Omega$ -trees [44]. Here, also the data structure needs to be implemented only once and the dual data structure for the dual propagator is derived.

## 4.2 Generalization

Common views for integer variables capture linear transformations of the integer values: an *offset view* for  $o \in \mathbb{Z}$  on  $x$  is defined as  $\varphi_x(v) = v + o$ , and a *scale view* for  $a \in \mathbb{Z}$  on  $x$  is defined as  $\varphi_x(v) = av$ .

Offset and scale views are useful for generalizing propagators. Generalization has two key advantages: simplicity and efficiency. A more specialized propagator is often simpler to implement (and simpler to implement *correctly*) than a generalized version. The specialized version can save memory and runtime during execution.

We can devise an efficient propagation algorithm for the common case of a linear equality constraint with unit coefficients  $\sum_{i=1}^n x_i = c$ . The more general case  $\sum_{i=1}^n a_i x_i = c$  can be derived by using scale views for  $a_i$  on  $x_i$  (the same technique of course applies to linear inequalities and disequality rather than equality). A common optimization for the general case is to determine the greatest common divisor of the coefficients  $a_i$  and simplify the equation accordingly. Using views does not preclude this optimization, as the analysis and simplification can be performed on the model before posting the derived propagator.

Similarly, a propagator for *all-different*( $x_1, \dots, x_n$ ) can be generalized to *all-different*( $c_1 + x_1, \dots, c_n + x_n$ ) by using offset views for  $c_i \in \mathbb{Z}$  on  $x_i$ . Likewise, from a propagator for the element constraint  $a[x] = y$  for integers  $a_1, \dots, a_n$  and integer variables  $x$  and  $y$ , we can derive the generalized version  $a[x + o] = y$  with an offset view, where  $o \in \mathbb{Z}$  provides a useful offset for the index variable  $x$ .

These generalizations can be applied to domain- as well as bounds-complete propagators. While most Boolean propagators are domain-complete, bounds completeness plays an important role for integer propagators. [Section 3.3](#) shows that, given appropriate hull-surjective and/or hull-injective views, the different notions of bounds consistency are preserved when deriving propagators.

The views for integer variables presented in this section have the following properties: minus and offset views are hull-bijective, whereas a scale view for  $a \in \mathbb{Z}$  on  $x$  is always hull-injective and only hull-surjective if  $a = 1$  or  $a = -1$  (in which cases it coincides with the identity view or a minus view, respectively).

## 4.3 Specialization

We employ *constant views* to specialize propagators. A constant view behaves like an assigned variable. In practice, specialization has two advantages. Fewer variables require less memory. And specialized propagators can be compiled to more efficient code, if the constants are known at compile time.

Examples for specialization are

- a propagator for binary linear inequality  $x + y \leq c$  derived from a propagator for  $x + y + z \leq c$  by using a constant 0 for  $z$ ;
- a reified propagator for  $(x = c) \leftrightarrow b$  from  $(x = y) \leftrightarrow b$  and a constant  $c$  for  $y$ ;
- propagators for the counting constraints  $|\{i \mid x_i = c\}| = z$  and  $|\{i \mid x_i = y\}| = c$  from a propagator for  $|\{i \mid x_i = y\}| = z$ ;
- a propagator for set disjointness from a propagator for  $x \cap y = z$  and a constant empty set for  $z$ ; and many more.

We have to extend the model in a straightforward way to accommodate constant views. Propagators may now be defined with respect to a superset of the variables,  $X' \supseteq X$ . A constant view for the value  $k$  on a variable  $z \in X' \setminus X$  translates between the two sets of variables:

$$\varphi(c) = \{a[k/z] \mid a \in c\} \quad \varphi^-(c) = \{a|_X \mid a \in c\}$$

Here,  $a[k/z]$  means augmenting the assignment  $a$  so that it maps  $z$  to  $k$ , and  $a|_X$  is the functional restriction of  $a$  to the set  $X$ .

It is important that this definition preserves failure. If a propagator returns a failed domain  $d$  that maps  $z$  to the empty set, then  $\varphi^-(d)$  is the empty set, too (recall that this is really  $\varphi^-(\text{con}(d))$ , and  $\text{con}(d) = \emptyset$  if  $d(z) = \emptyset$ ).

#### 4.4 Type Conversion

A type conversion view lets propagators for one type of variable work with a different type, by translating the underlying representation. Our model already accommodates for this, as a view  $\varphi_x$  maps elements between different sets  $V$  and  $V'$ .

**Integer views.** Boolean variables are essentially integer variables restricted to the values  $\{0, 1\}$ . Constraint programming systems may choose a more efficient implementation for Boolean variables and hence the types for integer and Boolean variables differ. By wrapping an efficient Boolean variable in an *integer view*, all integer propagators can be directly reused with Boolean variables. Note that in an implementation based on C++ templates as presented in [Section 6](#), a wrapper may not even be necessary if Boolean variables simply implement the same interface as integer variables. Integer views for Boolean variables can save substantial effort. For example, an implementation of the *regular*-constraint for Boolean variables can be derived which is actually useful in practice [26].

Care should be taken, however, as dedicated propagators for the Boolean case can often exploit special techniques. For example, a Boolean disjunction propagated using watched literals is far more efficient than instantiating a linear constraint  $\sum_{i=1}^n b_i \geq 1$ ; and the constraint  $\sum_{i=1}^n a_i b_i \leq c$  for constant coefficients  $a_i$  can be implemented efficiently by sorting the terms according to the  $a_i$ .

**Singleton set views.** A *singleton set view* on an integer variable  $x$ , defined as  $\varphi_x(v) = \{v\}$ , presents an integer variable as a set variable. Many constraints involve both integer and set variables, and some of them can be expressed with singleton set views. A simple constraint is  $x \in y$ , where  $x$  is an integer variable and  $y$  a set variable. Singleton set views derive it as  $\{x\} \subseteq y$ . This extends to  $\{x\} \diamond y$  for all other set relations  $\diamond$ .

Singleton set views can also be used to derive pure integer constraints from set propagators. For example, the constraint  $\text{same}(x_1, \dots, x_n, y_1, \dots, y_m)$  with integer variables  $x_i, y_i$

<i>Variable type</i>	<i>Implemented</i>	<i>Derived</i>	<i>Ratio</i>
Integer	93	377	4.05
Boolean	30	93	3.10
Set	31	146	4.71
<i>Overall</i>	154	616	4.00

**Table 1** Number of implemented vs. derived propagators

states that the variables  $x_i$  take the same values as the variables  $y_i$ . With singleton set views,  $\bigcup_{i=1}^n \{x_i\} = \bigcup_{j=1}^m \{y_j\}$  implements this constraint (albeit with weaker propagation than the algorithm presented in [4]).

**Set bounds and complete set domain variables.** Most systems approximate set variable domains as set intervals defined by lower and upper bounds [32, 18]. However, [22] introduces a representation for the complete domains of set variables, using ROBDDs. Type conversion views can translate between set interval and ROBDD-based implementations. We can derive a propagator on ROBDD-based variables from a set interval propagator, and thus reuse set interval propagators for which no efficient ROBDD representation exists.

#### 4.5 Applicability and Return on Investment

To get an overview of how applicable the presented techniques for propagator derivation are, let us consider the use of views in Gecode (version 3.7.2). Table 1 shows the number of propagator implementations and the number of propagators derived from the implementations. On average, every propagator implementation results in four derived propagators. Propagator implementations in Gecode account for almost 40 000 lines of code and 21 000 lines of documentation. As a rough estimate, deriving propagators using views thus saves around 120 000 lines of code and 60 000 lines documentation to be written, tested, and maintained. On the other hand, the views mentioned in this section are implemented in less than 8 000 lines of code, yielding a 1500% return on investment.

## 5 Related Work

The techniques presented in this paper are based on previous research and folklore ideas that have been known in the constraint programming community for a long time. Reusing functionality (like a propagator) by wrapping it in an adaptor (like a view) is of course a much more general technique – think of higher-order functions like *fold* or *map* in functional programming languages; or chaining command-line tools in Unix operating systems using pipes.

This section discusses related approaches in the constraint programming context, as well as systems that have used or extended the concept of views as presented here.

### 5.1 Indexicals and IBM ILOG Expressions

Views are closely related to both indexicals [5, 43] and IBM ILOG CP Optimizer expressions [24]. Indexicals are read-only or *unidirectional* but more expressive than views, Solver

expressions are *bidirectional* but give no guarantees about the result of update operations. Views are bidirectional but limited in expressivity, as discussed below.

An indexical is a propagator that prunes a single variable and is defined in terms of range expressions. In contrast to views, range expressions can involve multiple variables, but on the other hand only operate in one direction. For instance, in an indexical for the constraint  $\llbracket x = y + z \rrbracket$ , the range expression  $y + z$  would be used to prune the domain of  $x$ , but not for pruning the domains of  $y$  or  $z$ .

The `ILCIntExp` base class in IBM ILOG CP Optimizer (which has been around since early versions of ILOG Solver) is an extension of the indexical concept, where also domain update operations are supported. Expressions can therefore be used to implement derived propagators. However, no guarantees are given as to whether an update operation on a complex expression actually results in domain updates of the variables the expression consists of. The CP Optimizer documentation explicitly recommends decomposing expressions using intermediate variables in order to get the necessary guarantees. Views as presented in this paper follow an opposite approach. They extend indexicals with update operations but are *limited* in expressivity, just enough to get the strong correctness and completeness guarantees of derived propagators presented in [Section 3](#).

## 5.2 SAT Literals

Unit propagation in SAT solvers performs propagation for Boolean clauses, which are disjunctions of *literals*, which in turn are positive or negated Boolean variables. In implementations such as MiniSat [12], the Boolean clause propagator is in fact derived from a simple  $n$ -ary disjunction propagator and *literal views* of the variables that perform negation for the negative literals.

## 5.3 Constraint Composition

Instead of regarding a view  $\varphi$  as *transforming* a constraint  $c$ , one can regard  $\varphi$  as *additional* constraints, implementing the decomposition. Assuming  $\text{vars}(c) = x_1, \dots, x_n$ , we use additional variables  $x'_1, \dots, x'_n$ . Instead of  $c$ , we use  $c' = c[x_1/x'_1, \dots, x_n/x'_n]$ , which is the same relation as  $c$ , but on  $x'_1, \dots, x'_n$ . Finally,  $n$  *view constraints*  $c_{\varphi,i}$  link the original variables to the new variables, each  $c_{\varphi,i}$  being equivalent to the relation  $x'_i = \varphi_i(x_i)$ . The solutions of the decomposition model, restricted to the  $x_1, \dots, x_n$ , are exactly the solutions of the original view-based model.

Every view constraint  $c_{\varphi,i}$  shares exactly one variable with  $c$  and no variable with any other  $c_{\varphi,i}$ . Thus, the constraint graph is Berge-acyclic [3], and a fixed point can be computed by first propagating all the  $c_{\varphi,i}$ , then propagating  $c[x_1/x'_1, \dots, x_n/x'_n]$ , and then again propagating the  $c_{\varphi,i}$ . This is exactly what  $\varphi^- \circ p \circ \varphi$  does. Constraint solvers typically do not provide any means of specifying the propagator scheduling in such a fine-grained way (Lagerkvist and Schulte show how to use propagator groups to achieve this [27]). Thus, deriving propagators using views is also a technique for specifying perfect propagator scheduling.

On a more historical level, a derived propagator is related to the notion of *path consistency*. A domain is path-consistent for a set of constraints, if for any subset  $\{x, y, z\}$  of its variables,  $v_1 \in d(x)$  and  $v_2 \in d(y)$  implies that there is a value  $v_3 \in d(z)$  such that the pair  $(v_1, v_2)$  satisfies all the (binary) constraints between  $x$  and  $y$ , the pair  $(v_1, v_3)$  satisfies



all the (binary) constraints between  $x$  and  $z$ , and the pair  $(v_3, v_2)$  satisfies all the (binary) constraints between  $z$  and  $y$  [28]. If  $\widehat{\varphi}(p)$  is domain-complete for  $\varphi^-(c)$ , then it achieves path consistency for the constraint  $c[x_1/x'_1, \dots, x_n/x'_n]$  and all the  $c_{\varphi,i}$  in the decomposition model.

#### 5.4 Views in Other Systems

Several systems have adopted techniques for deriving propagators using views as presented in this paper. The following three explicitly reference our previous work as a basis for their implementations.

**MINION** [16] uses views much in the same way as Gecode, to derive propagators from generic implementations.

**Lazy Clause Generation** [13] is a combination of propagation-based finite-domain constraint solving and SAT-style clause learning. Views have proven essential for an efficient implementation, as they greatly reduce the number of Boolean literals generated, which saves memory, makes learned clauses more effective, and helps the VSIDS search heuristic.

**CaSPER** [9, 8] extends views as introduced here to include non-injective multi-variable views. This makes it possible to map entire arithmetic and Boolean expressions in the model to expressions of views and simple propagators. Non-injective views induce weaker guarantees for the derived propagators (as will be discussed in [Section 7](#)). Derived propagators in CaSPER are therefore restricted to bounds propagation.

## 6 Implementation

This section presents an implementation architecture for views and derived propagators, based on making propagators *parametric*. Deriving a propagator then means *instantiating* a parametric propagator with views. The presented architecture is an orthogonal layer of abstraction on top of any solver implementation.

### 6.1 Views

The model introduced views as functions that transform the input and output of a propagator, which maps domains to domains. In an object-oriented implementation of this model, a propagator is no longer a function, but an object with a `propagate` method that *accesses* and *modifies* a domain through the methods of variable objects. Such an object-oriented model is used for example by ILOG Solver [33] and Choco [25], and is the basis of most of the current propagation-based constraint solvers.

In the following examples, we use C++ templates to achieve propagator parametricity. [Section 6.3](#) will discuss an alternative implementation based on dynamic binding.

[Figure 2](#) shows C++ classes for a simple integer variable (just representing bounds information) and corresponding minus and constant integer views. The views have the same interface as the variable, so that they can be used in its place. A minus view contains a pointer to the underlying integer variable and *delegates* all the operations, performing the necessary transformations. A constant view simply returns the constant for domain access operations, and ignores update operations. For the sake of brevity, the sample code does not handle failure (when domains are emptied).

```

class IntVar {
private: int _min, _max;
public: int min(void) { return _min; }
       int max(void) { return _max; }
       void adjmin(int n) { if (n > _min) _min = n; }
       void adjmax(int n) { if (n < _max) _max = n; }
       void subscribe(EventSet e) { /* perform subscription */ }
};

class MinusView {
protected: IntVar* x;
public: MinusView(IntVar* x0) : x(x0) {}
       int min(void) { return -x->max(); }
       int max(void) { return -x->min(); }
       void adjmin(int n) { x->adjmax(-n); }
       void adjmax(int n) { x->adjmin(-n); }
       void subscribe(EventSet e) { x->subscribe(negate(e)); }
};

class ConstIntView {
protected: int c;
public: ConstIntView(int c0) : c(c0) {}
       int min(void) { return c; }
       int max(void) { return c; }
       void adjmin(int n) {}
       void adjmax(int n) {}
       void subscribe(Event e) {}
};

```

Fig. 2 Integer variable, minus and constant views

**Events.** Most constraint solvers schedule the execution of propagators according to *events* (for an overview see [38]). For example, a propagator  $p$  for  $\llbracket x < y \rrbracket$  can only prune the domain (and thus should only be executed) if either the lower bound of  $x$  or the upper bound of  $y$  changes, written  $\text{lbc}(x)$  and  $\text{ubc}(y)$ . We say that  $p$  *subscribes* to the *event set*  $\{\text{lbc}(x), \text{ubc}(y)\}$ , implemented as the `subscribe` method in Figure 2.

Now assume that  $p'$  is derived from  $p$  using minus views on  $x$  and  $y$ , thus implementing  $x > y$ . Obviously,  $p'$  should subscribe to the dual event set,  $\{\text{ubc}(x), \text{lbc}(y)\}$ . In the implementation, minus views first negate the event set before delegating the subscription to the underlying variable.

For some views, events like `ubc` or `lbc` may not map to corresponding bounds events on the view. For instance, when using a permutation view that maps  $i$  to  $\pi(i)$  for an arbitrary permutation  $\pi$ , an upper bound event on the variable may not correspond to a bound change on the view and vice versa. In this case, the `subscribe` method of the view must map any event to the `dmc` event, signaling an arbitrary domain change. In a system that permits views like this, propagators must be able to deal with spurious events, as they may be scheduled even if the event they are actually interested in (such as `ubc`) may not have happened.

## 6.2 Deriving Propagators

In order to derive a propagator using view objects like the above, we use *parametricity*, a mechanism provided by the implementation language that supports the instantiation of the same code (the propagator) with different parameters (the views).

```

template<class VX, class VY>
class LessThan : public Propagator {
protected: VX* x; VY* y;
public: LessThan(VX* x0, VY* y0) : x(x0), y(y0) {
    x->subscribe(LOWER_BOUND); y->subscribe(UPPER_BOUND);
}
    virtual void propagate(void) {
        x->adjmax(y->max()-1); y->adjmin(x->min()+1);
    }
};

```

**Fig. 3** Parametric less-than propagator

Figure 3 shows a simple less-than propagator. The propagator is based on C++ templates, it is *parametric* over the types of the two views it uses and can be *instantiated* with any view that provides the necessary operations. This type of parametricity is called *parametric polymorphism*, and is available in other programming languages for example in the form of Java generics [20] or Standard ML functors [29].

Given two pointers to integer variables  $x$  and  $y$ , the propagator can be instantiated to implement  $\llbracket x < y \rrbracket$  as follows (using the `IntVar` class from Figure 2):

```
new LessThan<IntVar, IntVar>(x, y);
```

The following instantiation yields a propagator for  $\llbracket x > y \rrbracket$ :

```
new LessThan<MinusView, MinusView>(new MinusView(x), new MinusView(y));
```

### 6.3 Parametricity

Independent of the concrete implementation, views form an orthogonal layer of abstraction on top of any propagation-based constraint solver. As long as the implementation language provides some kind of parametricity, and variable domains are accessed through some form of variable objects, propagators can be derived using views.

In addition to parametric polymorphism, two other forms of parametricity exist, *functional* parametricity and *dynamic binding*. Functional parametricity means that in languages such as Standard ML [29] or Haskell [31], a higher-order function is parametric over its arguments. Dynamic binding is typically coupled with inheritance in object-oriented languages (virtual function calls in C++, method calls in Java). Even in languages that lack direct support for parametricity, parametric behavior can often be achieved using other mechanisms, such as macros or function pointers in C.

**Choice of parametricity.** In C++, parametric polymorphism and dynamic binding have advantages and disadvantages as it comes to deriving propagators.

Templates are compiled by *monomorphization*: the code is replicated and specialized for each instance. The compiler can generate optimized code for each instance, for example by inlining the transformations that a view implements.

Achieving high efficiency in C++ with templates sacrifices some expressiveness. Instantiation can *only* happen at compile-time. Hence, either C++ must be used for modeling, or all potentially required propagator variants must be instantiated explicitly. The *choice* which propagator to use can however be made at runtime: for linear equations, for instance, if all coefficients are units, the optimized original propagator can be posted.

For  $n$ -ary constraints, compile-time instantiation can be a limitation, as all arrays must be monomorphic (contain only a single kind of view). For example, one cannot mix scale and minus views in linear constraints. For some propagators, we can work around this restriction using more than a single array of views. For example, a propagator for a linear constraint can employ two arrays of different view types, one of which may then be instantiated with identity views and the other with minus views. While this poses a limitation in principle, our experience from Gecode suggests that there are only few propagators in practice that suffer from this limitation.

Dynamic binding is more flexible than parametric polymorphism, as instantiation happens at runtime and arrays can be polymorphic. Views with virtual methods can therefore be used at the modeling level, and new views can be introduced and used with existing propagators without recompilation. Several systems that implement views or related approaches (for instance IBM ILOG CP Optimizer [24] or CaSPER [8]) make use of dynamic binding in order to take advantage of this added flexibility.

The flexibility of dynamic binding comes at the cost of reduced efficiency, for two reasons. First, dynamic binding is usually implemented using *virtual tables*, where methods are invoked indirectly, requiring a table lookup at runtime. Second, and more importantly, the indirection through a virtual table makes certain compiler optimizations impossible. For instance, the transformations done by view operations typically cannot be inlined and optimized. For a detailed discussion of the performance impact of dynamic binding, see [10] and [23]. Section 8 evaluates empirically how these virtual method calls affect performance.

A system with views based on C++ templates does not rule out dynamic binding. Propagators can be instantiated with all the views known to be useful at compile time, and, additionally, as a version with views based on virtual methods.

**Compile-time versus runtime constants.** Some views involve a parameter, such as the coefficient of a scale view or the constant of a constant view. Depending on how these views are used, the parameters may be known at compile time, or they may be part of the instance data of the constraint model.

For example, when deriving Boolean disjunction  $x \vee y$  from  $x \vee y \leftrightarrow z$  and a constant 1 for  $z$ , we can take advantage of the constant being known at compile-time, and the compiler can apply more aggressive optimizations. Other examples are a minus view that can be regarded as a compile-time specialization of a scale view with coefficient  $-1$ , and a zero view which specializes a constant integer view.

Conversely, using runtime parameters the programmer regains some flexibility; for example, an *affine view*  $\varphi_x(v) = av + o$  with integer parameters  $a$  and  $o$  can express any combination of offset, scale, minus, and constant views, and therefore (monomorphic) arrays of affine views express a limited form of array polymorphism. These types of views will be used for propagators where the parameters are part of the instance data of the model (as will often be the case with array coefficients, for instance).

## 6.4 Iterators

Typical domain operations involve a single integer value, for instance adjusting the minimum or maximum of an integer variable. These operations are not efficient if a propagator performs full domain reasoning on integer variables or deals with set variables. Therefore, set-valued operations, like updating a whole integer variable domain to a new set, or excluding a set of elements from a set variable domain, are important for efficiency. Many

constraint programming systems provide an abstract set-datatype for accessing and updating variable domains, as for example in Choco [6], ECL<sup>i</sup>PS<sup>e</sup> [11], SICStus Prolog [42], and Mozart [30]. IBM ILOG CP Optimizer [24] only allows access by iterating over the values of a variable domain.

This section develops *iterators* as one particular abstract datatype for set-valued operations on variables and views. There are two main reasons to discuss iterators in this paper. First, iterators provide simple, expressive, and efficient set-valued operations on variables. Second, and more importantly, iterators transparently perform the transformations needed for set-valued operations on views, and thus constitute a perfect fit for deriving propagators.

**Range sequences and range iterators.** We define a *range*  $[m .. n]$  to denote the set of integers  $\{l \in \mathbb{Z} \mid m \leq l \leq n\}$ . A *range sequence*  $\text{ranges}(S)$  for a finite set of integers  $S \subseteq \mathbb{Z}$  is the shortest sequence  $s = \langle [m_1 .. n_1], \dots, [m_k .. n_k] \rangle$  such that  $S = \bigcup_{i=1}^k [m_i .. n_i]$  and the ranges are ordered by their smallest elements ( $m_i \leq m_{i+1}$  for  $1 \leq i < k$ ). We thus define the set covered by a range sequence as  $\text{set}(s) = \bigcup_{i=1}^k [m_i .. n_i]$ . The above range sequence is also written as  $\langle [m_i .. n_i]_{i=1}^k \rangle$ . Clearly, the range sequence of a set is unique, none of its ranges is empty, and  $n_i + 1 < m_{i+1}$  for  $1 \leq i < k$ .

A *range iterator* for a range sequence  $s = \langle [m_i .. n_i]_{i=1}^k \rangle$  is an object that provides iteration over  $s$ : each of the  $[m_i .. n_i]$  can be obtained in sequential order but only one at a time. A range iterator  $r$  provides the following operations:  $r.\text{done}()$  tests whether all ranges have been iterated,  $r.\text{next}()$  moves to the next range, and  $r.\text{min}()$  and  $r.\text{max}()$  return the minimum and maximum value for the current range. By  $\text{set}(r)$  we refer to the set defined by an iterator  $r$  (which must coincide with  $\text{set}(s)$ ).

A range iterator naturally hides its implementation. It can iterate a sequence (for instance an array) directly by position, but it can just as well traverse a linked list or the leaves of a balanced tree, or for example iterate over the union of two other iterators.

Iterators are consumed by iteration. Hence, if the same sequence needs to be iterated twice, a fresh iterator is needed. If iteration is cheap, an iterator can support multiple iterations by providing a reset operation. Otherwise, a *cache iterator* takes an arbitrary range iterator as input, iterates it completely, and stores the obtained ranges in an array. Its operations then use the array. The cache iterator implements a reset operation, so that the possibly costly input iterator is used only once, while the cache iterator can be used as often as needed.

**Iterators for variables.** The two basic set-valued operations on integer variables are domain access and domain update. For an integer variable  $x$ , the operation  $x.\text{getdom}()$  returns a range iterator for  $\text{ranges}(d(x))$ . The operation  $x.\text{setdom}(r)$  updates the variable domain of  $x$  to  $\text{set}(r)$  given a range iterator  $r$ , provided that  $\text{set}(r) \subseteq d(x)$ . The responsibility for ensuring that  $\text{set}(r) \subseteq d(x)$  is left to the programmer.

In order to provide safer and richer operations, we can use *iterator combinators*. For example, an *intersection iterator*  $r = \text{iinter}(r_1, r_2)$  combines two range iterators  $r_1$  and  $r_2$  such that  $\text{set}(r) = \text{set}(r_1) \cap \text{set}(r_2)$ . Similarly, a *difference iterator*  $r = \text{iminus}(r_1, r_2)$  yields  $\text{set}(r) = \text{set}(r_1) \setminus \text{set}(r_2)$ .

Richer set-valued operations are then effortless. The operation  $x.\text{adjdom}(r)$  adjusts the domain  $d(x)$  by  $\text{set}(r)$ , yielding  $d(x) \cap \text{set}(r)$ , whereas  $x.\text{excdom}(r)$  excludes  $\text{set}(r)$  from  $d(x)$ , yielding  $d(x) \setminus \text{set}(r)$ :

$$\begin{aligned} x.\text{adjdom}(r) &= x.\text{setdom}(\text{iinter}(x.\text{getdom}(), r)) \\ x.\text{excdom}(r) &= x.\text{setdom}(\text{iminus}(x.\text{getdom}(), r)) \end{aligned}$$

In contrast to the  $x.\text{setdom}(\cdot)$  operation, the richer set-valued operations are inherently contracting, and thus safer to use when implementing a propagator.

Iterators also serve as the natural interface for operations on set variables, which are usually approximated as set intervals defined by a lower and an upper bound [32, 18]:

$$d(x) = [\text{glb}(d(x)) .. \text{lub}(d(x))] = \{s \mid \text{glb}(d(x)) \subseteq s, s \subseteq \text{lub}(d(x))\}$$

In order to access and update these set bounds, propagators use set-valued operations based on iterators:  $x.\text{glb}()$  returns a range iterator for  $\text{ranges}(\text{glb}(d(x)))$ ,  $x.\text{lub}()$  returns a range iterator for  $\text{ranges}(\text{lub}(d(x)))$ ,  $x.\text{adjglb}(r)$  updates the domain of  $x$  to  $[\text{glb}(d(x)) \cup \text{set}(r), \text{lub}(d(x))]$ , and  $x.\text{adjlub}(r)$  updates the domain of  $x$  to  $[\text{glb}(d(x)), \text{lub}(d(x)) \cap \text{set}(r)]$ .

Iterator combinators provide the operations that set propagators need: union, intersection, difference, and complement. Many propagators can thus be implemented directly using iterators and do not require any temporary data structures for storing set-valued intermediate results.

All set-valued operations are parametric with respect to the iterator  $r$ : any range iterator can be used. As for parametric propagators, an implementor has to decide on the kind of parametricity to use. Gecode uses template-based parametric polymorphism, with the performance benefits due to monomorphization and consequent code optimization mentioned previously.

**Advantages.** Range iterators provide essential advantages over an explicit set representation. First, any range iterator regardless of its implementation can be used in domain operations. This turns out to result in simple, efficient, and expressive domain updates. Second, no costly memory management is required to maintain a range iterator as it provides access to only one range at a time. Third, the abstractness of range iterators makes them compatible with views and derived propagators: the necessary view transformations can be encapsulated in an iterator, as discussed below.

**Iterators for views.** As iterators hide their implementation, they are perfectly suited for implementing the transformations required for set-valued operations on views.

Set-valued operations for constant integer views are straightforward. For a constant view  $v$  on constant  $k$ , the operation  $v.\text{getdom}()$  returns an iterator for the singleton range sequence  $\langle [k .. k] \rangle$ . The operation  $v.\text{setdom}(r)$  just checks whether the range sequence of  $r$  is empty (in order to detect failure).

Set-valued operations for an offset view are provided by an *offset iterator*. For a range sequence  $r = \langle [m_i .. n_i] \rangle_{i=1}^k$  and offset  $c$ ,  $\text{ioffset}(r, c)$  iterates  $\langle [m_i + c .. n_i + c] \rangle_{i=1}^k$ . An offset view on  $x$  with offset  $c$  then implements  $\text{getdom}$  as  $\text{ioffset}(x.\text{getdom}(), c)$  and  $\text{setdom}(r)$  as  $x.\text{setdom}(\text{ioffset}(r, -c))$ .

For minus views we just give the range sequence, iteration is obvious. For a given range sequence  $\langle [m_i .. n_i] \rangle_{i=1}^k$ , the negative sequence is obtained by reversal and sign change as  $\langle [-n_{k-i+1} .. -m_{k-i+1}] \rangle_{i=1}^k$ . The same iterator for this sequence can be used both for  $\text{setdom}$  and  $\text{getdom}$  operations. Note that implementing the iterator is involved as it changes direction of the range sequence. There are two different options for changing direction: either the set-valued operations accept iterators in both directions or a cache iterator is used to reverse the direction. Gecode uses the latter and Section 8.2 evaluates the overhead introduced by cache iterators.

A scale iterator provides the necessary transformations for scale views. Assume a scale view on a variable  $x$  with a coefficient  $a > 0$ , and let  $\langle [m_i .. n_i] \rangle_{i=1}^k$  be a range sequence

for  $d(x)$ . If  $a = 1$ , the scale iterator does not change the range sequence. Otherwise, the corresponding scaled range sequence is  $\langle \{a \times m_1\}, \{a \times (m_1 + 1)\}, \dots, \{a \times n_1\}, \dots, \{a \times m_k\}, \{a \times (m_k + 1)\}, \dots, \{a \times n_k\} \rangle$ . For the other direction, assume we want to update the domain using a set  $S$  through a scale view. Assume that  $\langle [m_i .. n_i]_{i=1}^k \rangle$  is a range sequence for  $S$ . Then for  $1 \leq i \leq k$  the ranges  $[[m_i/a] .. [n_i/a]]$  correspond to the required variable domain for  $x$ , however they do not necessarily form a range sequence as the ranges might be empty, overlapping, or adjacent. Iterating the range sequence is however simple by skipping empty ranges and merging overlapping or adjacent ranges. Scale views for a variable  $x$  and a coefficient  $a$  in Gecode are restricted to be strictly positive so as to not change the direction of the scaled range sequence. A negative coefficient can be obtained by using a scale view together with a minus view.

A complement view of a set variable  $x$  uses a *complement iterator*, which given a range iterator  $r$  iterates over  $\overline{\text{set}(r)}$ .

## 7 Limitations

Although views are widely applicable, they are no silver bullet. This section explores some limitations of the presented model.

**Beyond injective views.** Views are required to be injective, as otherwise  $\varphi^- \circ \varphi$  is no longer the identity function, and derived propagators would not necessarily be contracting. An example for this behavior is a view for the absolute value of an integer variable. Assuming a variable domain  $d(x) = \{1\}$ , an absolute value view  $\varphi$  would leave the domain as it is,  $\varphi(d)(x) = \{1\}$ , but the inverse would “invent” the negative value,  $\varphi^-(\varphi(d))(x) = \{-1, 1\}$ . With an adapted definition of derived propagators, such as  $\widehat{\varphi}(p)(d) = \varphi^-(p(\varphi(d))) \cap d$ , non-injective views could be used – however, many of the proofs in this paper rely on injectivity. Correia [8] shows that some theorems hold for non-injective views, for instance when propagators are restricted to bounds reasoning.

**Multi-variable views.** Some multi-variable views that seem interesting for practical applications do not preserve contraction, for instance a view on the sum or product of two variables. The reason is that removing a value through the view would have to result in removing a *tuple* of values from the domain. As domains can only represent Cartesian products, this is not possible in general. Such a view would have two main disadvantages. First, if propagation of the original constraint is strong but does not lead to an actual domain pruning through the views, then the potentially high computational cost for the pruning does not pay off. A cheaper but weaker, dedicated propagation algorithm or a different modeling with stronger pruning is then a better choice. Second, if views do not preserve contraction, then Proposition 5 does not hold. That means that a propagator  $p$  cannot easily detect subsumption any longer, as it would have to detect it for  $\widehat{\varphi}(p)$  instead of just for itself,  $p$ . Systems such as Gecode that disable subsumed propagators (as described in [40]) then lose this potential for optimization.

For contraction-preserving views on multiple variables, all the theorems still hold. Two such views we could identify are a set view of Boolean variables  $[b_1, \dots, b_n]$ , behaving like  $\{i \mid b_i = 1\}$ ; and an integer view of Boolean variables  $[b_1, \dots, b_n]$ , where  $b_i$  is 1 if and only if the integer has value  $i$ ; as well as the inverse views of these two.

Benchmark	time (ms)	mem. (KByte)	failures	propagations
All-Interval (50)	77.74	261	0	6 685
All-Interval (100)	1 703.14	837	0	25 866
Alpha (naive)	33.74	45	7 435	136 179
BIBD (7-3-60)	420.91	7 414	1 306	914 321
Eq-20	0.47	27	54	3 460
Golomb Rulers (Bnd, 10)	170.21	125	8 890	944 651
Golomb Rulers (Dom, 10)	198.61	837	8 890	936 096
Graph Coloring	97.44	6 730	1 100	125 639
Magic Sequence (500)	79.01	8 773	251	84 086
Magic Sequence (GCC, 500)	105.84	612	251	3 460
Partition (32)	564.87	282	42 534	2 253 419
Perfect Square	49.83	6 215	150	291 056
Queens (10)	14.98	29	4 992	43 448
Queens (Dom, 10)	34.55	101	3 940	59 508
Queens (100)	0.70	356	22	455
Queens (Dom,100)	12.67	2 572	8	693
Sorting (400)	507.59	268 510	0	459 501
Social Golfers (8-4-9)	71.03	14 423	32	181 290
Social Golfers (5-3-7)	431.97	3 785	1 174	836 201
Hamming Codes (20-3-32)	392.34	35 273	2 296	753 751
Steiner Triples (9)	39.56	1 768	1 067	297 501
Sudoku (Set, 1)	1.53	116	0	1 779
Sudoku (Set, 4)	2.91	181	1	3 180
Sudoku (Set, 5)	12.05	323	11	12 875

**Table 2** Results for the system with views (vanilla Gecode 3.7.2)

**Propagator invariants.** Propagators typically rely on certain invariants of a variable domain implementation. If idempotency or completeness of a propagator depend on these invariants, type conversion views lead to problems, as the actual variable implementation behind the view may not respect the same invariants.

For example, a propagator for set variables based on the set interval approximation can assume that adjusting the lower bound of a variable does not affect its upper bound. If this propagator is instantiated with a type conversion view for an ROBDD-based set variable (see [Section 4.4](#)), this invariant is violated: if, for instance, the current domain is  $\{\{1, 2\}, \{3\}\}$ , and 1 is added to the lower bound, then 3 is removed from the upper bound (in addition to 2 being added to the lower bound). If a propagator reports that it has computed a fixed point based on the assumption that the upper bound cannot have changed, it may actually not be at a fixed point. This potentially results in incorrect propagation, for instance if the propagator could detect failure if it were run again.

## 8 Evaluation

While [Section 3](#) proved that derived propagators are perfect with respect to the mathematical model, this section shows that in most cases one can also obtain perfect implementations of derived propagators, not incurring any performance penalties compared to dedicated, handwritten propagators.

**Experimental setup.** The experiments are based on Gecode 3.7.2 [14]. All examples were compiled using the GNU C++ compiler gcc 4.4.3, all with the same compiler options (generating 64 bit code with full optimization `-O3`). The experiments were done on an 8-core Intel



Benchmark	time %	mem. %	propagations %
Alpha (naive)	383.21	366.67	673.11
BIBD (7-3-60)	337.74	239.83	268.33
Eq-20	584.80	366.67	704.57
Partition (32)	172.35	179.08	227.77
Perfect Square	113.88	114.42	104.65
Queens (Dom, 10)	163.51	195.05	519.68
Queens (Dom,100)	124.29	102.49	2371.86
Social Golfers (8-4-9)	230.73	243.77	160.22
Social Golfers (5-3-7)	173.60	159.18	151.89
Hamming Codes (20-3-32)	112.12	105.26	99.65
Steiner Triples (9)	124.62	101.81	101.76
Sudoku (Set, 1)	142.30	113.79	110.62
Sudoku (Set, 4)	133.16	126.52	107.55
Sudoku (Set, 5)	132.44	109.91	110.55

**Table 3** Relative performance of decomposition, compared to views

Core i7 at 2.7 GHz running 64 bit Linux. During the experiments, only one example was run at a time, with no other processes except for the usual Linux background processes running, in order to minimize the influence of process scheduling on the results. Runtimes are the average of 20 runs, with a coefficient of deviation less than 2% for all benchmarks. All example programs are available in the Gecode distribution. Table 2 shows the figures for the unmodified Gecode 3.7.2 (pure integer models above, models with integer and set variables below the horizontal line). This version makes full use of views as presented in this paper, and the results of further experiments will be given relative to these numbers. For example, a runtime of 130% means that the example needs 30% more time, while 50% means that it is twice as fast as in vanilla Gecode 3.7.2. The column *time* shows the runtime, *mem.* the peak allocated memory, *failures* the number of failures during search, and *propagations* the number of propagator invocations.

As many of the experimental results rely on the optimization capabilities of the used C++ compiler, we verified that all experiments yield similar results with the Microsoft Visual Studio 2008 C++ compiler.

## 8.1 Views Versus Decomposition

In order to evaluate whether deriving propagators is worth the effort in the first place, this set of experiments compares derived propagators with their decompositions, revealing a significant overhead of the latter.

Table 3 shows the results of these experiments. For *Alpha* and *Eq-20*, linear equations with coefficients are decomposed. For *Queens 100*, we replace the special *all-different-with-offsets* by its decomposition into an *all-different* propagator and binary equality-with-offset propagators. In *BIBD* and *Perfect Square*, we decompose ternary Boolean propagators, implementing  $x \wedge y \leftrightarrow z$  as  $\neg x \vee \neg y \leftrightarrow \neg z$  in *BIBD*, and  $x \vee y \leftrightarrow z$  as  $\neg x \wedge \neg y \leftrightarrow \neg z$  in *Perfect Square*. In the remaining examples, we decompose a set intersection into complement and union propagators.

Some integer examples show a significant overhead of up to six times the runtime and memory when decomposed. The overhead of most set examples as well as *Perfect Square* is moderate, partly because no additional variable was introduced if the model already contained its complement or negation. As to be expected, decomposition often needs signifi-

Benchmark	time %	prop. %	Benchmark	time %	prop. %
All-Interval (50)	99.93	100.00	Partition (32)	107.42	108.66
All-Interval (100)	100.18	100.00	Queens (10)	100.18	100.00
Alpha (naive)	98.55	100.00	Queens (Dom, 10)	101.57	100.00
Golomb Rulers (Bnd, 10)	99.81	99.99	Queens (100)	97.09	100.00
Golomb Rulers (Dom, 10)	94.34	100.00	Queens (Dom,100)	97.22	100.00
Graph Coloring	100.56	100.00	Sorting (400)	103.62	100.00

**Table 4** Relative performance of minus views

cantly more propagation steps, but as the additional steps are performed by cheap propagators (like  $x = y + i$  or  $x = \neg y$ ), the runtime effect is less drastic. *Queens 100* is an extreme case, where 23 times the propagation steps only cause 24% more runtime. The reason is that the scheduling order does not take advantage of the fact that the decompositions are Berge-acyclic as discussed in [Section 5](#).

## 8.2 Impact of Derivation Techniques

The techniques presented in [Section 4](#) have different impacts on the performance of the derived propagators.

**Generalization and specialization.** These techniques can be implemented without any performance overhead compared to a handwritten propagator. This is not surprising as the only potential overhead could be that a function call is not resolved at compile time. For example, a thorough inspection of the code generated by the GNU C++ compiler and the Microsoft Visual Studio C++ compiler shows that they are able to fully inline the operations of offset and scale views.

**Transformation and type conversion.** These techniques can incur an overhead compared to a dedicated implementation, as the transformations performed by the views can sometimes not be removed by compiler optimizations, and type conversions may be costly if the data structures for the variable domains differ significantly.

For example, a propagator instantiated with two minus views of variables  $x$  and  $y$  may include a comparison,  $(-x) < (-y)$ . Due to the invariants guaranteed by views, this is equivalent to  $y < x$ , saving two negations. However, the asymmetry in the two’s complement representation of integers prevents the compiler from performing this optimization. As an experiment to evaluate this effect, we instantiated an *all-different* propagator with minus views. The resulting derived propagator of course implements the same constraint, but incurs the overhead of negation. Similarly, we replaced the max propagator in the *Sort* example with a min (where the propagator for min is derived from the propagator for max) and negated all parameters. According to the results in [Table 4](#), the overhead is mostly negligible.

It is interesting to note that the domain-complete *all-different* propagator, when instantiated with minus views, requires a cache iterator for sequence reversal (as discussed in [Section 6.4](#)). Surprisingly, the overhead of minus views is largely independent of the use of cache iterators which is confirmed in [Section 8.4](#).

Other transformations are translated optimally, such as turning  $(-x) - (-y)$  into  $y - x$ . Boolean negation views also lead to optimal code, as they do not compute  $1 - x$  for a Boolean variable  $x$ , but instead swap the positive and negative operations.

Benchmark	time %	prop. %	Benchmark	time %	prop. %
Social Golfers (8-4-9)	148.45	100.00	Sudoku (Set, 1)	141.83	110.62
Social Golfers (5-3-7)	130.30	100.60	Sudoku (Set, 4)	133.15	107.55
Hamming Codes (20-3-32)	119.98	100.00	Sudoku (Set, 5)	133.31	110.55
Steiner Triples (9)	120.83	100.00			

**Table 5** Relative performance of views compared to dedicated set propagators

Benchmark	time %	Benchmark	time %
All-Interval (50)	111.14	Social Golfers (8-4-9)	125.56
All-Interval (100)	107.09	Social Golfers (5-3-7)	119.52
Alpha (naive)	141.05	Hamming Codes (20-3-32)	120.20
BIBD (7-3-60)	140.48	Steiner Triples (9)	132.25
Eq-20	217.96	Sudoku (Set, 1)	109.39
Golomb Rulers (Bnd, 10)	163.43	Sudoku (Set, 4)	110.57
Golomb Rulers (Dom, 10)	146.19	Sudoku (Set, 5)	109.64
Graph Coloring	105.32		
Magic Sequence (500)	133.72		
Magic Sequence (GCC, 500)	147.63		
Partition (32)	151.28		
Perfect Square	119.07		
Queens (10)	110.52		
Queens (100)	116.39		

**Table 6** Relative performance of virtual method calls

Set-valued transformations can result in non-optimal code. For example, a propagator for ternary intersection,  $x = y \cap z$ , will include an inference  $x.\text{adjglb}(y.\text{glb}() \cap z.\text{glb}())$ . To derive a propagator for  $x = y \cup z$ , we instantiate the intersection propagator with complement views for  $x$ ,  $y$ , and  $z$ , yielding the following inference:

$$\bar{x}.\text{adjglb}(\bar{y}.\text{glb}() \cap \bar{z}.\text{glb}())$$

which amounts to computing

$$x.\text{adjlub}(\overline{y.\text{lub}() \cap z.\text{lub}()})$$

It would be more efficient to implement the equivalent  $x.\text{adjlub}(y.\text{lub}() \cup z.\text{lub}())$  because this requires three set operations less. Unfortunately, no compiler will find this equivalence automatically, as it requires knowledge about the semantics of the set operations. [Table 5](#) compares a dedicated propagator for the constraint  $x \cap y = z$  with a version using complement views and a propagator for  $x \cup y = z$ . The overhead of 20% to 48% does not render views useless for set variables, but it is nevertheless significant.

### 8.3 Templates Versus Virtual Methods

As suggested in [Section 6](#), in C++, compile-time polymorphism using templates results in more efficient code than virtual method calls. The reason is not only that a virtual method call needs an indirection compared to a normal method call, but also (and more importantly) that virtual calls cannot be inlined. To evaluate this, we changed the basic operations of integer variables into virtual methods, such that view operations need one virtual method

Benchmark	time %	Benchmark	time %
All-Interval (50)	102.57	Social Golfers (8-4-9)	462.24
All-Interval (100)	101.84	Social Golfers (5-3-7)	371.74
Golomb Rulers (Bnd, 10)	100.33	Hamming Codes (20-3-32)	245.40
Golomb Rulers (Dom, 10)	101.37	Steiner Triples (9)	290.40
Graph Coloring	102.73	Sudoku (Set, 1)	267.92
Magic Sequence (GCC, 500)	99.60	Sudoku (Set, 4)	261.30
Queens (Dom, 10)	128.27	Sudoku (Set, 5)	267.92
Queens (Dom,100)	110.16		

**Table 7** Relative performance of cache iterators

call. This is not possible for iterator-based operations, as template methods cannot be virtual in C++. In order to obtain a conservative approximation of virtual template methods, we prevented the inlining of iterator-based variable operations, so that the compiler at least has to generate normal function calls. Note that this only affects iterator-based variable operations and we still compile with full optimization including inlining. The results of these experiments appear in [Table 6](#). Virtual method calls and non-inlined iterator-based operations cause a runtime overhead between 5% and 117% for the integer examples (left table), and 9% to 32% for the set examples (right table). The runtime overhead for set examples is lower as the basic operations on set variables are considerably more expensive than the basic operations on integer variables.

#### 8.4 Iterators Versus Temporary Data Structures

The following experiments show that using range iterators improves the efficiency of propagators, compared to the use of explicit set data structures for temporary results.

For the experiments, temporary data structures have been emulated by wrapping all iterator combinators in a cache iterator as described in [Section 6.4](#). That way, every computation on iterators (such as the union of two iterators) is first computed, then stored in a cache iterator, and then read from the cache. The cache iterator thus serves as a set data structure for temporary results of set operations. [Table 7](#) shows the results. For integer propagators, computing with temporary data structures incurs very little overhead. Only the Queens examples, dominated by the *all-different* constraints, show an overhead of more than 10%. For set propagators, which make much more use of iterator combinators than integer propagators, the overhead becomes prohibitive, resulting in up to 4.6 times the runtime. The memory consumption (not shown in the table) does not increase, because iterators are not stored, and only few iterators are active at a time.

These results heavily depend on how the variable domains and temporary data structures are represented. The reported overhead therefore possibly does not carry over if, for instance, bit vectors are used instead of range sequences.

## 9 Conclusion

Folklore techniques for deriving propagator variants from parametric propagators have been used by constraint programming system implementors for decades. This paper has consolidated the existing approaches into the concept of using *views* to derive propagator variants.

Such variants are ubiquitous, and the paper has shown how to *systematically* derive propagators using different types of views, corresponding to techniques such as transformation, generalization, specialization, and type conversion.

Based on a formal, implementation independent model of propagators and views, the paper has identified fundamental properties of views that result in *perfect* derived propagators. The paper has shown that a derived propagator inherits correctness and domain completeness from its original propagator, and bounds completeness given additional properties of the used views.

The paper has presented an implementation architecture for views based on *parametricity*. The propagator implementation is kept parametric over the type of view that is used, so that deriving a propagator amounts to instantiating a parametric propagator with the proper views. This implementation architecture is an orthogonal layer of abstraction that can be implemented on top of any constraint solver.

An empirical evaluation has shown that views have proven invaluable for the implementation of Gecode, saving huge amounts of code to be written and maintained. Furthermore, deriving propagators using C++ templates has been shown to yield competitive (in many cases optimal) performance compared to dedicated handwritten propagators. The experiments have also clarified that deriving propagators is vastly superior to decomposing the constraints into additional variables and simple propagators.

**Acknowledgements** We thank Mikael Lagerkvist, Gert Smolka, and Thibaut Feydy for fruitful discussions, and the reviewers of earlier versions and drafts of this paper for their constructive comments. Christian Schulte has been partially funded by the Swedish Research Council (VR) under grant 621-2004-4953. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program. Most of the research for this paper was done while Guido Tack was at the Programming Systems Lab, Saarland University, Germany.

## References

1. Apt, K.R.: Principles of Constraint Programming. Cambridge University Press, Cambridge, UK (2003)
2. Baptiste, P., Le Pape, C., Nuijten, W.: Constraint-based Scheduling. International Series in Operations Research & Management Science. Kluwer Academic Publishers (2001)
3. Beerl, C., Fagin, R., Maier, D., Yannakakis, M.: On the desirability of acyclic database schemes. *Journal of the ACM* **30**(3), 479–513 (1983)
4. Beldiceanu, N., Katriel, I., Thiel, S.: Filtering algorithms for the same constraint. In: Régin and Rueher [36], pp. 65–79
5. Carlsson, M., Ottosson, G., Carlson, B.: An open-ended finite domain constraint solver. In: H. Glaser, P.H. Hartel, H. Kuchen (eds.) *Programming Languages: Implementations, Logics, and Programs, LNCS*, vol. 1292, pp. 191–206. Springer, Southampton, UK (1997)
6. Choi, C.W., Harvey, W., Lee, J.H.M., Stuckey, P.J.: Finite domain bounds consistency revisited. In: A. Sattar, B.H. Kang (eds.) *AI 2006: Advances in Artificial Intelligence, LNCS*, vol. 4304, pp. 49–58. Springer (2006)
7. Corporation, I.: IBM ILOG CP Optimizer V2.3 User's Manual (2009)
8. Correia, M.: Modern techniques for constraint solving: The CaSPER experience. Ph.D. thesis, Universidade Nova de Lisboa (2010)
9. Correia, M., Barahona, P.: Type parametric compilation of algebraic constraints. In: L.S. Lopes, N. Lau, P. Mariano, L.M. Rocha (eds.) *EPIA, LNCS*, vol. 5816, pp. 201–212. Springer (2009)
10. Driesen, K., Hölzle, U.: The direct cost of virtual function calls in C++. In: *OOPSLA*, pp. 306–323 (1996)
11. Eén, N., Sörensson, N.: An extensible SAT-solver. In: E. Giunchiglia, A. Tacchella (eds.) *Theory and Applications of Satisfiability Testing, LNCS*, vol. 2919, pp. 502–518. Springer, Santa Margherita Ligure, Italy (2004)
12. Feydy, T., Stuckey, P.J.: Lazy clause generation reengineered. In: Gent [15], pp. 352–366

13. Gent, I.P. (ed.): Fifteenth International Conference on Principles and Practice of Constraint Programming, *LNCS*, vol. 5732. Springer (2009)
14. Gent, I.P., Jefferson, C., Miguel, I.: Minion: A fast scalable constraint solver. In: G. Brewka, S. Coradeschi, A. Perini, P. Traverso (eds.) *ECAI, Frontiers in Artificial Intelligence and Applications*, vol. 141, pp. 98–102. IOS Press (2006)
15. Gent, I.P., Jefferson, C., Miguel, I.: Watched literals for constraint propagation in Minion. In: F. Benhamou (ed.) Twelfth International Conference on Principles and Practice of Constraint Programming, *LNCS*, vol. 4204, pp. 182–197. Springer, Nantes, France (2006)
16. Gervet, C.: Conjunto: Constraint logic programming with finite set domains. In: M. Bruynooghe (ed.) International Symposium on Logic Programming, pp. 339–358. MIT Press, Ithaca, NY, USA (1994)
17. Google: Google or-tools. <http://code.google.com/p/or-tools/> (2012)
18. Gosling, J., Joy, B., Steele, G., Bracha, G.: The Java Language Specification, third edn. Addison-Wesley Professional (2005)
19. Harvey, W., Stuckey, P.J.: Improving linear constraint propagation by changing constraint representation. *Constraints* **7**, 173–207 (2003)
20. Hawkins, P., Lagoon, V., Stuckey, P.J.: Solving set constraint satisfaction problems using ROBDDs. *Journal of Artificial Intelligence Research* **24**, 109–156 (2005)
21. Hölzle, U., Ungar, D.: Optimizing dynamically-dispatched calls with run-time type feedback. In: PLDI '94: Proceedings of the ACM SIGPLAN 1994 conference on Programming language design and implementation, pp. 326–336. ACM, New York, NY, USA (1994). DOI <http://doi.acm.org/10.1145/178243.178478>
22. Laburthe, F.: Choco: Implementing a CP kernel. In: N. Beldiceanu, W. Harvey, M. Henz, F. Laburthe, E. Monfroy, T. Müller, L. Perron, C. Schulte (eds.) Proceedings of TRICS: Techniques for Implementing Constraint programming Systems, a post-conference workshop of CP 2000, pp. 71–85 (2000)
23. Lagerkvist, M.Z., Pesant, G.: Modeling irregular shape placement problems with regular constraints. In: First Workshop on Bin Packing and Placement Constraints BPPC'08 (2008)
24. Lagerkvist, M.Z., Schulte, C.: Propagator groups. In: Gent [15], pp. 524–538
25. Mackworth, A.: Consistency in networks of relations. *Artificial Intelligence* **8**(1), 99–118 (1977)
26. Milner, R., Tofte, M., MacQueen, D.: The Definition of Standard ML. MIT Press, Cambridge, MA, USA (1997)
27. Peyton Jones, S.L.: Haskell 98. *Journal of Functional Programming* **13**(1) (2003)
28. Puget, J.F.: PECOS: A high level constraint programming language. In: Proceedings of the first Singapore international conference on Intelligent Systems (SPICIS), pp. 137–142. Singapore (1992)
29. Puget, J.F.: A C++ implementation of CLP. In: Proceedings of the Second Singapore International Conference on Intelligent Systems (SPICIS), pp. B256–B261. Singapore (1994)
30. Puget, J.F.: A fast algorithm for the bound consistency of alldiff constraints. In: Proceedings of the 15th National Conference on Artificial Intelligence, pp. 359–366. AAAI Press, Madison, WI, USA (1998)
31. Quimper, C.G.: Efficient propagators for global constraints. Ph.D. thesis, University of Waterloo, Canada (2006)
32. Régin, J.C., Rueher, M. (eds.): Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, *LNCS*, vol. 3011. Springer (2004)
33. Régin, J.C.: A filtering algorithm for constraints of difference in CSPs. In: Proceedings of the Twelfth National Conference on Artificial Intelligence, pp. 362–367. AAAI Press, Seattle, WA, USA (1994)
34. Schulte, C., Carlsson, M.: Finite domain constraint programming systems. In: F. Rossi, P. van Beek, T. Walsh (eds.) Handbook of Constraint Programming, chap. 14, pp. 495–526. Elsevier Science Publishers, Amsterdam, The Netherlands (2006)
35. Schulte, C., Lagerkvist, M., Tack, G.: Gecode. [www.gecode.org](http://www.gecode.org) (2012)
36. Schulte, C., Stuckey, P.J.: When do bounds and domain propagation lead to the same search space? *Transactions on Programming Languages and Systems* **27**(3), 388–425 (2005)
37. Schulte, C., Stuckey, P.J.: Efficient constraint propagation engines. *Transactions on Programming Languages and Systems* **31**(1), 2:1–2:43 (2008)
38. Schulte, C., Tack, G.: Weakly monotonic propagators. In: Gent [15], pp. 723–730
39. Swedish Institute of Computer Science: SICStus Prolog (2012). [www.sics.se/sicstus/](http://www.sics.se/sicstus/)
40. The CHOCO team: CHOCO. <http://www.emn.fr/z-info/choco-solver/> (2012)
41. The ECLiPS<sup>e</sup> project: ECLiPS<sup>e</sup>. [www.eclipse-clp.org](http://www.eclipse-clp.org) (2012)
42. The Mozart consortium: The Mozart programming system. [www.mozart-oz.org](http://www.mozart-oz.org) (2012)
43. Van Hentenryck, P., Saraswat, V.A., Deville, Y.: Design, implementation, and evaluation of the constraint language cc(FD). *Journal of Logic Programming* **37**(1-3), 293–316 (1998)
44. Vilím, P.:  $O(n \log n)$  filtering algorithms for unary resource constraint. In: Régin and Rueher [36], pp. 335–347