Take Command of Your Constraints!

Sung-Shik Jongmans, Farhad Arbab

Take Command of Your Constraints!

Sung-Shik T.Q. Jongmans and Farhad Arbab

Centrum Wiskunde & Informatica, Amsterdam, Netherlands
[jongmans,farhad]@cwi.nl

Abstract. Constraint automata (CA) are a coordination model based on finite automata on infinite words. Although originally introduced for compositional *modeling* of coordinators, an interesting new application of CA is actually *implementing* coordinators (i.e., compiling CA to executable code). Such an approach guarantees correctness-by-construction and can even yield code that outperforms hand-crafted code. The extent to which these two potential advantages arise depends on the smartness of CA-compilers and the existence of proofs of their correctness.
We present and prove the correctness of a critical optimization for CA-compilers: a sound and complete translation from declarative constraints in transition labels to imperative commands in a sequential language. This optimization avoids expensive calls to a constraint solver at runtime, otherwise performed each time a transition fires, and thereby significantly improves the performance of generated coordination code.

1 Introduction

Context. A promising application domain for coordination languages is programming protocols among threads in multicore applications. One reason for this is a classical software engineering advantage: coordination languages typically provide high-level constructs and abstractions that more easily compose into correct—with respect to programmers' intentions—protocol specifications than do conventional lower-level synchronization mechanisms (e.g., locks or semaphores). However, not only do coordination languages simplify programming protocols, but their high-level constructs and abstractions also leave more room for compilers to perform optimizations that conventional language compilers cannot apply. Eventually, sufficiently smart compilers for coordination languages should be capable of generating code (e.g., in Java or in C) that can compete with carefully hand-crafted code. Preliminary evidence for feasibility of this goal appears elsewhere [1,2]. A crucial step toward adoption of coordination languages for multicore programming, then, is the development of such compilers.

To study the performance advantages of using coordination languages for multicore programming, in ongoing work, we are developing compilation technology for *constraint automata* (CA) [3]. Constraint automata are a general coordination model based on finite automata on infinite words. Every CA models the behavior of a single coordinator; a product operator on CA models the synchronous composition of such coordinators (useful to construct complex coordinators out of simpler ones). Structurally, a CA consists of a finite set of states,

$\{x_1; ; x_2\},$
$d_{x_2} \approx d_{x_1}$

$\{x_1; ; \},$
$d_{y'} \approx d_{x_1}$

$\{; ; x_2\},$
$d_{x_2} \approx d_y$

$\{x_1; ; x_2, x_3\},$
$d_{x_2} \approx d_{x_1}$
$\wedge\, d_{x_3} \approx d_{x_1}$

$\{x_1, x_2; ; x_3\},$
$d_{x_3} \approx f(d_{x_1}, d_{x_2})$

$\{x_1; ; x_2\},$
$R(d_{x_1})$
$\wedge\, d_{x_1} \approx d_{x_2}$

$\{x_1; ; \}, \neg R(d_{x_1})$

$\mathsf{Sync}(x_1; ; x_2)$ $\mathsf{Fifo}[y](x_1; x_2)$ $\mathsf{Repl}(x_1; ; x_2, x_3)$ $\mathsf{BinOp}_f(x_1, x_2; ; x_3)$ $\mathsf{Filter}_R(x_1; ; x_2)$

(a) Abstract/parametric primitives

$\{A, B; C, D; E, F\},$
$d_C \approx \mathtt{add}(d_A, d_B) \wedge d_D \approx d_C$
$\wedge\, d_E \approx d_C \wedge \mathtt{Odd}(d_D) \wedge d_D \approx d_F$

$\{A, B; C, D; E\},$
$d_C \approx \mathtt{add}(d_A, d_B) \wedge d_D \approx d_C$
$\wedge\, d_E \approx d_C \wedge \neg\mathtt{Odd}(d_D)$

Example

$\{A; ; \},\, d_{m'} \approx d_A$

λ_1
λ_2
λ_3
λ_4

Languages2

$\lambda_1 : d_m \approx d_B \wedge d_{C1} \approx d_B \wedge d_{C2} \approx d_B$
$\wedge\, \mathtt{Mandarin}(d_{C1}) \wedge d_{D1} \approx d_{C1}$
$\wedge\, \mathtt{Spanish}(d_{C2}) \wedge d_{D2} \approx d_{C2}$
$\lambda_2 : d_m \approx d_B \wedge d_{C1} \approx d_B \wedge d_{C2} \approx d_B$
$\wedge\, \mathtt{Mandarin}(d_{C1}) \wedge d_{D1} \approx d_{C1}$
$\wedge\, \neg\mathtt{Spanish}(d_{C2})$
$\lambda_3 : d_m \approx d_B \wedge d_{C1} \approx d_B \wedge d_{C2} \approx d_B$
$\wedge\, \neg\mathtt{Mandarin}(d_{C1})$
$\wedge\, \mathtt{Spanish}(d_{C2}) \wedge d_{D2} \approx d_{C2}$
$\lambda_4 : d_m \approx d_B \wedge d_{C1} \approx d_B \wedge d_{C2} \approx d_B$
$\wedge\, \neg\mathtt{Mandarin}(d_{C1}) \wedge \neg\mathtt{Spanish}(d_{C2})$
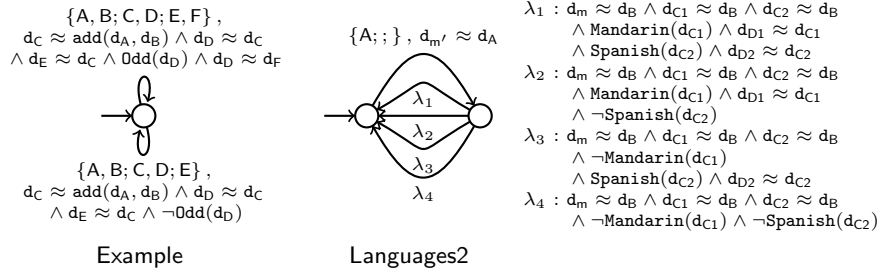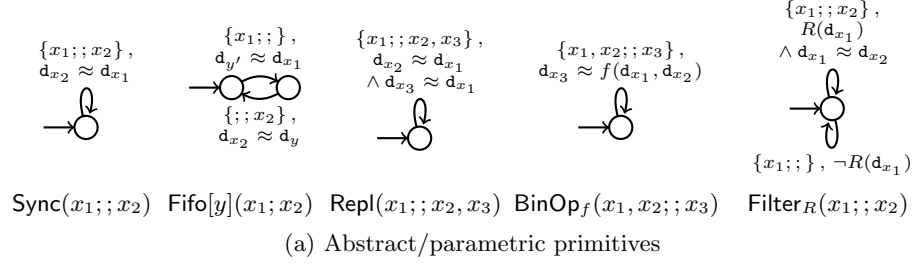
(b) Concrete/instantiated composites

Fig. 1: Example CA. Semicolons separate input/internal/output ports.

a finite set of transitions, a set of directed *ports*, and a set of local *memory cells*. Ports represent the boundary/interface between a coordinator and its co-ordinated agents (e.g., computation threads). Such agents can perform blocking I/O-operations on ports: a coordinator's input ports admit `put` operations, while its output ports admit `get` operations. Memory cells represent internal buffers in which a coordinator can temporarily store data items. Different from classical automata, transition labels of CA consist of two elements: a set of ports, called a *synchronization constraint*, and a logical formula over ports and memory cells, called a *data constraint*. A synchronization constraint specifies which ports need an I/O-operation for its transition to fire (i.e., those ports synchronize in that transition and their pending I/O-operations complete), while a data constraint specifies which particular data items those I/O-operations may involve. Figure 1 already shows some examples; details follow shortly. Essentially, a CA constrains *when* I/O-operations may complete on *which* ports. As such, CA quite literally materialize Wegner's definition of coordination as "constrained interaction" [4].

Given a library of "small" CA, each of which models a primitive coordinator with its own local behavior, programmers can compositionally construct "big" CA, each of which models a composite coordinator with arbitrarily complex global behavior, fully tailored to the needs of these programmers and their programs. Our current CA-compilers can subsequently generate Java/C code. Afterward, these compilers either automatically blend their generated code into programs' computation code or provide programmers the opportunity to do this manually. At run-time, the code generated for a big CA (i.e., a composite co-

Sync$(x_1;;x_2)$ Fifo$[y](x_1;x_2)$ Repl$(x_1;;x_2,x_3)$ BinOp$_f(x_1,x_2;;x_3)$ Filter$_R(x_1;;x_2)$

(a) Abstract/parametric primitives (cf. Figure 1a)

Example (in black) and Fibonacci (in black+gray)   Languages2
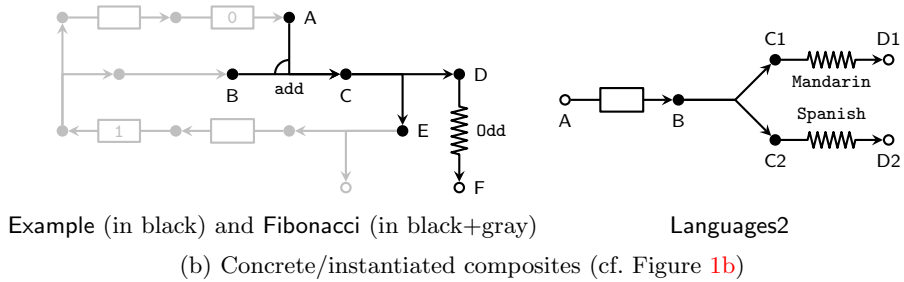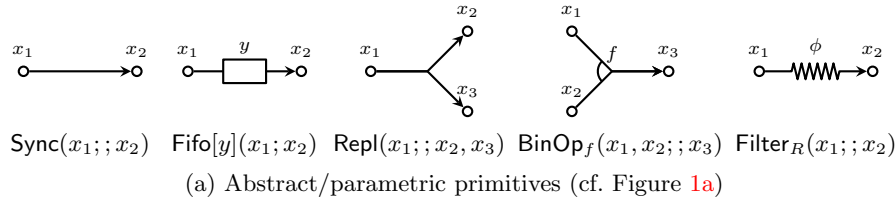
(b) Concrete/instantiated composites (cf. Figure 1b)

Fig. 2: Reo syntax for the CA in Figure 1. White vertices represent input/output ports; black vertices represent internal ports.

ordinator) executes a state machine that simulates that CA, repeatedly firing transitions as computation threads perform I/O-operations. Straightforward as this may seem, one needs to overcome a number of serious issues before this approach can yield practically useful code. Most significantly, these issues include exponential explosion of the number of states or transitions of CA, and oversequentialization or overparallelization of generated code. We have already reported our work on these issues along with promising results elsewhere [5,6,1,7].

Instead of programming with CA directly, one can adopt a more programmer-friendly syntax for which CA serve as semantics. In our work, for instance, we adopted the syntax of Reo [8,9], a graphical calculus of channels. Figure 2 already shows some examples; details follow shortly. (Other CA syntaxes beside Reo exist though [10,11,12,13], which may be at least as programmer-friendly.)

Problem. To fire a transition at run-time, code generated for a CA must evaluate the data constraint of that transition: it must ensure that the data involved in blocking I/O-operations pending on the transition's ports satisfy that constraint.

A straightforward evaluation of data constraints requires expensive calls to a constraint solver. Such calls cause high run-time overhead. In particular, because transitions fire sequentially, avoiding constraint solving to reduce this sequential bottleneck is crucial in getting good performance for the whole program.

Contribution & organization. In this paper, we introduce a technique for statically translating a data constraint, off-line at compile-time, into a *data command*: an imperative implementation (in a sequential language with assignment

and guarded failure) of a data constraint that avoids expensive calls to a constraint solver at run-time. As with our previous optimization techniques [5,1,7], we prove that the translation in this paper is sound and complete. Such correctness proofs are important, because they ensure that our compilation approach guarantees *correctness-by-construction* (e.g., model-checking results obtained for pre-optimized CA also hold for their generated, optimized implementations). We also give preliminary performance results to show our optimization's potential.

In Section 2, we discuss data constraints and CA. In Sections 3 and 4, we discuss our translation algorithm. In Section 5, we give preliminary performance results. Section 6 concludes this paper. Some relatively lengthy formal definitions and detailed proofs of Theorems 1 and 2 appear in a technical report [14].

2 Preliminaries: Data Constraints, Constraint Automata

Data constraints. Let $\mathbb{D}$ denote the finite set of all *data items*, typically ranged over by d. Let $\mathsf{nil} \notin \mathbb{D}$ denote a special object for the *empty data item*. Let $\mathbb{P}$ denote the finite set of all *places* where data items can reside, typically ranged over by x or y; every place models either a port or a memory cells. We model atomic coordination steps—the letters in the alphabet of CA—with elements from the partial function space $\mathbb{DISTR} = \mathbb{P} \rightharpoonup \mathbb{D} \cup \{\mathsf{nil}\}$, called *distributions*, typically ranged over by δ. Informally, a distribution δ associates every place x involved in the step modeled by δ with the data item $\delta(x)$ observable in x.

Let $\mathbb{F} = \bigcup\{\mathbb{D}^k \to \mathbb{D} \mid k > 0\}$ and $\mathbb{R} = \bigcup\{\wp(\mathbb{D}^k) \mid k > 0\}$ denote the sets of all *data functions* and *data relations* of finite arity. Let $\mathbb{DATA}$, $\mathbb{FUN}$, and $\mathbb{REL}$ denote the sets of all *data item symbols*, *data function symbols* and *data relation symbols*, typically ranged over by d, f, and R. Let $\mathsf{arity} : \mathbb{FUN} \cup \mathbb{REL} \to \mathbb{N}_+$ denote a function that associates every data function/relation symbol with its positive arity. Let $\mathcal{I} : \mathbb{DATA} \cup \mathbb{FUN} \cup \mathbb{REL} \to \mathbb{D} \cup \mathbb{F} \cup \mathbb{R}$ denote a bijection that associates every data item/function/relation symbol with its interpretation. A *data term* is a word t generated by the following grammar:

$$t \quad ::= \quad \mathsf{d}_x \mid \mathsf{nil} \mid d \mid f(t_1, \ldots, t_k) \ \textbf{if} \ \mathsf{arity}(f) = k$$

Let $\mathbb{TERM}$ denote the set of all data terms. Let $\mathsf{eval} : \mathbb{DISTR} \times \mathbb{TERM} \to \mathbb{D} \cup \{\mathsf{nil}\}$ denote a function that evaluates every data term t to a data item $\mathsf{eval}_\delta(t)$ under distribution δ. For instance, $\mathsf{eval}_\delta(\mathsf{d}_x) = \delta(x)$—if δ is defined for x—and $\mathsf{eval}_\delta(d) = \mathcal{I}(d)$. If a data term t contains nil or d_x for some $x \notin \mathrm{Dom}(\delta)$, we have $\mathsf{eval}_\delta(t) = \mathsf{nil}$. This ensures that eval is a total function, even though the deltas in $\mathbb{DISTR}$ are partial functions. See also [14, Definition 7]. We call a term of the form d_x a *free variable*. Intuitively, d_x represents the data item residing in place x. Let $\mathsf{Free} : \mathbb{TERM} \to \wp(\mathbb{TERM})$ denote a function that maps every data term t to its set of free variables.

A *data constraint* is a word ϕ generated by the following grammar:

$$
\begin{aligned}
a \quad &::= \quad \bot \mid \top \mid t \approx t \mid t \not\approx \mathsf{nil} \mid R(t_1, \ldots, t_k) \ \textbf{if} \ \mathsf{arity}(R) = k \\
\phi \quad &::= \quad a \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi
\end{aligned}
$$

Let $\mathbb{DC}$ denote the set of all data constraints. We often call $t_1 \approx t_2$ atoms *equalities*. We define the semantics of data constraints over distributions. Let $\models^{\text{dc}} \subseteq \mathbb{D}\text{ISTR} \times \mathbb{DC}$ denote the satisfaction relation on data constraints. Its definition is standard for $\bot$ (contradiction), $\top$ (tautology), $\neg$ (negation), $\vee$ (disjunction), and $\wedge$ (conjunction). For other atoms, we have the following:

$$\delta \models^{\text{dc}} t_1 \approx t_2 \qquad \textbf{iff} \;\; \mathsf{eval}_\delta(t_1) = \mathsf{eval}_\delta(t_2) \neq \mathsf{nil}$$
$$\delta \models^{\text{dc}} t \not\approx \mathtt{nil} \qquad \textbf{iff} \;\; \mathsf{eval}_\delta(t) \neq \mathsf{nil} \qquad \textit{(i.e., notation for } \delta \models^{\text{dc}} t \approx t)$$
$$\delta \models^{\text{dc}} R(t_1, \dots, t_k) \;\; \textbf{iff} \;\; \mathcal{I}(R)(\mathsf{eval}_\delta(t_1), \dots, \mathsf{eval}_\delta(t_k))$$

In the second rule, if a t_i evaluates to nil, the right-hand side is undefined—hence false—because the domain of data relation $\mathcal{I}(R)$ excludes nil. If $\delta \models^{\text{dc}} \phi$, we call δ a *solution* for ϕ. Let $[\![\cdot]\!] : \mathbb{DC} \to \wp(\mathbb{D}\text{ISTR})$ denote a function that associates every data constraint ϕ with its meaning $[\![\phi]\!] = \{\delta \mid \delta \models^{\text{dc}} \phi\}$. We write $\phi \Rightarrow \phi'$ iff $[\![\phi]\!] \subseteq [\![\phi']\!]$. We also extend function Free from data terms to data constraints.

Constraint automata. A constraint automaton (CA) is a tuple $(Q, \mathcal{X}, \mathcal{Y}, \longrightarrow, \imath)$ with Q a set of states, $\mathcal{X} \subseteq \mathbb{P}$ a set of ports, $\mathcal{Y} \subseteq \mathbb{P}$ a set of memory cells, $\longrightarrow \subseteq Q \times (\wp(\mathcal{X}) \times \mathbb{DC}) \times Q$ a transition relation labeled with pairs (X, ϕ), and $\imath \in Q$ an initial state. For every label (X, ϕ), no ports outside X may occur in ϕ. Set $\mathcal{X}$ consists of three disjoint subsets of input ports $\mathcal{X}_{\text{in}}$, *internal ports* $\mathcal{X}_{\text{int}}$, and output ports $\mathcal{X}_{\text{out}}$. We call a CA for which $\mathcal{X}_{\text{int}} = \emptyset$ a *primitive*; otherwise, we call it a *composite*.

Although generally important, we skip the definition of the product operator on CA, because it does not matter in this paper. Every CA accepts infinite sequences of distributions [3]: $(Q, \mathcal{X}, \mathcal{Y}, \longrightarrow, \imath)$ accepts $\delta_0 \delta_1 \cdots$ if an infinite sequence of states $q_0 q_1 \cdots$ exists such that $q_0 = \imath$ and for all $i \geq 0$, a transition $(q_i, (X, \phi), q_{i+1})$ exists such that $\mathrm{Dom}(\delta_i) = X$ and $\delta_i \models^{\text{dc}} \phi$.

Without loss of generality, we assume that all data constraints occur in disjunctive normal form. Moreover, because replacing a transition $(q, (X, \phi_1 \vee \phi_2), q')$ with two transitions $(q, (X, \phi_1), q')$ and $(q, (X, \phi_2), q')$ preserves behavioral congruence on CA [3], without loss of generality, we assume that the data constraint in every label is a conjunction of *literals*, typically ranged over by ℓ.

Figure 1a shows example primitives; Figure 2a shows their Reo syntax. Sync models a synchronous channel from an input x_1 to an output x_2. Fifo models an asynchronous channel with a 1-capacity buffer y from x_1 to x_2. Repl models a coordinator that, in each of its atomic coordination steps, replicates the data item on x_1 to both x_2 and x_3. BinOp models a coordinator that, in each of its atomic coordination steps, applies operation f to the data items on x_1 and x_2 and passes the result to x_3. Filter models a lossy synchronous channel from x_1 to x_2; data items pass this channel only if they satisfy predicate R.

Figure 1b shows example composites; Figure 2b shows their Reo syntax. Example—our running example in this paper—consists of instantiated primitives $\mathsf{BinOp}_{\mathsf{add}}(\mathsf{A}, \mathsf{B}; ; \mathsf{C})$, $\mathsf{Repl}(\mathsf{C}; ; \mathsf{D}, \mathsf{E})$, and $\mathsf{Filter}_{\mathsf{Odd}}(\mathsf{D}; ; \mathsf{F})$, where $\mathtt{add}$ and $\mathtt{Odd}$ have the obvious interpretation. In each of its atomic coordination steps, if the sum of the data items (supposedly integers) on its inputs A and B is odd,

Example passes this sum to its outputs E and F. Otherwise, if the sum is even, Example passes this value only to E. Figure 2b shows that Example constitutes Fibonacci. Fibonacci coordinates two consumers by generating the Fibonacci sequence. Whenever Fibonacci generates an even number, it passes that number to *only one* consumer; whenever it generates an odd number, it passes that number to *both* consumers. Finally, Languages2 consists of instantiated primitives Fifo[m](A; ; B), Repl(B; ; C1, C2), Filter$_{\texttt{Mandarin}}$(C1; ; D1), and Filter$_{\texttt{English}}$(C2; ; D2). Languages2 coordinates a producer and two consumers. If the producer puts a Mandarin (resp. English) data item on input A, Languages2 asynchronously passes this data item only to the consumer on output D1 (resp. D2). Languages2 easily generalizes to Languages$i$, for i different languages; we do so in Section 5.

3 From Data Constraints to Data Commands

At run-time, compiler-generated code executes in one or more CA-*threads*, each of which runs a state machine that simulates a CA. (We addressed the challenge of deciding the number of CA-threads elsewhere [5,6,7].) The *context* of a CA-thread is the collection of put/get operations on implementations of its input/output ports, performed by computation threads. Every time the context of a CA-thread changes, that CA-thread examines whether this change enables a transition in its current state q: for each transition $(q, (X, \phi), q')$, it checks whether every port $x \in X$ has a pending I/O-operation and if so, whether the data items involved in the pending put operations and the current content of memory cells can constitute a solution for ϕ. For the latter, the CA-thread calls a constraint solver, which searches for a distribution δ such that $\delta \stackrel{\text{dc}}{\models} \phi$ and $\delta_{\text{init}} \subseteq \delta$, where:

$$\delta_{\text{init}} = \{x \mapsto d \mid \text{the put pending on input port } x \text{ involves data item } d\} \\ \cup \{y \mapsto d \mid \text{memory cell } y \text{ contains data item } d\} \tag{1}$$

Constraint solving over a finite discrete domain (e.g., $\mathbb{D}$) is NP-complete [15]. Despite carefully and cleverly optimized backtracking searches, using general-purpose constraint solving techniques for solving a data constraint ϕ inflicts not only overhead proportional to ϕ's size but also a constant overhead for preparing, making, and processing the result of the call itself. Although we generally cannot escape using conventional constraint solving techniques, a practically relevant class of data constraints exists for which we can: the data constraints of many CA in practice are in fact declarative specifications of sequences of imperative instructions (including those in Figure 1). In this section, we therefore develop a technique for statically translating such a data constraint ϕ, off-line at compile-time, into a *data command*: a little imperative program that computes a distribution δ such that $\delta \stackrel{\text{dc}}{\models} \phi$ and $\delta_{\text{init}} \subseteq \delta$, without conventional constraint solving hassle. Essentially, we formalize and automate what a programmer would do if he/she were to write an imperative implementation of a declarative specification expressed as a data constraint. By the end of Section 4, we make the class of data constraints supported by our translation precise.

3.1 Data Commands

A data command is a word P generated by the following grammar:

$$P \quad ::= \quad \texttt{skip} \mid x := t \mid \texttt{if } \phi \texttt{ -> } P \texttt{ fi} \mid P \texttt{ ; } P$$

(We often write "value of x" instead of "the data item assigned to x".)

We adopt the following operational semantics of Apt et al. [16]. True to the idea that data commands compute solutions for data constraints, the *state* that a data command executes in is either a function from places to data items—a distribution!—or the distinguished symbol fail, which represents abnormal termination. A *configuration* is a pair of a data command and a state to execute that data command in. Let ε denote the *empty data command*, and equate $\varepsilon \texttt{ ; } P$ with P. Let $\delta[x := \mathsf{eval}_\delta(t)]$ denote an update of δ as usual. The following rules define the transition relation on configurations, denoted by $\Longrightarrow$.

$$\overline{(\texttt{skip}, \delta) \Longrightarrow (\varepsilon, \delta)} \quad \overline{(x := t, \delta) \Longrightarrow (\varepsilon, \delta[x := \mathsf{eval}_\delta(t)])}$$

$$\frac{\delta \models^{\mathrm{dc}} \phi}{(\texttt{if } \phi \texttt{ -> } P \texttt{ fi}, \delta) \Longrightarrow (P, \delta)} \quad \frac{\delta \not\models^{\mathrm{dc}} \phi}{(\texttt{if } \phi \texttt{ -> } P \texttt{ fi}, \delta) \Longrightarrow (\varepsilon, \mathsf{fail})}$$

$$\frac{(P, \delta) \Longrightarrow (P', \delta')}{(P \texttt{ ; } P'', \delta) \Longrightarrow (P' \texttt{ ; } P'', \delta')}$$

Note that $\texttt{if } \phi \texttt{ -> } P \texttt{ fi}$ commands are *failure* statements rather than *conditional* statements: if the current state violates the *guard* ϕ, execution abnormally terminates. The *partial correctness semantics*, which ignores abnormal termination, of a data command P in a state δ is the set of final states $\mathcal{M}(P, \{\delta\}) = \{\delta' \mid (P, \delta) \Longrightarrow^* (\varepsilon, \delta')\}$; its *total correctness semantics* is the set consisting of fail or its final states $\mathcal{M}_{\mathrm{tot}}(P, \{\delta\}) = \{\mathsf{fail} \mid (P, \{\delta\}) \Longrightarrow^* (\varepsilon, \mathsf{fail})\} \cup \mathcal{M}(P, \{\delta\})$.

Shortly, to prove the correctness of our translation from data constraints to data commands, we use Hoare logic [17], where *triples* $\{\phi\} P \{\phi'\}$ play a central role. In such triples, ϕ characterizes the set of input states, P denotes the data command to execute in those states, and ϕ' characterizes the set of output states. A triple $\{\phi\} P \{\phi'\}$ holds in the sense of partial (resp. total) correctness, if $\mathcal{M}(P, \llbracket \phi \rrbracket) \subseteq \llbracket \phi' \rrbracket$ (resp. $\mathcal{M}_{\mathrm{tot}}(P, \llbracket \phi \rrbracket) \subseteq \llbracket \phi' \rrbracket$). To prove properties of data commands, we use the following sound proof systems for partial (resp. total) correctness, represented by $\vdash$ (resp. $\vdash_{\mathrm{tot}}$) and adopted from Apt et al. [16].

$$\overline{\vdash \{\phi\} \texttt{skip} \{\phi\}} \qquad \frac{\begin{array}{c} \vdash \{\phi'\} P \{\phi''\} \\ \textbf{and } \phi \Rightarrow \phi' \\ \textbf{and } \phi'' \Rightarrow \phi''' \end{array}}{\vdash \{\phi\} P \{\phi'''\}} \quad \frac{\vdash \{\phi\} P \{\phi'\} \\ \textbf{and } \vdash \{\phi'\} P' \{\phi''\}}{\vdash \{\phi\} P \texttt{ ; } P' \{\phi''\}}$$

$$\overline{\vdash \{\phi[\mathsf{d}_x := t]\} \, x := t \, \{\phi\}}$$

$$\frac{\vdash \{\phi \wedge \phi_{\mathrm{g}}\} P \{\phi'\}}{\vdash \{\phi\} \texttt{if } \phi_{\mathrm{g}} \texttt{ -> } P \texttt{ fi} \{\phi'\}} \qquad \frac{\phi \Rightarrow \phi_{\mathrm{g}} \textbf{ and } \vdash_{\mathrm{tot}} \{\phi\} P \{\phi'\}}{\vdash_{\mathrm{tot}} \{\phi\} \texttt{if } \phi_{\mathrm{g}} \texttt{ -> } P \texttt{ fi} \{\phi'\}}$$

The first four rules apply not only to $\vdash$ but also to $\vdash_{\mathrm{tot}}$. We use $\vdash$ to prove the soundness of our upcoming translation; we use $\vdash_{\mathrm{tot}}$ to prove its completeness.

3.2 Precedence

Recall the following typical data constraint over ports A, B, C, D, and E, where A and B are inputs, from $\mathsf{Example}$ in Figure 1b (its lower transition):

$$\phi = \mathsf{d_C} \approx \mathtt{add}(\mathsf{d_A}\,,\,\mathsf{d_B}) \wedge \mathsf{d_D} \approx \mathsf{d_C} \wedge \mathsf{d_E} \approx \mathsf{d_C} \wedge \neg\mathtt{Odd}(\mathsf{d_D}) \tag{2}$$

To translate data constraints to data commands, the idea is to enforce equalities, many of which occur in practice, with assignments and to check all remaining literals with failure statements. In the case of ϕ, for instance, we first assign the data items involved in their pending $\mathtt{put}$ operations to A and B, whose symbols are denoted by $\mathcal{I}^{-1}(\delta_{\mathrm{init}}(\mathsf{A}))$ and $\mathcal{I}^{-1}(\delta_{\mathrm{init}}(\mathsf{B}))$, with δ_{init} as defined in (1), page 6. Next, we assign the evaluation of $\mathtt{add}(\mathsf{d_A}\,,\,\mathsf{d_B})$ to C. The order in which we subsequently assign the value of C to D and E does not matter. After the assignment to D, we check $\neg\mathtt{Odd}(\mathsf{d_D})$ with a failure statement. The following data command corresponds to one possible order of the last three steps.

$$\begin{aligned} P = \; &\mathsf{A} := \mathcal{I}^{-1}(\delta_{\mathrm{init}}(\mathsf{A}))\,;\,\mathsf{B} := \mathcal{I}^{-1}(\delta_{\mathrm{init}}(\mathsf{B}))\,;\,\mathsf{C} := \mathtt{add}(\mathsf{d_A}\,,\,\mathsf{d_B})\,; \\ &\mathsf{D} := \mathsf{d_C}\,;\,\mathtt{if}\,\neg\mathtt{Odd}(\mathsf{d_D})\,\mathtt{->}\,\mathtt{skip}\,\mathtt{fi}\,;\,\mathsf{E} := \mathsf{d_C} \end{aligned}$$

If execution of P in an empty initial state successfully terminates, the resulting final state δ should satisfy ϕ (soundness). Moreover, if a δ' exists such that $\delta' \stackrel{\mathrm{dc}}{\models} \phi$ and $\delta_{\mathrm{init}} \subseteq \delta'$, execution of P should successfully terminate (completeness).

Soundness and completeness crucially depend on the order in which assignments and failure statements occur in P. For instance, changing the order of $\mathsf{D} := \mathsf{d_C}$ and $\mathtt{if}\,\neg\mathtt{Odd}(\mathsf{d_D})\,\mathtt{->}\,\mathtt{skip}\,\mathtt{fi}$ yields a data command whose execution always fails (because D does not have a value yet on evaluating the guard of the failure statement). Such a data command is trivially sound but incomplete. Another complication is that not every equality can become an assignment. In a first class of cases, no operand matches d_x. An example is $\mathtt{add}(\mathsf{d_A}\,,\,\mathsf{d_B}) \approx \mathtt{mult}(\mathsf{d_A}\,,\,\mathsf{d_B})$: this equality should become a failure statement, because neither of its two operands can be assigned to the other. In a second class of cases, multiple equality literals have an operand that matches d_x. An example is $\mathsf{C} \approx \mathtt{add}(\mathsf{d_A}\,,\,\mathsf{d_B}) \wedge \mathsf{C} \approx \mathtt{mult}(\mathsf{d_A}\,,\,\mathsf{d_B})$: only one of these equalities should become an assignment, while the other should become a failure statement, to avoid conflicting assignments to C.

To deal with these complications, we define a *precedence relation* on literals that formalizes their dependencies. Recall that the data constraint in every transition label $(X\,,\,\phi)$ is a conjunction of literals. Let L_ϕ denote the set of literals in ϕ, and let $X_{\mathrm{in}} \subseteq X$ denote the set of *input places* (i.e., input ports and memory cells) involved in the transition. From L_ϕ and X_{in}, we construct a set of literals L to account for (i) symmetry of $\approx$ and (ii) the initial values of input places.

$$L = L_\phi \cup \{t_2 \approx t_1 \mid t_1 \approx t_2 \in L_\phi\} \cup \{\mathsf{d}_x \approx \mathcal{I}^{-1}(\delta_{\mathrm{init}}(x)) \mid x \in X_{\mathrm{in}}\} \tag{3}$$

Obviously, $\delta \stackrel{\mathrm{dc}}{\models} \bigwedge L$ implies $\delta \stackrel{\mathrm{dc}}{\models} \phi$ for all δ (i.e., extending L_ϕ to L is sound). Now, let $\prec_L$ denote the precedence relation on L defined by the following rules:

$$\frac{\mathsf{d}_x \approx t\,,\,\ell \in L \;\textbf{and}\; \mathsf{d}_x \in \mathsf{Free}(\ell)}{\mathsf{d}_x \approx t \prec_L \ell} \qquad \frac{\ell_1 \prec_L \ell_2 \prec_L \ell_3 \;\textbf{and}\; \ell_2 \notin \{\ell_1\,,\,\ell_3\}}{\ell_1 \prec_L \ell_3} \tag{4}$$
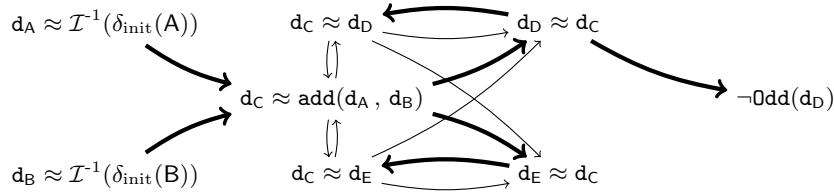
Fig. 3: Fragment of a digraph for an example precedence relation $\prec_L$ (e.g., without loops and without $\mathtt{add}(\mathsf{d_A}\,,\mathsf{d_B}) \approx \mathsf{d_C}$, for simplicity). An arc $(\ell\,,\ell')$ corresponds to $\ell \prec_L \ell'$. Bold arcs represent a strict partial order extracted from $\prec_L$.

Informally, $\mathsf{d}_x \approx t \prec_L \ell$ means that assignment $x := t$ must precede ℓ (i.e., ℓ depends on x). Note also that the first rule deals with the first class of equalities-that-cannot-become-assignments; shortly, we comment on the second class.

For the sake of argument—generally, this is *not* the case—suppose that $\prec_L$ is a strict partial order on L. In that case, we can linearize $\prec_L$ to a total order $<$ on L (i.e., embedding $\prec_L$ into $<$ such that $\prec_L \subseteq <$) with a topological sort on the digraph $(L\,,\prec_L)$ [18,19]. Intuitively, such a linearization gives us an order in which we can translate literals in L to data commands in a sound and complete way. In Section 3.3, we give an algorithm for doing so and indeed prove its correctness. Problematically, however, $\prec_L$ is generally not a strict partial order on L: it is generally neither asymmetric nor irreflexive (i.e., graph-theoretically, it contains cycles). For instance, Figure 3 shows a fragment of the digraph $(L\,,\prec_L)$ for ϕ in (2), page 8, which contains cycles. For now, we defer this issue to Section 4, because it forms a concern orthogonal to our translation algorithm and its correctness. Until then, we simply assume the existence of a procedure for extracting a strict partial order from $\prec_L$, represented by bold arcs in Figure 3.

Henceforth, we assume that every $\mathsf{d}_{x_i} \approx t_i$ literal precedes all differently shaped literals in a linearization of $\prec_L$. Although this assumption is conceptually unnecessary, it simplifies some of our notation and proofs. Formally, we can enforce it by adding a third rule to the definition of $\prec_L$:

$$\frac{\mathsf{d}_x \approx t\,,\ell \in L \ \textbf{ and } \ \left[\ell \neq \mathsf{d}_{x'} \approx t' \ \textbf{ for all } \ x'\,,t'\right]}{\mathsf{d}_x \approx t \prec_L \ell} \tag{5}$$

Proposition 1. *The rule in (5) introduces no cycles.*

(A proof appears in the technical report [14].)

3.3 Algorithm

We start by stating the precondition of our translation algorithm. Suppose that L as defined in (3), page 8, contains n $\mathsf{d}_x \approx t$ literals and m differently shaped literals. Let $\prec_L$ denote a strict partial order on L such that for every $\mathsf{d}_x \approx t \in L$ and for every $\mathsf{d}_y \in \mathsf{Free}(t)$, a $\mathsf{d}_y \approx t'$ literal precedes $\mathsf{d}_x \approx t$ according to $\prec_L$.

Then, let $\ell_1 < \cdots < \ell_n < \ell_{n+1} < \cdots < \ell_{n+m}$ denote a linearization of $\prec_L$, where $\ell_i = \mathtt{d}_{x_i} \approx t_i$ for all $1 \leq i \leq n$. The three rules of $\prec_L$ in Section 3.2 induce precedence relations for which all previous conditions hold, *except* that $\prec_L$ does not necessarily denote a strict partial order; we address this issue in the next section. The previous conditions aside, we also assume $\{\mathtt{d}_{x_1}, \ldots, \mathtt{d}_{x_n}\} = \bigcup \{$ $\mathsf{Free}(\ell_i) \mid 1 \leq i \leq n+m\}$. This extra condition means that for every free variable $\mathtt{d}_{x_i}$ in every literal in L, a $\mathtt{d}_{x_i} \approx t_i$ literal exists in the linearization. If this condition fails, some places can get a value only through search—exactly what we try to avoid—and not through assignment. In such cases, the data constraint is underspecified, and our translation algorithm is fundamentally inapplicable. Finally, we trivially assume that `nil` does not occur syntactically in any literal. A formal definition of this precondition appears in the technical report [14, Figure 10].

Figure 4 shows our algorithm. It first loops over the first n (according to $<$) $\mathtt{d}_x \approx t$ literals. If an assignment for x already exists in data command P, the algorithm translates $\mathtt{d}_x \approx t$ to a failure statement; if not, it translates $\mathtt{d}_x \approx t$ to an assignment. This approach resolves issues with the second class of equalities-that-cannot-become-assignments. After the first loop, the algorithm uses a second loop to translate the remaining m differently shaped literals to failure state-

```
P ← skip
i ← 1
while i ≤ n do
    if d_{x_i} ∈ {d_{x_j} | 1 ≤ j < i} then
        P ← P ; if d_{x_i} ≈ t_i -> skip fi
    else
        P ← P ; x_i := t_i
    i ← i + 1
while i ≤ n + m do
    P ← P ; if ℓ_i -> skip fi
    i ← i + 1
```

Fig. 4: Algorithm to translate data constraints to data commands

ments. The algorithm runs in time linear in $n + m$, and it clearly terminates.

The desired postcondition of the algorithm consists of its soundness and completeness. We define soundness as $\vdash \{\top\} \, P \, \{\ell_1 \wedge \cdots \wedge \ell_{n+m}\}$: after running the algorithm, execution of data command P yields a state that satisfies all literals in L on successful termination. We define completeness as $\big[\big[\delta' \models^{\mathrm{dc}} \ell_1 \wedge \cdots \wedge \ell_{n+m} \ \textbf{implies} \ \vdash_{\mathrm{tot}} \{\top\} \, P \, \{\top\}\big] \ \textbf{for all} \ \delta'\big]$: after running the algorithm, if a distribution δ' exists that satisfies all literals in L, data command P successfully terminated. Although soundness subsequently guarantees that the final state δ satisfies all literals in L, generally, $\delta \neq \delta'$. We use a different proof system for soundness (partial correctness, $\vdash$) than for completeness (total correctness, $\vdash_{\mathrm{tot}}$).

Theorem 1 ([14, Theorem 3]). *The algorithm is sound and complete.*

(A proof appears in the technical report [14].)

4 Handling Cycles

Our algorithm assumes that a precedence relation $\prec_L$ as defined in Section 3.2 is a strict partial order. However, this is generally not the case. In this section, we describe a procedure for extracting a strict partial order from $\prec_L$ without
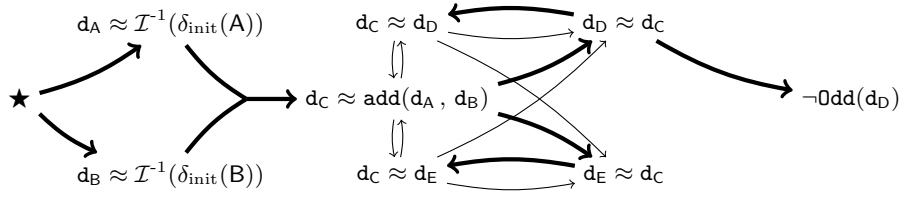
$$d_A \approx \mathcal{I}^{-1}(\delta_{\mathrm{init}}(A)) \qquad d_C \approx d_D \qquad d_D \approx d_C$$

$$\bigstar \qquad d_C \approx \mathsf{add}(d_A\,,\,d_B) \qquad \neg\mathsf{Odd}(d_D)$$

$$d_B \approx \mathcal{I}^{-1}(\delta_{\mathrm{init}}(B)) \qquad d_C \approx d_E \qquad d_E \approx d_C$$

Fig. 5: Fragment of the B-graph corresponding to the digraph in Figure 3 (e.g., without looping B-arcs and without $\mathsf{add}(d_A\,,\,d_B) \approx d_C$, for simplicity). Bold B-arcs represent an arborescence.

losing essential dependencies. We start by adding a distinguished symbol $\bigstar$ to the domain of $\prec_L$, and we extend the definition of $\prec_L$ with the following rules:

$$\frac{\ell \in L \ \textbf{and} \ \mathsf{Free}(\ell) = \emptyset}{\bigstar \prec_L \ell} \qquad \frac{d_x \approx t \in L \ \textbf{and} \ \mathsf{Free}(t) = \emptyset}{\bigstar \prec_L d_x \approx t} \tag{6}$$

These rules state that literals without free variables (e.g., $d_x \approx \mathcal{I}^{-1}(\delta_{\mathrm{init}}(x))$) do not depend on other literals. Now, $\prec_L$ is a strict partial order if the digraph $(L \cup \{\bigstar\}\,,\, \prec_L)$ is a $\bigstar$-*arborescence*: a digraph consisting of $n-1$ arcs such that each of its n vertices is reachable from $\bigstar$ [20]. Equivalently, in a $\bigstar$-arborescence, $\bigstar$ has no incoming arcs, every other vertex has exactly one incoming arc, and the arcs form no cycles [20]. The first formulation is perhaps most intuitive here: every path from $\bigstar$ to some literal ℓ represents an order in which our algorithm should translate the literals on that path to ensure the correctness of the translation of ℓ. The second formulation simplifies observing that arborescences correspond to strict partial orders (by their cycle-freeness).

A naive approach to extract a strict partial order from $\prec_L$ is to compute a $\bigstar$-arborescence of the digraph $(L \cup \{\bigstar\}\,,\, \prec_L)$. Unfortunately, however, this approach generally fails for $d_x \approx t$ literals where t has more than one free variable. For instance, by definition, every arborescence of the digraph in Figure 3 has only one incoming arc for $d_C \approx \mathsf{add}(d_A\,,\,d_B)$, even though assignments to both A and B must precede an assignment to C. Because these dependencies exist as two separate arcs, no arborescence of a digraph can capture them. To solve this, we should somehow represent the dependencies of $d_C \approx \mathsf{add}(d_A\,,\,d_B)$ with a single incoming arc. We can do so by allowing arcs to have multiple tails (i.e., one for every free variable). In that case, we can replace the two separate incoming arcs of $d_C \approx \mathsf{add}(d_A\,,\,d_B)$ with a single two-tailed incoming arc as in Figure 5. The two tails make explicit that to evaluate an add-term, we need values for both its arguments: multiple tails represent a conjunction of dependencies of a literal.

By replacing single-tail-single-head arcs with multiple-tails-single-head arcs, we effectively transform the digraphs considered so far into B-*graphs*, a special kind of hypergraph with only B-*arcs* (i.e., *b*ackward hyperarcs, i.e., hyperarcs with exactly one head) [21]. Deriving a B-graph over literals from a precedence relation as defined in Section 3.2 is generally impossible though: their richer

structure makes B-graphs more expressive—they give more information—than digraphs. In contrast, one can easily transform a B-graph to a precedence relation by splitting B-arcs into single-tailed arcs in the obvious way. Deriving precedence relations from more expressive B-graphs is therefore a correct way of obtaining strict partial orders that satisfy the precondition of our algorithm. Doing so just eliminates information that this algorithm does not care about anyway.

Thus, we propose the following. Instead of formalizing dependencies among literals in a set $L \cup \{\bigstar\}$ directly as a precedence relation, we first formalize those dependencies as a B-graph. If the resulting B-graph is a $\bigstar$-arborescence, we can directly extract a precedence relation $\prec_L$. Otherwise, we compute a $\bigstar$-arborescence of the resulting B-graph and extract a precedence relation $\prec_L$ afterward. Either way, because $\prec_L$ is extracted from a $\bigstar$-arborescence, it is a strict partial order whose linearization satisfies the precondition of our algorithm.

Let $\blacktriangleleft_L$ denote a set of B-arcs on $L \cup \{\bigstar\}$ defined by the following rules, plus the straightforward B-arcs adaptation of the rules in (6), page 11:

$$
\frac{\begin{array}{l} \ell \in L \\ \textbf{and } \mathsf{Free}(\ell) = \{\mathsf{d}_{x_1}, \ldots, \mathsf{d}_{x_k}\} \\ \textbf{and } \mathsf{d}_{x_1} \approx t_1, \ldots, \mathsf{d}_{x_k} \approx t_k \in L \end{array}}{\{\mathsf{d}_{x_1} \approx t_1, \ldots, \mathsf{d}_{x_k} \approx t_k\} \blacktriangleleft_L \ell} \qquad \frac{\begin{array}{l} \mathsf{d}_x \approx t \in L \\ \textbf{and } \mathsf{Free}(t) = \{\mathsf{d}_{x_1}, \ldots, \mathsf{d}_{x_k}\} \\ \textbf{and } \mathsf{d}_{x_1} \approx t_1, \ldots, \mathsf{d}_{x_k} \approx t_k \in L \end{array}}{\{\mathsf{d}_{x_1} \approx t_1, \ldots, \mathsf{d}_{x_k} \approx t_k\} \blacktriangleleft_L \mathsf{d}_x \approx t} \quad (7)
$$

The first rule generalizes the first rule in (4), page 8, by joining sets of dependencies of a literal in a single B-arc. The second rule states that $\mathsf{d}_x \approx t$ literals do not necessarily depend on d_x (as implied by the first rule) but only on the free variables in t: intuitively, a value for x can be derived from values of the free variables in t (cf. assignments). Note that literals can have multiple incoming B-arcs. Such multiple incoming B-arcs represent a disjunction of conjunctions of dependencies. Importantly, as long as all dependencies represented by *one* incoming B-arc are satisfied, the other incoming B-arcs do not matter. An arborescence, which contains one incoming B-arc for every literal, therefore preserves enough dependencies. Shortly, Theorem 2 makes this more precise. Figure 5 shows a fragment of the B-graph for data constraint ϕ in (2), page 8.

One can straightforwardly compute an arborescence of a B-graph $(L \cup \{\bigstar\}, \blacktriangleleft_L)$ with a graph exploration algorithm reminiscent of breadth-first search. Let $\blacktriangleleft_L^{\mathrm{arb}} \subseteq \blacktriangleleft_L$ denote the aborescence under computation, and let $L_{\mathrm{done}} \subseteq L$ denote the set of vertices (i.e., literals in L) that have already been explored; initially, $\blacktriangleleft_L^{\mathrm{arb}} = \emptyset$ and $L_{\mathrm{done}} = \{\bigstar\}$. Now, given some L_{done}, compute a set of vertices L_{next} that are connected only to vertices in L_{done} by a B-arc in $\blacktriangleleft_L$. Then, for every vertex in L_{next}, add an incoming B-arc to $\blacktriangleleft_L^{\mathrm{arb}}$.[1] Afterward, add L_{next} to

---

[1] If a vertex ℓ in L_{next} has multiple incoming B-arcs, the choice among them matters not: the choice is local, because every B-arc has only one head (i.e., adding an ℓ-headed B-arc to $\blacktriangleleft_L^{\mathrm{arb}}$ cannot cause another vertex to get multiple incoming B-arcs, which would invalidate the arborescence). General hypergraphs, whose hyperarcs can have multiple heads, violate this property (i.e., the choice of which hyperarc to add is global instead of local). Computing arborescences of such hypergraphs is NP-complete [22], whereas one can compute aborescences of B-graphs in linear time.

L_{done}. Repeat this process until L_{next} becomes empty. Once that happens, either $\blacktriangleleft_L^{\text{arb}}$ contains an arborescence (if $L_{\text{done}} = L$) or no arborescence exists. This computation runs in linear time, in the size of the B-graph. See also Footnote 1. Given $\blacktriangleleft_L^{\text{arb}}$, the following rules yield a cycle-free precedence relation on $L \cup \{\bigstar\}$:

$$\frac{\{\ell_1, \ldots, \ell_k\} \blacktriangleleft_L^{\text{arb}} \ell \text{ and } 1 \le i \le k}{\ell_i \prec_L \ell} \quad \frac{\ell_1 \prec_L \ell_2 \prec_L \ell_3 \text{ and } \ell_2 \notin \{\ell_1, \ell_3\}}{\ell_1 \prec_L \ell_3} \quad (8)$$

Theorem 2 ([14, **Theorem 4**]). *$\prec_L$ as defined by the rules in (5)(8), pages 9 and 13, is a strict partial order and a large enough subset of $\prec_L$ as defined by the rules in (4)(5)(6), pages 8, 9, and 11, to satisfy the precondition of our translation algorithm in Section 3.3.*

(A proof appears in the technical report [14].) For instance, the bold arcs in Figure 3 represent the precedence relation for the arborescence in Figure 5.

If a $\bigstar$-arborescence of $(L \cup \{\bigstar\}, \blacktriangleleft_L)$ does *not* exist, every $|L|$-cardinality subset of $\blacktriangleleft_L$ has at least one vertex ℓ that is unreachable from $\bigstar$. In that case, by the rules in (6), page 11, ℓ depends on at least one free variable (otherwise, $\{\bigstar\} \blacktriangleleft_L \ell$). Because no B-graph equivalent of a path [23] exists from $\bigstar$ to ℓ, the other literals in L fail to resolve at least one of ℓ's dependencies. This occurs, for instance, when ℓ depends on d_y, while L contains no $\mathsf{d}_y \approx t$ literal. Another example is a recursive literal $\mathsf{d}_x \approx t$ with $\mathsf{d}_x \in \mathsf{Free}(t)$: unless another literal $\mathsf{d}_x \approx t'$ with $t \neq t'$ exists, all its incoming B-arcs contain loops to itself, meaning that no arborescence exists. In practice, such cases inherently require constraint solving techniques to find a value for d_x. Nonexistence of a $\bigstar$-arborescence thus signals a fundamental boundary to the applicability of our translation algorithm (although more advanced techniques of translating some parts of a data constraint to a data command and leaving other parts to a constraint solver are imaginable and left for future work). Thus, the set of data constraints to which our translation algorithm is applicable contains exactly those (i) whose B-graph has a $\bigstar$-arborescence, which guarantees linearizability of the induced precedence, and (ii) that satisfy also the rest of the precondition of our algorithm in Section 3.3.

5 Preliminary Performance Results

In the work that we presented in this paper, we focused on the formal definition of our translation and its proof of correctness. A comprehensive quantitative evaluation remains future work. Indeed, constructing a set of representative examples, identifying independent variables that may influence the outcome (e.g., number of cores, memory architecture, etc.), setting up and performing the corresponding experiments, processing/analyzing the measurements, and eventually presenting the results is a whole other challenge. Still, presenting an optimization technique and not shedding any light on its performance may leave the reader with an unsatisfactory feeling. Therefore, in this section, we provide preliminary performance results to give a rough indication of our translation's merits.

We extended our most recent CA-to-Java compiler and used this compiler to

	Constr.	*Comm.*	$\times$		*Constr.*	*Comm.*	$\times$
Sync	33119333	39800986	1.20	Language2	17278247	24646838	1.43
Fifo	33050122	41398084	1.25	Language4	4423326	11512506	2.60
Replicator	17961129	21803913	1.21	Language6	1062306	5294838	4.98
Example	10573857	12687767	1.20	Language8	194374	1746440	8.98
Fibonacci	1818671	88947751	48.91	Language10	25649	362050	14.12

Fig. 6: Preliminary performance results for ten coordinators. Column "*Constr.*" shows results for constraint-based implementations (in number of coordination steps completed in four minutes); column "*Comm.*" shows restults for command-based implementations; column "$\times$" shows the ratio of the second over the first.

generate both *constraint-based* coordination code (i.e., generated without our translation) and *command-based* coordination code (i.e., generated with our translation) for ten coordinators modeled as CA: three elementary primitives from Figures 1a and 2a (to see how our optimization affects such basic cases) and seven more complex composites, including those in Figures 1b and 2b. See Section 2 for a discussion of these coordinators' behavior. The constraint-based implementations use a custom constraint solver with constraint propagation [24], tailored to our setting of data constraints. The data commands in the generated command-based implementations are imperative Java code, very similar to what programmers would hand-craft (modulo style).

In total, thus, we generated twenty coordinators in Java. We ran each of those implementations ten times on a quadcore machine at 2.4 GHz (no Hyper-Threading; no Turbo Boost) and averaged our measurements. In every run, we warmed up the JVM for thirty seconds before starting to measure the number of coordination steps that an implementation could finish in the subsequent four minutes. Figure 6 shows our results. The command-based implementations outperform their constraint-based versions in all cases. The Language$i$ coordinators furthermore show that the speed-up achieved by their command-based implementations increases as i increases. This may suggest that our optimization becomes relatively more effective as the size/complexity of a coordinator increases, as also witnessed by Fibonacci. Figure 6 shows first evidence for the effectiveness of our translation in practice, although further study is necessary.

6 Discussion

In constraint programming, it is well-known that "if domain specific methods are available they should be applied *instead* [sic] of the general methods" [24, page 2]. The work presented in this paper takes this guideline to an extreme: essentially, every data command generated for a data constraint ϕ by our translation algorithm is a little constraint solver capable of solving only ϕ, with good performance. This good performance comes from the fact that the order of performing assignments and failure statements has already been determined at compile-time.

Moreover, this precomputed order guarantees that backtracking is unnecessary: the data constraint for ϕ finds a solution if one exists without search (i.e., Theorem 1). In contrast, general constraint solvers need to do this work, which our approach does at compile-time, as part of the solving process at run-time.

Execution of data commands bears similarities with *constraint propagation* techniques [24], in particular with *forward checking* [25,26]. Generally, in constraint propagation, the idea is to reduce the search space of a given *constraint satisfaction problem* (CSP) by transforming it into an equivalent "simpler" CSP, where variables have smaller domains, or where constraints refer to fewer variables. With forward checking, whenever a variable x gets a value v, a constraint solver removes values from the domains of all subsequent variables that, together with v, violate a constraint. In the case of an equality $x = y$, for instance, forward checking reduces the domain of y to the singleton $\{v\}$ after an assignment of v to x. That same property of equality is implicitly used in executing our data commands (i.e., instead of representing the domain of a variable and the reduction of this domain to a singleton explicitly, we directly make an assignment).

Our translation from data constraints to data commands may also remind one of classical *Gaussian eliminination* for solving systems of linear equations over the reals [24]: there too, variables are ordered and values/expressions for some variables are substituted into other expressions. The difference is that we have functions, relations, and our data domain may include other data types, which makes solving data constraints directly via Gaussian elimination at least not obvious. However, Gaussian elimination does seem useful as a preprocessing step for translating certain data constraints to data commands that our current algorithm does not support. Future work should clarify this possibility.

Clarke et al. worked on purely constraint-based implementations of coordinators [27]. Essentially, they specify not only the transition labels of a CA as boolean constraints but also its state space and transition relation. In recent work, Proença and Clarke developed a variant of compile-time *predicate abstraction* to improve performance [28]. They used this technique also to allow a form of interaction between the constraint solver and external components during constraint solving [29]. The work of Proença and Clarke resembles ours in the sense that we all try to "simplify" constraints at compile-time. Main differences are that (i) we fully avoid constraint solving and (ii) we consider a richer language of data constraints. For instance, Proença and Clarke have only unary functions in their language, which would have cleared our need for B-graphs.

References

1. Jongmans, S.S., Halle, S., Arbab, F.: Automata-based Optimization of Interaction Protocols for Scalable Multicore Platforms. In: COORDINATION 2014. Volume 8459 of LNCS. Springer (2014) 65–82
2. Jongmans, S.S., Halle, S., Arbab, F.: Reo: A Dataflow Inspired Language for Multicore. In: DFM 2013, IEEE (2014) 42–50
3. Baier, C., Sirjani, M., Arbab, F., Rutten, J.: Modeling component connectors in Reo by constraint automata. SCP **61**(2) (2006) 75–113

4. Wegner, P.: Coordination as Constrained Interaction (Extended Abstract). In: COORDINATION 1996. Volume 1061 of LNCS. Springer (1996) 28–33
5. Jongmans, S.S., Arbab, F.: Global Consensus through Local Synchronization. In: FOCLASA 2013. Volume 393 of CCIS. Springer (2013) 174–188
6. Jongmans, S.S., Arbab, F.: Toward Sequentializing Overparallelized Protocol Code. In: ICE 2014. Volume 166 of EPTCS. CoRR (2014) 38–44
7. Jongmans, S.S., Santini, F., Arbab, F.: Partially-Distributed Coordination with Reo. In: PDP 2014, IEEE (2014) 697–706
8. Arbab, F.: Reo: a channel-based coordination model for component composition. MSCS **14**(3) (2004) 329–366
9. Arbab, F.: Puff, The Magic Protocol. In: Talcott Festschrift. Volume 7000 of LNCS. Springer (2011) 169–206
10. Arbab, F., Kokash, N., Meng, S.: Towards Using Reo for Compliance-Aware Business Process Modeling. In: ISoLA 2008. Volume 17 of CCIS. Springer (2008) 108–123
11. Changizi, B., Kokash, N., Arbab, F.: A Unified Toolset for Business Process Model Formalization. In: FESCA 2010. (2010) 147–156
12. Meng, S., Arbab, F., Baier, C.: Synthesis of Reo circuits from scenario-based interaction specifications. SCP **76**(8) (2011) 651–680
13. Bliudze, S., Sifakis, J.: Causal semantics for the algebra of connectors. FMSD **36**(2) (2010) 167–194
14. Jongmans, S.S., Arbab, F.: Take Command of Your Constraints! (Technical Report). Technical Report FM-1501, CWI (2015)
15. Russell, S., Norvig, P.: Artificial Intelligence. 2nd edn. Prentice Hall (2003)
16. Apt, K., de Boer, F., Olderog, E.R.: Verification of Sequential and Concurrent Programs. 3rd edn. Springer (2009)
17. Hoare, T.: An Axiomatic Basis for Computer Programming. CACM **12**(10) (1969) 576–580
18. Kahn, A.: Topological Sorting in Large Networks. CACM **5**(11) (1962) 558–562
19. Knuth, D.: Fundamental Algorithms. 3rd edn. Volume 1 of The Art of Computer Programming. Addison-Wesley (1997)
20. Korte, B., Vygen, J.: Combinatorial Optimization: Theory and Algorithms. 4th edn. Volume 21 of Algorithms and Combinatorics. Springer (2008)
21. Gallo, G., Longo, G., Pallottino, S., Nguyen, S.: Directed hypergraphs and applications. DAM **42**(2–3) (1993) 177–201
22. Woeginger, G.: The complexity of finding arborescences in hypergraphs. IPL **44**(3) (1992) 161–164
23. Ausiello, G., Franciosa, P., Frigioni, D.: Directed Hypergraphs: Problems, Algorithmic Results, and a Novel Decremental Approach. In: ICTCS 2001. Volume 2202 of LNCS. Springer (2001) 312–328
24. Apt, K.: Principles of Constraint Programming. Cambridge University Press (2009)
25. Bessière, C., Meseguer, P., Freuder, E., Larrosa, J.: On forward checking for non-binary constraint satisfaction. Artificial Intelligence **141**(1–2) (2002) 205–224
26. McGregor, J.: Relational consistency algorithms and their application in finding subgraph and graph isomorphism. Information Science **19** (1979) 229–250
27. Clarke, D., Proença, J., Lazovik, A., Arbab, F.: Channel-based coordination via constraint satisfaction. SCP **76**(8) (2011) 681–710
28. Proença, J., Clarke, D.: Data Abstraction in Coordination Constraints. In: FOCLASA 2013. Volume 393 of CCIS. Springer (2013) 159–173
29. Proença, J., Clarke, D.: Interactive Interaction Constraints. In: COORDINATION 2013. Volume 7890 of LNCS. Springer (2013) 211–225