

# Coordination via Interaction Constraints I: Local Logic

Dave Clarke

Dept. Computer Science, Katholieke Universiteit Leuven,  
Celestijnenlaan 200A,  
3001 Heverlee, Belgium  
dave.clarke@cs.kuleuven.be

José Proença\*

CWI,  
Science Park 123,  
1098 XG Amsterdam, The Netherlands  
jose.proenca@cw.nl

Wegner describes coordination as constrained interaction. We take this approach literally and define a coordination model based on *interaction constraints* and partial, iterative and interactive constraint satisfaction. Our model captures behaviour described in terms of *synchronisation* and *data flow constraints*, plus various modes of interaction with the outside world provided by *external constraint symbols*, *on-the-fly constraint* generation, and *coordination variables*. Underlying our approach is an engine performing (partial) constraint satisfaction of the sets of constraints. Our model extends previous work on three counts: firstly, a more advanced notion of external interaction is offered; secondly, our approach enables local satisfaction of constraints with appropriate partial solutions, avoiding global synchronisation over the entire constraints set; and, as a consequence, constraint satisfaction can finally occur concurrently, and multiple parts of a set of constraints can be solved and interact with the outside world in an asynchronous manner, unless synchronisation is required by the constraints.

This paper describes the underlying logic, which enables a notion of *local solution*, and relates this logic to the more global approach of our previous work based on classical logic.

## 1 Introduction

Coordination models and languages [15] address the complexity of systems of concurrent, distributed, mobile and heterogeneous components, by separating the parts that perform the computation (the components) from the parts that “glue” these components together. The glue code offers a layer between components to intercept, modify, redirect, synchronise communication among components, and to facilitate monitoring and managing their resource usage, typically separate from the resources themselves.

Wegner describes coordination as constrained interaction [16]. We take this approach literally and represent coordination using constraints. Specifically, we take the view that a component connector specifies a (series of) constraint satisfaction problems, and that valid interaction between a connector and its environment corresponds to the solutions of such constraints.

In previous work [5] we took the channel-based coordination model Reo [1], extracted constraints underlying each channel and their composition, and formulated behaviour as a constraint satisfaction problem. There we identified that interaction consisted of two phases: solving and updating constraints. Behaviour depends upon the current state. The semantics were described *per-state* in a series of *rounds*. Behaviour in a particular step is phrased in terms of *synchronisation and data flow constraints*, which describe the synchronisation and the data flow possibilities of participating ports. *Data flow* on the end of a channel occurs when a single datum is passed through that end. Within a particular round data flow may occur on some number of ends; this is equated with the notion of *synchrony*. The constraints were

---

\*Supported by FCT grant 22485 - 2005, Portugal.

based on a synchronisation and a data flow variable for each port. Splitting the constraints into synchronisation and data flow constraints is very natural, and it closely resembles the constraint automata model [3]. These constraints are solved during the solving phase. Evolution over time is captured by incorporating state information into the constraints, and updating the state information between solving phases. Stronger motivation for the use of constraint-based techniques for the Reo coordination model can be found in our previous work [5]. By abstracting from the channels metaphor and using only the constraints, the implementation is free to optimise constraints, eliminating costly infrastructure, such as unnecessary channels. Furthermore, constraint-solving techniques are well studied in the literature, and there are heuristics to search efficiently for solution, offering significant improvement of other models underlying Reo implementations. To increase the expressiveness and usefulness of the model, we added external state variables, external function symbols and external predicates to the model. These external symbols enable modelling of a wider range of primitives whose behaviour cannot be expressed by constraints, either because the internal constraint language is not expressive enough, or to wrap external entities, such as those with externally maintained state. The constraint satisfaction process was extended with means for interacting with external entities to resolve external function symbols and predicates.

In this paper, we make three major contributions to the model:

**Partiality** Firstly, we allow solutions for the constraints and the set of known predicates and functions to be partial [4]. We introduce a minimal notion of partial solution which admits solutions only on relevant parts (variables) of a connector. External symbols that are only discovered on-the-fly are more faithfully modelled in a partial setting.

**Locality** Secondly, we assume a *do nothing* solution for the constraints of each primitive exists, where no data is communicated. This assumption, in combination with partiality, allows certain solutions for part of a connector to be consistently extended to solutions for the full connector. Furthermore, our notion of locality enables independent parts of the connector to evolve concurrently.

**Interaction** Thirdly, we formalise the constraint satisfaction process with respect to the interaction with the external world, and we introduce external constraint symbols. These can be seen as lazy constraints, which are only processed by the engine on demand, by consulting an external source. These can be used to represent, for example, a stream of possible choices, which are requested on demand, such as the pages of different flight options available on an airline booking web page.

**Organization** The next section gives an overview of the approach taken in this paper, providing a global picture and relating the different semantics we present for our constraints. The rest of the paper is divided into two main parts. The first part describes how constraints are defined, and defines four different semantics for variants of the constraint language and relates them. We present a classical semantics in § 3 and two partial semantics in § 4, and exploit possible concurrency by allowing local solutions in § 5. The second part introduces a constraint-based engine to perform the actual coordination, search and applying solutions for the constraints. We describe stateful primitives in § 6, and add interaction in § 7. We give some conclusions about this work in § 8.

## 2 Coordination = Scheduling + Data Flow

We view coordination as a constraint satisfaction problem, where solutions for the constraints yield how data should be communicated among components. More specifically, solutions to the constraints describe *where* and *which* data flow. Synchronisation variables describe the where, and data flow variables

describe the which. With respect to our previous work [5], we move from a classical semantics to a local semantics, where solutions address only part of the connector, as only a relevant subset of the variables of the constraints are required for solutions. We do this transformation from classical to local semantics in a stepwise manner, distinguishing four different semantics that yield different notions of valid solution  $\sigma$ , mapping synchronisation and data flow variables to appropriate values:

#### Classical semantics

- $\sigma$  are always total (for the variables of the connector under consideration);
- an explicit value NO-FLOW is added to the data domain to represent the data value when there is no data flow;
- an explicit *flow axiom* is added to constraints to ensure the proper relationship between synchronisation variables and data flow variables; and
- constraints are solved globally across the entire ‘connector’.

#### Partial semantics

- $\sigma$  may be partial, not binding all variables in a constraint;
- the NO-FLOW value is removed and modelled by leaving the data flow variable undefined; and
- as the previous flow axiom is no longer expressible, the relationship between synchronisation and data flow variables is established by a new meta-flow axiom, which acts after constraints have been solved to filter invalid solutions.

#### Simple semantics

- $\sigma$  is partial, and the semantics is such that only certain “minimal” solutions, which define only the necessary variables, are found; and
- the meta-flow axiom is expressible in this logic, so a *simple* flow axiom can again be added to the constraints.

#### Local semantics

- formulæ are partitioned into blocks, connected via shared variables;
- each block is required to always admit a *do nothing* solution;
- some solutions in a block can be found without looking at its neighbours, whenever there is no-flow on its *boundary* synchronisation variables;
- two or more such solutions are locally compatible;
- blocks can be merged in order to find more solutions, in collaboration, when existing solutions do not ensure the no-flow condition over the *boundary* synchronisation variables; and
- the search space underlying constraints is smaller than in the previous semantics, and there is a high degree of locality and concurrency.

We present formal embeddings between these logics, with respect to solutions that obey the various (meta-) flow axioms (linking solutions for synchronisation and data flow variables). We call such solutions *firings*. The first step is from a classical to a partial semantics. The number of solutions increases, as new (smaller) solutions also become valid. We then move to a simple semantics to regain an expressible flow axiom, where only some “minimal” partial solutions are accepted. In the last step we present a local semantics, where we avoid the need to inspect even more constraints, namely, we avoid visiting constraints added conjunctively to the system, by introducing some requirements on solutions to blocks of constraints.

### 3 Coordination via Constraint Satisfaction

In previous work we described coordination in terms of constraint satisfaction. The main problem with that approach is that the constraints needed to be solved globally, which means that it is not scalable as the basis of an engine for coordination. In this section, we adapt the underlying logic and notion of solution to increase the amount of available locality and concurrency in the constraints. Firstly, we move from the standard classical interpretation of the logic to a partial interpretation. This offers some improvement, but the solutions of a formula need to be filtered using a semantic variant of the flow axiom, which is undesirable because filtering them out during the constraint satisfaction process could be significantly faster. We improve on this situation by introducing a simpler notion of solution for formulæ, requiring only relevant variables to be assigned. This approach avoids post hoc filtering of solutions. Unfortunately, even simple solutions still require more or less global constraint satisfaction. Although it is the case that many constraints may correspond to no behaviour within parts of the connector—indeed all constraints admit such solutions—the constraint satisfier must still visit the constraints to determine this fact. In the final variant, we simply assume that the no behaviour solution can be chosen for any constraint not visited by the constraint solver, and thus the constraint solver can find solutions to constraints without actually visiting all constraints. This means that more concurrency is available and different parts of the implicit constraint graph can be solved independently and concurrently.

We start by motivating our work via an example, and we then describe the classical approach to constraint satisfaction and its problems, before gradually refining our underlying logic to be more amenable to scalable satisfaction.

#### 3.1 Coordination of a complex data generator

We introduce a motivating example, depicted in Figure 1, where a *Complex Data Generator (CDG)* sends data to *Client*. Data communication is controlled via a coordinating connector. The connector consists of a set of composed coordination building blocks, each with some associated constraints describing their behavioural possibilities. We call these building blocks simply *primitives*. The CDG and the Client are also primitives, and play the same coordination game by providing some constraints reflecting their intentions to read or write data.




Figure 1: Network of constraints: coordinating a complex data generator.

Figure 1 uses a graphical notation to depict how the different primitives are connected. Each box represents a primitive with some associated constraints, connected to each other via shared variables. For example, the CDG shares variable  $a$  with  $FIFO_1$ ,  $Filter(p)$ , and  $SyncDrain$ , indicating that the same data flows through the lines connecting these primitives in the figure. The arrows represent the direction of data flow, thus the value of  $a$  is given by  $CDG$  and further constrained by the other attached primitives.

Most of the coordination primitives are channels from the Reo coordination language [1]. Previous work [5] described a constraint-based approach to modelling Reo networks. Here we forego the graphical representation to emphasise the idea that a coordinating connector can be seen as a soup of constraints linked by shared variables. One optimisation we immediately employ is using the same name for ends which are connected by synchronous channels or replicators.<sup>1</sup> Note that in the original description of Reo, nodes act both as mergers and replicators. This behaviour can be emulated using merger and replicator primitives, as we have done. The result is a simpler notion of node, a 1:1 node which both synchronises and copies data from the output port to the input port. Primitives act as constraint providers, which are combined until they reach a consensus regarding where and which data will be communicated. Only then a possible communication of data takes place, and the primitives update their constraints.

In the particular case of the example in Figure 1, there is a complex data generator (CDG) that can write one of several possible values, a filter that can only receive (and send) data if it validates a given predicate  $p$ , a component (User approval) that approves values based on interaction with a user, a destination client that receives the final result, and some other primitives that impose further constraints. We will come back to this example after introducing some basic notions about the constraints.

**Notation** We write  $\mathcal{Data}$  to denote a global set of possible data that can flow on the network. NO-FLOW is a special constant not in  $\mathcal{Data}$  that represents no data flow.  $\mathcal{X}$  denotes a set of *synchronisation variables* over  $\{\mathbf{tt}, \mathbf{ff}\}$ ,  $\widehat{\mathcal{X}} = \{\widehat{x} \mid x \in \mathcal{X}\}$  a set of *data flow variables* over  $\mathcal{Data} \cup \{\text{NO-FLOW}\}$ ,  $\mathcal{P}$  a set of predicate symbols, and  $\mathcal{F}$  a set of function symbols such that  $\mathcal{Data} \subseteq \mathcal{F}$ . (Actually,  $\mathcal{Data}$  is the Herbrand universe over function symbols  $\mathcal{F}$ .) We use the following variables to range over various domains:  $x \in \mathcal{X}$ ,  $\widehat{x} \in \widehat{\mathcal{X}}$ ,  $f \in \mathcal{F}$ , and  $p \in \mathcal{P}$ . Recall that synchronisation variables  $\mathcal{X}$  and data flow variables  $\widehat{\mathcal{X}}$  are intimately related, as one describes whether data flows and the other describes what the data is.

### 3.2 Classical Semantics

Consider the logic with the following syntax of formulæ ( $\psi$ ) and terms ( $t$ ):

$$\begin{aligned} \psi &::= \mathbf{tt} \mid x \mid \psi_1 \wedge \psi_2 \mid \neg \psi \mid p(t_1, \dots, t_n) \\ t &::= \widehat{x} \mid f(t_1, \dots, t_n) \end{aligned}$$

$\mathbf{tt}$  is *true*. We assume that one of the internal predicates in  $\mathcal{P}$  is equality, which is denoted using the standard infix notation  $t_1 = t_2$ . The other logical connectives can be encoded as usual:  $\mathbf{ff} \hat{=} \neg \mathbf{tt}$ ;  $\psi_1 \vee \psi_2 \hat{=} \neg(\neg \psi_1 \wedge \neg \psi_2)$ ;  $\psi_1 \rightarrow \psi_2 \hat{=} \neg \psi_1 \vee \psi_2$ ; and  $\psi_1 \leftrightarrow \psi_2 \hat{=} (\psi_1 \rightarrow \psi_2) \wedge (\psi_2 \rightarrow \psi_1)$ . Constraints can be easily extended with an existential quantifier, provided that it does not appear in a negative position, or alternatively, that it is used only at the top level.

The semantics is based on a relation  $\sigma, \mathcal{I} \models_C \psi$ , where  $\sigma$  is a total map from  $\mathcal{X}$  to  $\{\mathbf{tt}, \mathbf{ff}\}$  and from  $\widehat{\mathcal{X}}$  to  $\mathcal{Data} \cup \{\text{NO-FLOW}\}$ , and  $\mathcal{I}$  is an arity-indexed total map from  $\mathcal{P}_n \times \mathcal{T}^n$  to  $\{\mathbf{tt}, \mathbf{ff}\}$ , for each  $n \geq 0$ , where  $\mathcal{P}_n$  is the set of all predicate symbols of arity  $n$ ,  $\mathcal{T}$  is the set of all possible ground terms (terms with no variables) plus the constant NO-FLOW. The semantics is defined by a satisfaction relation  $\models_C$  defined as follows. The function  $Val_\sigma$  replaces all variables  $v$  by  $\sigma(v)$ , and we assume that  $Val_\sigma(f(t_1, \dots, t_n)) = \text{NO-FLOW}$  whenever  $t_i = \text{NO-FLOW}$ , for some  $i \in 1..n$ .

<sup>1</sup>Semantically, this view of synchronous channels and replicators is valid.

**Definition 1 (Classical Satisfaction)**

$$\begin{array}{ll}
\sigma, \mathcal{I} \models_C \text{tt} & \text{always} \\
\sigma, \mathcal{I} \models_C x & \text{iff } \sigma(x) = \text{tt} \\
\sigma, \mathcal{I} \models_C \psi_1 \wedge \psi_2 & \text{iff } \sigma, \mathcal{I} \models_C \psi_1 \text{ and } \sigma, \mathcal{I} \models_C \psi_2 \\
\sigma, \mathcal{I} \models_C \neg\psi & \text{iff } \sigma, \mathcal{I} \not\models_C \psi \\
\sigma, \mathcal{I} \models_C p(t_1, \dots, t_n) & \text{iff } p(\text{Val}_\sigma(t_1), \dots, \text{Val}_\sigma(t_n)) \mapsto \text{tt} \in \mathcal{I}
\end{array}$$

◁

The following axiom relates synchronisation and data flow variables, stating that a synchronisation variable being set to `ff` corresponds to the corresponding data flow being `NO-FLOW`.

**Axiom 1 (Flow Axiom)**

$$\neg x \leftrightarrow \hat{x} = \text{NO-FLOW} \quad (\text{flow axiom})$$

We introduced this in our previous approach for coordination via constraints [5]. Every pair of variables,  $x$  and  $\hat{x}$ , is expected to obey this axiom. Write  $FA(x)$  for the flow axiom for variables  $x, \hat{x}$  and  $FA(X)$  for the conjunction  $\bigwedge_{x \in X} FA(x)$ . Also write  $fv(\psi)$  for the free variables of  $\psi$ , i.e., variables from  $\mathcal{X}$  and  $\widehat{\mathcal{X}}$  that occur in  $\psi$ .

**Definition 2 (Classical Firing)** A solution  $\sigma$  to constraint  $\psi$  which satisfies the meta-flow axiom is called a classical firing. That is,  $\sigma$  is a classical firing for  $\psi$  if and only if  $\sigma, \mathcal{I} \models_C \psi \wedge FA(fv(\psi))$ . ◁

**Example 1** Recall the example from § 3.1. We define the constraints for each primitive in Table 1. The client does not impose any constraints on the input data (`tt`), and the  $FIFO_1$  primitive is empty so its constraints only say that no data can be output. *UserAppr* is an external predicate symbol, which must be resolved using external interaction (See § 7.1). Later we extend some of these constraints to capture the notion of state and interaction (see Table 2). The behaviour of the full system is given by the firings for the conjunction of all constraints.



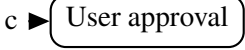
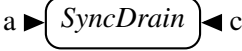
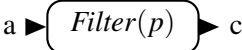

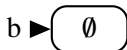
Primitive	Constraint
	$\psi_1 = a \rightarrow (\hat{a} = d_1 \vee \hat{a} = d_2 \vee \hat{a} = d_3)$
	$\psi_2 = \text{tt}$
	$\psi_3 = c \rightarrow \text{UserAppr}(\hat{c})$
	$\psi_4 = a \leftrightarrow c$
	$\psi_5 = c \rightarrow a \wedge c \rightarrow (p(\hat{c}) \wedge \hat{a} = \hat{c})$ $\wedge (a \wedge p(\hat{a})) \rightarrow c$
	$\psi_6 = \neg b$
	$\psi_7 = \neg b$

Table 1: List of primitives and their associated constraints, where  $d_1, d_2, d_3 \in \mathcal{Data}$ .

Consider the  $Filter(p)$  primitive in Table 1. The flow axiom is given by  $FA(a, c)$ . The constraint  $c \rightarrow a$  can be read as: if there is data flowing on  $c$ , then there must also be data flowing on  $a$ . The second part,  $c \rightarrow (p(\widehat{c}) \wedge \widehat{a} = \widehat{c})$ , says that when there is data flowing on  $c$  its value must validate the predicate  $p$ , and data flowing in  $a$  and  $c$  must be the same. Finally, the third part  $(a \wedge p(\widehat{a})) \rightarrow c$  states that data that validates the predicate  $p$  cannot be lost, i.e., flow on  $a$  but not on  $c$ . A classical firing for the interpretation  $\mathcal{I}$  is  $\{a \mapsto \text{tt}, c \mapsto \text{tt}, \widehat{a} \mapsto d, \widehat{c} \mapsto d\}$  whenever  $d \in \mathcal{D}ata$  is such that  $p(d) \mapsto \text{tt} \in \mathcal{I}$ . The assignment  $\{c \mapsto \text{tt}, \widehat{c} \mapsto \text{NO-FLOW}\}$  is not a classical firing because it violates the flow axiom, and because it is not a total map (it refers to neither  $a$  nor  $\widehat{a}$ ).

## 4 Partiality

The first step towards increasing the amount of available concurrency and the scalability of our approach is to make the logic partial. This means that solutions no longer need to be total, so for some  $x \in \mathcal{X}$  or  $\widehat{x} \in \widehat{\mathcal{X}}$ ,  $\sigma(x)$  or  $\sigma(\widehat{x})$  may not be defined. In addition, we drop the NO-FLOW value and so  $\sigma(\widehat{x})$  may either map to a value from  $\mathcal{D}ata$  or be undefined. The semantics is defined by a satisfaction relation  $\models_P$  and a dissatisfaction relation  $\models_P^{\perp}$  defined below. These state, respectively, when a formula is definite true or definitely false. Partiality is introduced in either the clause for  $x$  or for  $p(t_1, \dots, t_n)$ , whenever some variable is not defined in  $\sigma$ . An assignment  $\sigma$  now is a *partial* map from synchronisation variables to  $\{\text{tt}, \text{ff}\}$ , and from data flow variables to  $\mathcal{D}ata$ . Similarly, an interpretation  $\mathcal{I}$  is now an arity-indexed family of partial map from  $\mathcal{P}_n \times \mathcal{T}^n$  to  $\{\text{tt}, \text{ff}\}$ , where  $\mathcal{T}$  is the set of all possible ground terms, and  $Val_{\sigma}(f(t_1, \dots, t_n)) = \perp$  whenever  $Val_{\sigma}(t_i) = \perp$ , for some  $i \in 1..n$ . We use  $\perp$  to indicate when such a map is undefined.

### Definition 3 (Partial Satisfaction)

$\sigma, \mathcal{I} \models_P \text{tt}$	<i>always</i>
$\sigma, \mathcal{I} \models_P x$	<i>iff</i> $\sigma(x) = \text{tt}$
$\sigma, \mathcal{I} \models_P \psi_1 \wedge \psi_2$	<i>iff</i> $\sigma, \mathcal{I} \models_P \psi_1$ and $\sigma, \mathcal{I} \models_P \psi_2$
$\sigma, \mathcal{I} \models_P \neg \psi$	<i>iff</i> $\sigma, \mathcal{I} \not\models_P \psi$
$\sigma, \mathcal{I} \models_P p(t_1, \dots, t_n)$	<i>iff</i> $p(Val_{\sigma}(t_1), \dots, Val_{\sigma}(t_n)) \mapsto \text{tt} \in \mathcal{I}$
$\sigma, \mathcal{I} \models_P^{\perp} x$	<i>iff</i> $\sigma(x) = \text{ff}$
$\sigma, \mathcal{I} \models_P^{\perp} \psi_1 \wedge \psi_2$	<i>iff</i> $\sigma, \mathcal{I} \models_P^{\perp} \psi_1$ or $\sigma, \mathcal{I} \models_P^{\perp} \psi_2$
$\sigma, \mathcal{I} \models_P^{\perp} \neg \psi$	<i>iff</i> $\sigma, \mathcal{I} \models_P \psi$
$\sigma, \mathcal{I} \models_P^{\perp} p(t_1, \dots, t_n)$	<i>iff</i> $p(Val_{\sigma}(t_1), \dots, Val_{\sigma}(t_n)) \mapsto \text{ff} \in \mathcal{I}$

◁

**Lemma 1** *If  $\sigma$  and  $\mathcal{I}$  are total, then either  $\sigma, \mathcal{I} \models_P \psi$  or  $\sigma, \mathcal{I} \models_P^{\perp} \psi$ , but it is never undefined.*

We need to adapt the flow axiom as it refers explicitly to NO-FLOW, which is no longer available. The obvious change would be to replace NO-FLOW by partiality, giving (semantically)  $\sigma(x) = \text{ff} \iff \sigma(\widehat{x}) = \perp$ . But we can do better, permitting  $\sigma(x) = \perp$  to also represent no data flow. In addition, it is feasible that  $\sigma(x) = \text{tt}$  with  $\sigma(\widehat{x}) = \perp$  is a valid combination, to cover the case where the actual value of the data does not matter. Together, these give the following *meta-flow axiom*, which is a semantic and not syntactic condition.

**Axiom 2 (Meta-Flow Axiom)** *An assignment  $\sigma$  obeys the meta-flow axiom whenever for all  $x \in \mathcal{X}$ :*

$$\sigma(\widehat{x}) \neq \perp \implies \sigma(x) = \text{tt}$$

Write  $MFA(\sigma)$  whenever  $\sigma$  obeys the meta-flow axiom.

The following table gives all solutions to the meta-flow axiom:

	possible				forbidden	
$x$	tt	tt	ff	$\perp$	ff	$\perp$
$\widehat{x}$	$d$	$\perp$	$\perp$	$\perp$	$d$	$d$

For comparison, the following table gives the solutions for the flow axiom:

	possible		forbidden	
$x$	tt	ff	tt	ff
$\widehat{x}$	$d$	NO-FLOW	NO-FLOW	$d$

**Definition 4 (Partial Firing)** *A partial solution  $\sigma$  to a constraint  $\psi$  that satisfies the meta-flow axiom is called a partial firing. That is, whenever  $\sigma, \mathcal{I} \models_P \psi$  and  $MFA(\sigma)$ .  $\triangleleft$*

Consider again the constraints of the  $Filter(p)$  primitive in Table 1. A possible firing for it is  $\{a \mapsto \text{tt}, c \mapsto \text{ff}, \widehat{a} \mapsto d, \widehat{c} \mapsto \text{NO-FLOW}\}$  where  $d \in \mathcal{Data}$  does not validate the predicate  $p$ . The equivalent partial solution can be obtained by replacing NO-FLOW by  $\perp$ . Therefore,  $\{a \mapsto \text{tt}, c \mapsto \text{ff}, \widehat{a} \mapsto d\}$  is also a partial firing, whenever  $p(d)$  does not hold. Note also that  $\{c \mapsto \text{ff}\}, \mathcal{I} \models_P a \rightarrow c$  holds in the partial setting, yet  $\{c \mapsto \text{ff}\}, \mathcal{I} \models_C a \rightarrow c$  does not hold in the classical setting, because the classical satisfaction requires the solutions to be total mappings of all variables involved.

#### 4.1 Embeddings: Classical $\leftrightarrow$ Partial

We can move from an explicit representation of no-flow, namely  $\sigma(x) = \text{ff}$  and  $\sigma(\widehat{x}) = \text{NO-FLOW}$ , to an implicit representation using partiality, namely either  $\sigma(\widehat{x}) = \perp$  and  $\sigma(x) = \perp$  or  $\sigma(x) = \text{ff}$ , which means that the constraint solver need not find a value for  $x$  or  $\widehat{x}$ .

**Lemma 2 (Classical to Partial)** *Let  $\psi$  be a constraint where NO-FLOW does not occur in  $\psi$ , and  $\sigma$  be an assignment where  $\text{dom}(\sigma) = \text{fv}(\psi)$ . We write  $\mathcal{I}^\circ$  to represent the interpretation obtained by replacing in  $\mathcal{I}$  the constant NO-FLOW by  $\perp$ . If  $\sigma$  is a classical firing for  $\psi$  and the interpretation  $\mathcal{I}$ , then  $\sigma^\circ$  is a partial firing for  $\psi$  and the interpretation  $\mathcal{I}^\circ$ , where  $\sigma^\circ$  is defined as follows:*

$$\begin{aligned} \sigma^\circ(x) &= \sigma(x) \\ \sigma^\circ(\widehat{x}) &= \begin{cases} \perp, & \text{if } \sigma(\widehat{x}) = \text{NO-FLOW} \\ \sigma(\widehat{x}), & \text{otherwise} \end{cases} \end{aligned}$$

*Proof.* Assume that  $\sigma, \mathcal{I} \models_C \psi \wedge FA(\text{fv}(\psi))$ . Then (1)  $\sigma, \mathcal{I} \models_C \psi$  and (2)  $\sigma, \mathcal{I} \models_C FA(\text{fv}(\psi))$ . It can be seen by straightforward induction that (1) implies that  $\sigma^\circ, \mathcal{I}^\circ \models_P \psi$ , because  $\mathcal{I}^\circ$  maps the same values than  $\mathcal{I}$  after replacing NO-FLOW by  $\perp$ , and  $\sigma^\circ$  is defined for all free synchronisation variables. Since (2) holds for every  $x \in \text{fv}(\psi)$  and  $\text{dom}(\sigma^\circ) \cap \mathcal{X} = \text{dom}(\sigma) \cap \mathcal{X} = \text{fv}(\psi) \cap \mathcal{X}$ , we can safely conclude that  $MFA(\sigma^\circ)$ : when  $\sigma^\circ(x) \neq \text{tt}$  then  $\sigma^\circ(x) = \sigma(x) = \text{ff}$ , implying by the flow axiom that  $\sigma^\circ(\widehat{x}) = \text{NO-FLOW}$ , where we conclude that  $\sigma^\circ(\widehat{x}) = \perp$  (and the meta-flow axiom holds).  $\square$



**Lemma 3 (Partial to Classical)** *If  $\sigma$  is a partial firing for  $\psi$  and for a total interpretation  $\mathcal{I}$ , then  $\sigma^\dagger$  is a classical firing for  $\psi$  and for an interpretation  $\mathcal{I}^\dagger$ , where  $\mathcal{I}^\dagger$  results from replacing  $\perp$  by NO-FLOW in  $\mathcal{I}$ , and  $\sigma^\dagger$  is defined as follows:*

$$\begin{aligned} \sigma^\dagger(x) &= \sigma(x), & \text{if } \sigma(x) \neq \perp \\ \sigma^\dagger(x) &= \text{ff}, & \text{if } \sigma(x) = \perp \\ \\ \sigma^\dagger(\hat{x}) &= \sigma(\hat{x}), & \text{if } \sigma(\hat{x}) \neq \perp \\ \sigma^\dagger(\hat{x}) &= \text{NO-FLOW}, & \text{if } \sigma(\hat{x}) = \perp \text{ and } \sigma(x) \neq \text{tt} \\ \sigma^\dagger(\hat{x}) &= 42, & \text{if } \sigma(\hat{x}) = \perp \text{ and } \sigma(x) = \text{tt} \end{aligned}$$

*Proof.* Assume that (1)  $\sigma, \mathcal{I} \models_P \psi$  and (2)  $MFA(\sigma)$ . Note that  $\sigma^\dagger$  is total,  $\mathcal{I}^\dagger$  is total, and  $\sigma \subseteq \sigma^\dagger$ . It can be seen by Lemma 1 that  $\sigma \subseteq \sigma^\dagger$  and (1) imply  $\sigma^\dagger, \mathcal{I}^\dagger \models_C \psi$ . We show that  $\sigma^\dagger, \mathcal{I}^\dagger \models_C FA(\text{fv}(\psi))$  by considering all possible cases. (i) If  $\sigma(\hat{x}) \neq \perp$  then  $\sigma^\dagger(\hat{x}) = \sigma(\hat{x})$ , and from (2) we conclude that  $\sigma^\dagger(x) = \sigma(x) = \text{tt}$ . (ii) If  $\sigma(\hat{x}) = \perp$  and  $\sigma(x) = \text{tt}$ , then  $\sigma^\dagger(\hat{x}) = 42$  and  $\sigma^\dagger(x) = \text{tt}$ . (iii) If  $\sigma(\hat{x}) = \perp$  and  $\sigma(x) \neq \text{tt}$ , then  $\sigma^\dagger(x) = \text{ff}$  and  $\sigma^\dagger(\hat{x}) = \text{NO-FLOW}$ .  $\square$

## 4.2 Inexpressibility of Meta-Flow Axiom

Unfortunately, the meta-flow axiom is not expressible in partial logic. A consequence of this is that if partial logic is used as the constraint language, solutions may be found which do not satisfy this axiom; such solutions subsequently need to be filtered after performing constraint satisfaction, which clearly is not ideal, as the constraint engine would need to continue to find a real solution, having wasted time finding this non-solution.

The following lemma will help prove that the meta-flow axiom is not expressible.

**Lemma 4** *If  $\sigma, \mathcal{I} \models_P \psi$  and  $\sigma \subseteq \sigma'$ , then  $\sigma', \mathcal{I} \models_P \psi$ .*

*Proof.* By straightforward induction on  $\psi$ .  $\square$

**Lemma 5** *No formula  $\psi$  exists such that  $\sigma, \mathcal{I} \models_P \psi$  if and only if  $MFA(\sigma)$ .*

*Proof.* Assume that  $\psi_{MFA}$  is such a formula over variables  $\{x, \hat{x}\}$ . Then for  $\sigma = \{x \mapsto \text{ff}\}$  we have that  $\sigma, \mathcal{I} \models_P \psi_{MFA}$ . Now  $\sigma \subseteq \sigma' = \{x \mapsto \text{ff}, \hat{x} \mapsto 42\}$ . Hence by Lemma 4, we have that  $\sigma', \mathcal{I} \models_P \psi_{MFA}$ . But  $\sigma'$  does not satisfy the meta-flow axiom. Contradiction.  $\square$

## 4.3 Simple Logic

Using partial logic as the basis of a coordination engine is not ideal, as constraint satisfaction for this logic could find solutions which do not satisfy the meta-flow axiom (due to Lemma 5). Such solutions would need to be filtered in a post-processing phase, resulting in an undesirable overhead.

We resolve this problem by modifying the semantics so that only certain ‘minimal’ solutions are found. These solutions define only the necessary variables—which has the consequence that the constraint solver needs only to satisfy variables mentioned in the (relevant branch of a) constraint. We extend also the syntax of formulæ by distinguishing two kinds of conjunctions. The *overlapping conjunction* ( $\wedge$ ) of two constraints accepts two compatible solutions and joins them together, while an *additive conjunction* ( $\bar{\wedge}$ ) accepts only solutions which satisfy both constraints. Both kinds of conjunction are present, firstly, to talk about the joining of solutions for (partially) independent parts of a connector (overlapping conjunction), and to enforce overarching constraints, such as the flow axiom (additive conjunction). The semantics for the logic is formalised in Definition 5. In this logic, the meta-flow axiom is expressible.

**Definition 5 (Simple satisfaction)** We define inductively a simple satisfaction relation  $\sigma, \mathcal{I} \models_S \psi$  and a simple dissatisfaction relation  $\sigma, \mathcal{I} \models_S \neg \psi$ , where the assignment  $\sigma$  and the interpretation  $\mathcal{I}$  may be partial.

$\emptyset, \mathcal{I} \models_S \text{tt}$	<i>always</i>
$\{[x \mapsto \text{tt}]\}, \mathcal{I} \models_S x$	<i>always</i>
$\sigma_1 \cup \sigma_2, \mathcal{I} \models_S \psi_1 \wedge \psi_2$	<i>iff</i> $\sigma_1, \mathcal{I} \models_S \psi_1$ and $\sigma_2, \mathcal{I} \models_S \psi_2$ and $\sigma_1 \frown \sigma_2$
$\sigma, \mathcal{I} \models_S \psi_1 \bar{\wedge} \psi_2$	<i>iff</i> $\sigma, \mathcal{I} \models_S \psi_1$ and $\sigma, \mathcal{I} \models_S \psi_2$
$\sigma, \mathcal{I} \models_S \neg \psi$	<i>iff</i> $\sigma, \mathcal{I} \models_S \psi$
$\sigma, \mathcal{I} \models_S p(t_1, \dots, t_n)$	<i>iff</i> $p(\text{Val}_\sigma(t_1), \dots, \text{Val}_\sigma(t_n)) \mapsto \text{tt} \in \mathcal{I}$ and $\text{dom}(\sigma) = \text{fv}(p(t_1, \dots, t_n))$
$\{[x \mapsto \text{ff}]\}, \mathcal{I} \models_S x$	<i>always</i>
$\sigma, \mathcal{I} \models_S \psi_1 \wedge \psi_2$	<i>iff</i> for all $\sigma_1, \sigma_2$ s.t. $\sigma_1 \frown \sigma_2$ and $\sigma = \sigma_1 \cup \sigma_2$ we have $\sigma_1, \mathcal{I} \models_S \psi_1$ or $\sigma_2, \mathcal{I} \models_S \psi_2$
$\sigma, \mathcal{I} \models_S \psi_1 \bar{\wedge} \psi_2$	<i>iff</i> $\sigma, \mathcal{I} \models_S \psi_1$ or $\sigma, \mathcal{I} \models_S \psi_2$
$\sigma, \mathcal{I} \models_S \neg \psi$	<i>iff</i> $\sigma, \mathcal{I} \models_S \psi$
$\sigma, \mathcal{I} \models_S p(t_1, \dots, t_n)$	<i>iff</i> $p(\text{Val}_\sigma(t_1), \dots, \text{Val}_\sigma(t_n)) \mapsto \text{ff} \in \mathcal{I}$ and $\text{dom}(\sigma) = \text{fv}(p(t_1, \dots, t_n))$

$$\text{where } \sigma_1 \frown \sigma_2 \hat{=} \forall x \in \text{dom}(\sigma_1) \cap \text{dom}(\sigma_2). \sigma_1(x) = \sigma_2(x)$$

◁

The additive conjunction  $\bar{\wedge}$  of  $\psi_1$  and  $\psi_2$  is satisfied by  $\sigma$  if  $\sigma$  satisfies both  $\psi_1$  and  $\psi_2$ . The overlapping conjunction  $\wedge$  is more relaxed, and simply merges any pair of solutions for  $\psi_1$  and  $\psi_2$  that do not contradict each other. For the constraints of the primitives, the conjunctions that appear in a positive position are regarded as overlapping conjunctions ( $\wedge$ ), while the conjunctions that appear in a negative position are regarded as additive conjunctions ( $\bar{\wedge}$ ).<sup>2</sup> As a consequence, the rule for  $\sigma, \mathcal{I} \models_S \psi_1 \bar{\wedge} \psi_2$  is only used when applying the flow axiom, as we will soon see, and the rule for  $\sigma, \mathcal{I} \models_S \psi_1 \wedge \psi_2$  is present mainly for the completeness of the definition.

**Notation** In the following we write  $\psi^S$  to represent the constraints obtained by replacing all conjunctions in  $\psi$  in negative positions by  $\bar{\wedge}$ , and we write  $\psi^P$  to represent the constraints obtained by replacing all occurrences of  $\bar{\wedge}$  in  $\psi$  by  $\wedge$ . We also encode  $\psi_1 \vee \psi_2$  as  $\neg(\neg\psi_1 \bar{\wedge} \neg\psi_2)$ .

When specifying constraints in simple logic, we never use  $-\vee-$  in a positive position, which would correspond to  $\neg(\neg-\wedge\neg)$ , as this means satisfying the clause  $\sigma, \mathcal{I} \models_S \psi_1 \wedge \psi_2$  in order to find a given assignment. Therefore we do not require the use of universal quantification over solution sets. In the partial satisfaction relation, we define how to verify that a given pair  $\sigma, \mathcal{I}$  satisfies a constraint. The simple satisfaction relation aims at *constructing*  $\sigma$  such that the pair  $\sigma, \mathcal{I}$  satisfies the constraints. Assuming the universal quantifier is never used, we believe that the simple satisfaction relation describes a constructive process to obtain a solution that is not more complex than searching for a solution in partial logic. Note that we still lack experimental verification of this intuition.

The following axiom is the syntactic counterpart of the meta-flow axiom, modified slightly to be laxer about what it considers to be a solution (namely allowing data flow variable to be satisfied, without requiring that the corresponding synchronisation variable are defined).

<sup>2</sup>A positive position is inside the scope of an even number of negations, and a negative position is inside the scope of an odd number of negations. For example, in  $(\neg(a \wedge \neg b)) \wedge c$ ,  $a$  is in a negative position, while  $b$  and  $c$  are in a positive position.

**Definition 6 (Simple Flow Axiom)**

$$SFA(x) \hat{=} \text{tt} \vee x \vee \neg x \vee (x \wedge \hat{x} = \hat{x}) \vee \hat{x} = \hat{x} \quad (\text{simple flow axiom})$$

◁

We write  $SFA(X)$  for the conjunction  $\bigwedge_{x \in X} SFA(x)$ . We also write  $SFA(\psi)$  as a shorthand for  $SFA(\text{fv}(\psi))$ .

**Lemma 6**  $\sigma, \mathcal{S} \models_S SFA(x)$  if and only if  $\sigma^P$  satisfies the meta-flow axiom, where  $\sigma^P$  extends  $\sigma$  as follows:

$$\sigma^P = \sigma \cup \{x \mapsto \text{tt} \mid \hat{x} \in \text{dom}(\sigma)\}$$

*Proof.* We have  $\emptyset, \mathcal{S} \models_S \text{tt}$ ;  $\{x \mapsto \text{tt}\}, \mathcal{S} \models_S x$ ;  $\{x \mapsto \text{ff}\}, \mathcal{S} \models_S \neg x$ ;  $\{x \mapsto \text{tt}, \hat{x} \mapsto t\}, \mathcal{S} \models_S x \wedge \hat{x} = \hat{x}$ ; and  $\{\hat{x} \mapsto t\}, \mathcal{S} \models_S \hat{x} = \hat{x}$ , for an arbitrary ground term  $t$ , and no other  $\sigma$ . Extending  $\sigma$  to  $\sigma^P$  we obtain precisely the solutions to the meta-flow axiom.  $\square$

Observe that simple logic clearly does not preserve classical or even partial equivalences, as  $\text{tt} \vee x \vee \neg x \vee (x \wedge \hat{x} = \hat{x}) \vee \hat{x} = \hat{x} \equiv_C \text{tt}$ , classically, but this is not the case in simple logic.

**Definition 7 (Simple Firing)**

An assignment  $\sigma$  is called a simple firing whenever  $\sigma, \mathcal{S} \models_S \psi \bar{\wedge} SFA(\psi)$ . ◁

Note that the simple flow axiom differs from the meta-flow axiom because it also accepts solutions where  $\hat{x}$  is defined, but  $x$  is not. This is because the simple flow axiom is designed to filter from a set of minimal solutions (i.e., solutions in the simple logic), while the meta-flow axiom is designed to filter good solutions from all solutions the constraint engine finds, namely, the ones that include additional assignments to make the flow axiom hold. As a consequence, the assignment  $\{\hat{a} \mapsto d\}$  is a simple firing for the formula  $\hat{a} = d$ , and the assignment  $\{\hat{a} \mapsto d, a \mapsto \text{tt}\}$  is a partial firing for the same formula, but not the other way around.

With simple logic we can check a formula to ensure that all of its solutions satisfy the meta-flow axiom. This means that we do not need to filter solutions to such a constraint. Furthermore, as the simple flow axiom is preserved through composition ( $\bar{\wedge}$ ), we are guaranteed to have simple firings without having to perform a post hoc filter phase.

Note that implication in simple logic does not have the exact same meaning as in the other logics.  $c \rightarrow a$ , whenever in a positive position, is regarded in simple logic as  $\neg(c \bar{\wedge} \neg a)$ , which has only two firings:  $\{c \mapsto \text{ff}\}$  and  $\{a \mapsto \text{tt}\}$ . The union of these two firings is not a simple firing because it is not “minimal enough”. That is, the resulting union is not satisfied by the simple satisfaction relation because it contains too many elements. Recall the constraints of the  $Filter(p)$  in Table 1, and let  $d$  be such that  $p(d)$  does not hold. The assignment  $\{a \mapsto \text{tt}, c \mapsto \text{ff}, \hat{a} \mapsto d\}$  is both a partial and a simple firing. However,  $\{z \mapsto \text{tt}, a \mapsto \text{tt}, c \mapsto \text{ff}, \hat{a} \mapsto d\}$  is also a partial firing but not a simple firing, since  $z \notin \text{fv}(c \rightarrow a)$ , therefore the firing is not minimal enough.

**Lemma 7** Let  $\psi_1$  and  $\psi_2$  be constraints defined for the simple logic. Then

$$(\psi_1 \bar{\wedge} SFA(\psi_1)) \wedge (\psi_2 \bar{\wedge} SFA(\psi_2)) \equiv (\psi_1 \wedge \psi_2) \bar{\wedge} (SFA(\psi_1 \wedge \psi_2))$$

The equivalence between the left and the right hand formulae mean that they have the same solutions according to the simple satisfaction.

*Proof.* Let  $\text{sols}(\psi)$  denote the set of solutions of  $\psi$  according to the simple satisfaction relation, and let  $\text{sols}(\psi_1) = S_1$ ,  $\text{sols}(\psi_2) = S_2$ ,  $\text{sols}(SFA(\psi_1)) = S_{F1}$ , and  $\text{sols}(SFA(\psi_2)) = S_{F2}$ . The proof follows from the expansion of the definition of simple satisfaction.

$$\begin{aligned}
& \text{sols}((\psi_1 \bar{\wedge} SFA(\psi_1)) \wedge (\psi_2 \bar{\wedge} SFA(\psi_2))) \\
&= \{ \sigma_1 \cup \sigma_2 \mid \sigma_1 \in S_1 \cap S_{F1}, \sigma_2 \in S_2 \cap S_{F2}, \sigma_1 \frown \sigma_2 \} \\
&= \{ \sigma_1 \cup \sigma_2 \mid \sigma_1 \in S_1, \sigma_1 \in S_{F1}, \sigma_2 \in S_2, \sigma_2 \in S_{F2}, \sigma_1 \frown \sigma_2 \} \\
&= \{ \sigma_1 \cup \sigma_2 \mid \sigma_1 \in S_1, \sigma_2 \in S_2, \sigma_1 \frown \sigma_2 \} \cap \{ \sigma_1 \cup \sigma_2 \mid \sigma_1 \in S_{F1}, \sigma_2 \in S_{F2}, \sigma_1 \frown \sigma_2 \} \\
&= \text{sols}(\psi_1 \wedge \psi_2) \cap \text{sols}(SFA(\psi_1) \wedge SFA(\psi_2)) \\
&= \text{sols}((\psi_1 \wedge \psi_2) \bar{\wedge} (SFA(\psi_1) \wedge SFA(\psi_2))) \\
&= \text{sols}((\psi_1 \wedge \psi_2) \bar{\wedge} (SFA(\psi_1 \wedge \psi_2)))
\end{aligned}$$

□

### Lemma 8 (Partial to Simple)

- If  $\sigma, \mathcal{I} \models_P \psi$  and  $MFA(\sigma)$ , then there exists  $\sigma^\ddagger$  such that  $\sigma^\ddagger \subseteq \sigma$  and  $\sigma^\ddagger, \mathcal{I} \models_S \psi^S \bar{\wedge} SFA(\psi)$ .
- If  $\sigma, \mathcal{I} \models_P \psi$  and  $MFA(\sigma)$ , then there exists  $\sigma^\ddagger$  such that  $\sigma^\ddagger \subseteq \sigma$  and  $\sigma^\ddagger, \mathcal{I} \models_S \psi^S \bar{\wedge} SFA(\psi)$ .

*Proof.* Proof is by straightforward induction on  $\psi$ . Note that  $\bar{\wedge}$  cannot occur in  $\psi$ , and in each step  $\sigma^\ddagger$  is guaranteed to exist and to obey the simple flow axiom:

**Case tt** —  $\sigma^\ddagger = \emptyset$ .

**Case x** —  $\sigma^\ddagger = \{x \mapsto \sigma(x)\}$ .

**Case  $\psi_1 \wedge \psi_2$**  — For the  $\models_P$  case: Assume that  $\sigma, \mathcal{I} \models_P \psi_1 \wedge \psi_2$  and  $MFA(\sigma)$ . Therefore  $\sigma, \mathcal{I} \models_P \psi_1$  and  $\sigma, \mathcal{I} \models_P \psi_2$ . By the induction hypothesis, we have  $\sigma_1^\ddagger \subseteq \sigma_1$  and  $\sigma_2^\ddagger \subseteq \sigma_2$  such that  $\sigma_1^\ddagger, \mathcal{I} \models_S \psi_1^S \bar{\wedge} SFA(\psi_1)$  and  $\sigma_2^\ddagger, \mathcal{I} \models_S \psi_2^S \bar{\wedge} SFA(\psi_2)$ . Clearly we have  $\sigma_1^\ddagger \frown \sigma_2^\ddagger$  and  $\sigma_1^\ddagger \cup \sigma_2^\ddagger \subseteq \sigma$ , and by Lemma 7 we conclude that  $\sigma_1^\ddagger \cup \sigma_2^\ddagger, \mathcal{I} \models_S (\psi_1 \wedge \psi_2)^S \bar{\wedge} SFA(\psi_1 \wedge \psi_2)$ .

For the  $\models_P$  case: Assume that  $\sigma, \mathcal{I} \models_P \psi_1 \wedge \psi_2$  and  $MFA(\sigma)$ . Therefore  $\sigma, \mathcal{I} \models_P \psi_1$  or  $\sigma, \mathcal{I} \models_P \psi_2$ . By the induction hypothesis, we have  $\sigma_1^\ddagger \subseteq \sigma_1$  and  $\sigma_2^\ddagger \subseteq \sigma_2$  such that  $\sigma_1^\ddagger, \mathcal{I} \models_S \psi_1^S \bar{\wedge} SFA(\psi_1)$  and  $\sigma_2^\ddagger, \mathcal{I} \models_S \psi_2^S \bar{\wedge} SFA(\psi_2)$ . Note that  $\wedge$  is in a negative position, therefore  $(\psi_1 \wedge \psi_2)^S = \psi_1 \bar{\wedge} \psi_2$ .

Clearly we have that when  $\sigma^\ddagger = \sigma_1^\ddagger$  or  $\sigma^\ddagger = \sigma_2^\ddagger$ ,  $\sigma^\ddagger, \mathcal{I} \models_S (\psi_1 \wedge \psi_2)^S \bar{\wedge} SFA(\psi_1 \wedge \psi_2)$  and  $\sigma^\ddagger \subseteq \sigma$ .

**Case  $\neg\psi$**  — by induction hypothesis.

**Case  $p(t_1, \dots, t_n)$**  — in both cases  $\sigma^\ddagger = \{v \mapsto \sigma(v) \mid v \in \text{fv}(p(t_1, \dots, t_n))\}$

□

**Lemma 9** If  $\sigma, \mathcal{I} \models_S \psi$ , then  $\sigma^P, \mathcal{I} \models_P \psi^P$ . If  $\sigma, \mathcal{I} \models_S \psi$ , then  $\sigma^P, \mathcal{I} \models_P \psi^P$ .

*Proof.* By straightforward induction on  $\psi$ .

□

**Lemma 10 (Simple to Partial)** If  $\sigma$  is a simple firing for  $\psi$ , then  $\sigma^P$  is a partial firing for  $\psi^P$ . Furthermore, for all  $\sigma'$  such that  $\sigma^P \subseteq \sigma'$  and  $\sigma'$  satisfies the meta-flow axiom,  $\sigma'$  is a partial firing for  $\psi^P$ .

*Proof.* Follows from Lemmas 4 and 9.

□

The key difference between simple and partial is that simple finds the kernel of a solution by examining only the relevant variables. All partial solutions can be reconstructed by filling in arbitrary values (satisfying the meta-flow axiom) for the unspecified variables. Note that the classical model is faithful to existing semantics of Reo. By shifting to a partial logic, we can model *pure synchronisation*, which is when a synchronisation variable is true and the corresponding data flow variable is  $\perp$ . In the simple logic, if the data flow variable is not mentioned, it will never be assigned a value, reflecting that there is no data flowing in the corresponding ports, i.e., it is a pure synchronisation port.

## 5 Locality

With simple logic, there is still a single set of constraints and thus it is not clear how to exploit this to extract any inherent concurrency nor is it clear how to partition the constraints to distribute them. Our motivation is to use (distributed) constraint satisfaction as the basis of a coordination language between geographically distributed components and services.

The local semantics is based on a *configuration* consisting of constraints partitioned into blocks of constraints, denoted by  $\Psi = \langle \psi_1 \rangle, \dots, \langle \psi_n \rangle$ , or simply  $\Psi = \langle \psi_i \rangle^{i \in 1..n}$ .

**Definition 8 (No-Flow Assignment)** *An assignment  $\sigma$  is called a no-flow assignment whenever  $\text{dom}(\sigma) \subseteq \mathcal{X}$  and for all  $x \in \text{dom}(\sigma)$  we have  $\sigma(x) = \text{ff}$ .*  $\triangleleft$

**Axiom 3 (No-Flow Axiom)** *We say that a constraint  $\psi$  obeys the no-flow axiom whenever there is some no-flow assignment  $\sigma$  with  $\text{dom}(\sigma) \subseteq \text{fv}(\psi) \cap \mathcal{X}$  such that  $\sigma, \mathcal{S} \models_S \psi$ .*

*A configuration  $\Psi = \langle \psi_i \rangle^{i \in 1..n}$  obeys the no-flow axiom iff each  $\psi_i$  obeys the no-flow axiom.*

From now on, we assume that all configurations satisfy the no-flow axiom.

**Definition 9 (Boundary)** *Given a configuration  $\Psi = \langle \psi_1 \rangle, \dots, \langle \psi_n \rangle$ , define  $\text{boundary}_\Psi(\langle \psi_i \rangle)$  as  $\text{fv}(\psi_i) \cap \text{fv}(\Psi_{-i})$ , where  $\Psi_{-i} = \langle \psi_1 \rangle, \dots, \langle \psi_{i-1} \rangle, \langle \psi_{i+1} \rangle, \dots, \langle \psi_n \rangle$ .*

*We drop the  $\Psi$  subscript from  $\text{boundary}_\Psi(-)$  when it is clear from the context.*  $\triangleleft$

**Definition 10 (Local Firing)** *Given a configuration  $\Psi = \langle \psi_1 \rangle, \dots, \langle \psi_n \rangle$ . We say that:*

- $\sigma$  is a local firing for a block  $\langle \psi_i \rangle$  if and only if  $\sigma$  is a simple firing for  $\psi_i$  and for all  $x \in \text{boundary}_\Psi(\langle \psi_i \rangle)$  we have  $\sigma^P(x) \neq \text{tt}$ —we call this the boundary condition.<sup>3</sup>
- $\sigma$  is a local firing for  $\Psi$  if and only if  $\sigma = \sigma_1 \cup \dots \cup \sigma_{m'}$  such that
  1.  $I_1, \dots, I_{m'}, \dots, I_m$  is a partition of  $\{1..n\}$ ;
  2.  $\varphi_i = \bigwedge_{j \in I_i} \psi_j$  where  $i \in 1..m$ ; and
  3.  $\sigma_i$  is a local firing for block  $\langle \varphi_i \rangle$  where  $i \in 1..m'$ .

The intuition behind this definition is:

1. a local firing can occur in some isolated block *or* the conjunction of some blocks *or* in independent (conjunctions of) blocks; and
2. within each block a simple firing occurs that makes the assumption that there is no-flow on its boundary ports.




Figure 2: Simple network of constraints: two competing data producers.

<sup>3</sup> $\sigma^P$  is defined in Lemma 6.

**Example 2** We introduce a small example in Figure 2 that we use to illustrate the definition of local firings. Let  $\phi_i$ , where  $i \in 1..6$ , be the constraints for  $CDG_1$ ,  $CDG_2$ ,  $Lossy_1$ ,  $Lossy_2$ ,  $Merger$ , and  $Client$ , respectively. We define  $\phi_3$  for  $Lossy_1$  and  $\phi_5$  for  $Merger$  as follows.

$$\begin{aligned}\phi_3 &= b \mapsto a \wedge b \mapsto (\widehat{a} = \widehat{b}) \\ \phi_5 &= e \leftrightarrow (b \vee d) \wedge \neg(b \wedge d) \wedge b \mapsto (\widehat{e} = \widehat{b}) \wedge d \mapsto (\widehat{e} = \widehat{d})\end{aligned}$$

The remaining constraints can be derived similarly.

The  $Lossy_1$  can arbitrary lose data flowing on  $a$ , or pass data from  $a$  to  $b$ ; and  $Merger$  can pass data either from  $b$  to  $e$  or from  $d$  to  $e$ . The configuration that captures the behaviour of the full system is given by  $\Psi_{merge} = \langle \phi_i \bar{\wedge} SFA(\phi_i) \rangle_{i \in 1..6}$ .

We present some simple firings for the  $Lossy$  primitive, and show which of these are also local firings for  $\Psi_{merge}$ , and we then describe some more complex local firings for  $\Psi_{merge}$ . We omit the formal proof that the conditions for local firings hold for these firings. Formula  $\phi_3$  can be written as  $\neg(b \bar{\wedge} \neg a) \wedge \neg(b \bar{\wedge} \neg(\widehat{a} = \widehat{b}))$ , and the boundary of  $\langle \phi_3 \rangle$  is  $\{a, \widehat{a}, b, \widehat{b}\}$ . Valid simple firings for  $\phi_3$  are  $\{b \mapsto \mathbf{ff}\}$ ,  $\{a \mapsto \mathbf{tt}, b \mapsto \mathbf{ff}\}$ , and  $\{a \mapsto \mathbf{tt}, \widehat{a} \mapsto v, \widehat{b} \mapsto v\}$ , for any possible data value  $v \in \mathcal{Data}$ . The only simple firing that is also a local firing is then  $\{b \mapsto \mathbf{ff}\}$ . This means that  $\{b \mapsto \mathbf{ff}\}$  is also a local firing of  $\Psi_{merge}$ .

It is also possible to show that the solution  $\sigma_{top}$  corresponding to the flow of some data value  $v$  from  $CDG_1$  to  $Client$  is a simple firing for  $\phi_1 \wedge \phi_3 \wedge \phi_5 \wedge \phi_6$ , and that the solution  $\sigma_{bottom}$  corresponding to data being sent from  $CDG_2$  to  $Lossy_2$  and being lost is a simple firing for  $\phi_2 \wedge \phi_4$ . More precisely, we define  $\sigma_{top} = \{a \mapsto \mathbf{tt}, b \mapsto \mathbf{tt}, d \mapsto \mathbf{ff}, e \mapsto \mathbf{tt}, \widehat{a} \mapsto v, \widehat{b} \mapsto v, \widehat{e} \mapsto v\}$  and  $\sigma_{bottom} = \{c \mapsto \mathbf{ff}, d \mapsto \mathbf{ff}\}$ . The boundary for both sets of primitives is just  $\{d\}$ . In the solutions  $\sigma_{top}$  and  $\sigma_{bottom}$  the value of  $d$  is never  $\mathbf{tt}$ , so the boundary conditions hold. Therefore  $\sigma_{top}$ ,  $\sigma_{bottom}$ , and  $\sigma_{top} \cup \sigma_{bottom}$  are also local firings of  $\Psi_{merge}$ .

The local semantics is based on the simple semantics under the no-flow axiom assumption. Thus, a simple firing can be trivially seen as a local solution, but a local firing needs to be extended to be seen as a simple firing. This extension corresponds exactly to the unfolding of the no-flow axiom for the blocks of constraints not involved in the local firing. The embedding between these two semantics is formalised below.

**Lemma 11 (Simple to Local)** *If  $\sigma$  is a simple firing for  $\psi$ , then  $\sigma$  is a local firing for  $\Psi = \langle \psi \rangle$ .*

*Proof.* As  $boundary_{\Psi}(\langle \psi \rangle) = \emptyset$ , a simple firing for  $\psi$  is also a local firing for  $\psi$ . □

**Lemma 12 (Local to Simple)** *Let  $\sigma$  be a local firing for  $\Psi = \langle \psi_1 \rangle, \dots, \langle \psi_n \rangle$ , then there exists a  $\sigma^*$  such that (1)  $\sigma \subseteq \sigma^*$ , (2) for all  $x \in \text{dom}(\sigma^*) \setminus \text{dom}(\sigma)$  we have  $\sigma^*(x) = \mathbf{ff}$ , and (3)  $\sigma^*$  is a simple firing for  $\bigwedge_{i \in 1..n} \psi_i$ .*

*Proof.* Assume that  $\sigma$  is a local firing for  $\Psi = \langle \psi_1 \rangle, \dots, \langle \psi_n \rangle$ . Without loss of generality, we can assume that  $\sigma = \sigma_1 \cup \dots \cup \sigma_m$  where (1)  $m \leq n$  and for each  $k \in 1..m$  we have that  $\sigma_k$  is a local firing for  $\langle \psi_k \rangle$ . From the no-flow axiom, we can have a no-flow assignment  $\sigma_j^*$  for each  $j \in m+1..n$  such that  $\sigma_j^*, \mathcal{S} \models_S \psi_j$ . From the boundary condition, we can infer that for each  $\sigma_k$  and  $\sigma_j^*$  the condition  $\sigma_k \frown \sigma_j^*$  holds. Thus, for  $\sigma^* = \sigma \cup \bigcup_j \sigma_j^*$ , we have  $\sigma^*, \mathcal{S} \models_S \bigwedge_{i \in 1..n} \psi_i$ . □

## 6 State


### 6.1 State-machine

We follow the encoding of stateful connectors presented by Clarke et al. [5]. To encode stateful connectors we add  $state_p$  and  $state'_p$  to the term variables, for each  $p \in P$  corresponding to a stateful primitive. A state machine with states  $q_1, \dots, q_n$  is encoded as a formula of the form:

$$\Psi = state_p = q_1 \rightarrow \Psi_1 \wedge \dots \wedge state_p = q_n \rightarrow \Psi_n$$

where  $\Psi_1, \dots, \Psi_n$  are constraints representing the transitions from each state. For each firing  $\sigma$ , the value of  $\sigma(state'_p)$  determines how the connector evolves, giving the value of the next state.

We illustrate this encoding by an example, presenting the constraints encoding the state machine of a  $FIFO_1$  buffer from Table 1. The state can be either empty or  $full(d)$ , where  $d \in \mathcal{D}ata$ , and empty and full are function symbols in  $\mathcal{F}$ . We then define *fifo*-constraint to be  $state_{fifo} = empty \rightarrow \Psi_e \wedge state_{fifo} = full(d) \rightarrow \Psi_f$ , where  $\Psi_e$  and  $\Psi_f$  are the upper and lower labels of the following diagram, respectively:



To complete the encoding, we add a formula describing the present state to the mix. In the example, the formula  $state_{fifo} = empty$  records the fact that the  $FIFO_1$  is in the empty state, whereas  $state_{fifo} = full(d)$  records that it is in the full state, containing data  $d$ . The full constraint for the  $FIFO_1$  primitive is now (refining the constraints in Table 1):

$$state_{fifo} = empty \rightarrow \Psi_e \wedge state_{fifo} = full(d) \rightarrow \Psi_f \wedge state_{fifo} = empty,$$

### 6.2 Constraint satisfaction-based engine

A constraint satisfaction-based engine holds a configuration with the current set of constraints and operates in *rounds*, each of which consists of a *solve* phase and an *update* phase, which uses the firing to update the constraints and to model the transition to a new state. This is depicted in Figure 3.




Figure 3: Phases of the constraint satisfaction-based engine.

Each block of the configuration is a conjunction of two constraints  $\langle \rho \wedge \varepsilon \rangle$ , where  $\rho$  is persistent and  $\varepsilon$  is ephemeral. Persistent constraints are eternally true, and can be either (normal) stateless constraints, stateful constraints, or the conjunction of persistent constraints. Ephemeral constraints describe

the present state of the stateful constraints. Configurations are updated at each round. Let  $\mathcal{I}$  be an interpretation and, for each  $i \in 1..n$ , let  $P_i$  be a (possibly empty) set of names of stateful constraints. A full round can be represented as follows, where the superscript indicates the round number:

$$\langle \rho_i \wedge \varepsilon_i^m \rangle^{i \in 1..n} \xrightarrow{\text{solve}} \langle \sigma^m \rangle \xrightarrow{\text{update}} \langle \rho_i \wedge \varepsilon_i^{m+1} \rangle^{i \in 1..n}$$

satisfying the following:

$$\sigma^m, \mathcal{I} \models_S \bigwedge_{i \in 1..n} \rho_i \wedge \varepsilon_i^m \quad (\text{solve})$$

$$\varepsilon_i^{m+1} \equiv \bigwedge \{ \text{state}_p = \sigma^m(\text{state}'_p) \mid p \in P_i \text{ and } \sigma^m(\text{state}'_p) \neq \perp \} \wedge \bigwedge \{ \varepsilon_i^m \mid p \in P_i \text{ and } \sigma^m(\text{state}'_p) = \perp \} \quad (\text{update})$$

In round  $m$ , the solve phase finds a solution  $\sigma^m$ , and the update phase replaces the definition of  $\varepsilon^m$  by  $\varepsilon^{m+1}$  for round  $m$ , whenever the variable  $\text{state}'$  is defined. A correctness result of this approach with respect to Reo, for the classical semantics, has been presented by Clarke et al. [5]. The authors use the constraint automata semantics for Reo [3] as the reference for comparison. The present approach adapts the previous one by accounting for partial solutions of the constraints, which means that only some of the state variables are updated.

## 7 Interaction

We now extend the model with means for external interaction.

### 7.1 External functions, predicates and constraints

We defined in § 3.2 a core syntax for logic formulæ, extended with state variables in § 6. Satisfaction of formulæ is defined with respect to an assignment  $\sigma$  defining the values of variables, and an interpretation  $\mathcal{I}$  giving meaning to predicates. We now extend the syntax of the logic and the definition of interpretation  $\mathcal{I}$ , by introducing new symbols whose interpretation is also given by  $\mathcal{I}$ . These symbols are *external predicate* symbols  $\mathbf{p} \in \mathbf{P}$ , *external function* symbols  $\mathbf{f} \in \mathbf{F}$ , and *external constraints*  $\mathbf{c} \in \mathbf{C}$ . We also introduce *communication variables*  $\mathbf{k} \in \mathbf{K}$  whose value in the solution of a round can be communicated to the outside world. Formulæ are now given by the following syntax:

$$\begin{aligned} \psi &::= \text{tt} \mid x \mid \psi_1 \wedge \psi_2 \mid \neg \psi \mid p(\bar{t}) \mid \mathbf{p}(\bar{t}) \mid \mathbf{c}(\bar{\psi}, \bar{t}) \\ t &::= \hat{x} \mid \text{state}_p \mid \text{state}'_p \mid \mathbf{k} \mid f(\bar{t}) \mid \mathbf{f}(\bar{t}) \end{aligned}$$

Use  $\bar{t}$  as a shorthand for  $t_1, \dots, t_n$ . We extend the definition of interpretation to be an arity-indexed family of *partial* map from  $P_n \times \mathcal{T}^n$  to  $\{\text{tt}, \text{ff}\}$ ; from  $\mathbf{P}_n \times \mathcal{T}^n$  to  $\{\text{tt}, \text{ff}\}$ ; from  $\mathbf{F}_n \times \mathcal{T}^n$  to ground terms; and from  $\mathbf{C}$  to a term with  $l$  formulæ parameters and  $k$  term parameters.

We also extend the  $\text{Val}_{\sigma, \mathcal{I}}$  function, which is now parameterized on  $\sigma$  and  $\mathcal{I}$ . This function replaces variables  $v$  by  $\sigma(v)$  and  $\mathbf{f}(t_1, \dots, t_n)$  by  $\mathcal{I}(\mathbf{f}, \text{Val}_{\sigma, \mathcal{I}}, \dots, \text{Val}_{\sigma, \mathcal{I}})$ , or is undefined if any component is undefined.



The extension of the syntax of the logic requires the addition of two new (dis)satisfaction rules:

$$\begin{array}{ll}
\sigma, \mathcal{I} \models_S \mathbf{p}(t_1, \dots, t_n) & \text{iff } \mathbf{p}(\text{Val}_{\sigma, \mathcal{I}}(t_1), \dots, \text{Val}_{\sigma, \mathcal{I}}(t_n)) \mapsto \mathbf{tt} \in \mathcal{I} \\
& \text{and } \text{dom}(\sigma) = \text{fv}(p(t_1 \dots, t_n)) \\
\sigma, \mathcal{I} \models_S \mathbf{c}(\Psi_1, \dots, \Psi_m, t_1, \dots, t_n) & \text{iff } \sigma, \mathcal{I} \models_S \Psi[\Psi_1/v_1, \dots, \Psi_m/v_m, t_1/v_{m+1}, \dots, t_m/v_{m+n}] \\
& \text{where } \mathbf{c} \mapsto \lambda(v_1, \dots, v_{m+n}).\Psi \in \mathcal{I} \\
\\
\sigma, \mathcal{I} \models_S \mathbf{p}(t_1, \dots, t_n) & \text{iff } \mathbf{p}(\text{Val}_{\sigma, \mathcal{I}}(t_1), \dots, \text{Val}_{\sigma, \mathcal{I}}(t_n)) \mapsto \mathbf{ff} \in \mathcal{I} \\
& \text{and } \text{dom}(\sigma) = \text{fv}(p(t_1 \dots, t_n)) \\
\sigma, \mathcal{I} \models_S \mathbf{c}(\Psi_1, \dots, \Psi_m, t_1, \dots, t_n) & \text{iff } \sigma, \mathcal{I} \models_S \Psi[\Psi_1/v_1, \dots, \Psi_m/v_m, t_1/v_{m+1}, \dots, t_m/v_{m+n}] \\
& \text{where } \mathbf{c} \mapsto \lambda(v_1, \dots, v_{m+n}).\Psi \in \mathcal{I}
\end{array}$$

The notation  $\lambda(v_1, \dots, v_n).\psi$  denotes that  $\psi$  is a formula where  $\{v_1, \dots, v_n\} \subseteq \text{fv}(\psi)$ . Each  $v_i$  is a variable that acts as a placeholder for  $\psi$ , that is substituted when evaluating the external variable mapped to the  $\lambda$ -term, hence the  $\lambda$ -notation.

## 7.2 External world

The constraint-based engine introduced in § 6.2 describes the evolution of a configuration (a set of blocks of constraints). We now assume the existence of a set of primitives  $P$ , each of which provides a single block of constraints to the engine. These primitives can be one of three kinds [5]:

**internal, stateless** denoted by  $P_{no}$ . The underlying constraints involve neither state variables nor communication variables in  $\mathbf{K}$ , and all constraints are persistent—represented by setting the ephemeral constraints to  $\varepsilon_p = \mathbf{tt}$ , where  $p \in P_{no}$ .

**internal, stateful** denoted by  $P_{int}$ . Such primitives have constraints over the state variable pair  $state_p$  and  $state'_p$ , where  $state_p$  represents the value of the current state of  $p \in P_{int}$  and  $state'_p$  the value of the next state. The ephemeral constraint denotes the current state and is always of form  $\varepsilon_p \equiv state_p = t$ , for some ground term  $t$ . No communication variables may appear in the constraints.

**external** denoted by  $P_{ext}$ . Such primitives express constraints in terms of a communication variable  $\mathbf{k}$  through which data is passed from a primitive  $p \in P_{ext}$  to the outside world. The outside world then sends a new set of constraints to represent  $p$ 's next step behaviour. No state variables can appear in the constraints, as it is assumed that the state information is handled externally and incorporated into the constraints sent during the update phase.

We assume that the constraints  $\psi_p$  provided by each primitive  $p \in P$  can only have a fixed set of free variables, denoted by  $\text{fv}(p)$ . Note that  $\text{fv}(\psi_p) \subseteq \text{fv}(p)$ . The relation between external symbols, communication variables and the external primitives in  $P_{ext}$  is made via an *ownership* relation. That is, each external symbol and each communication variable is *owned* by a unique primitive in  $P_{ext}$ .

**Definition 11 (Ownership)** Let  $\mathcal{O} = \mathbf{F} \cup \mathbf{P} \cup \mathbf{C} \cup \mathbf{K}$ . Each  $o \in \mathcal{O}$  is managed by exactly one  $p \in P_{ext}$ . This is denoted using function  $\text{own} : \mathcal{O} \rightarrow P_{ext}$ . We may write  $\mathbf{k}_p$  to indicate that  $\text{own}(\mathbf{k}) = p$ .

We write  $\langle \psi \rangle_Q$  to indicate that the constraints in  $\psi$  are owned by primitives  $Q$ , where  $Q \subseteq P$ .  $\triangleleft$

**Example 3** We extend the constraints of our running example from Table 1, presented in Table 2. Using the updated constraints from Table 2, the global constraint is given by the configuration  $\langle \psi_i \rangle^{i \in 1..7}$ , the synchronous variables are  $\mathcal{X} = \{a, b, c\}$ , the only uninterpreted predicate symbol is equality, **more**  $\in \mathbf{C}$  is an external constraint symbol, **result**  $\in \mathbf{K}$  is a communication variable, and **UserAppr**  $\in \mathbf{P}$  is an external predicate symbol. Furthermore,  $\text{own}(\text{more}) = \text{CDG}$ ,  $\text{own}(\text{result}) = \text{Client}$ , and  $\text{own}(\text{UserAppr}) = \text{User approval}$ .



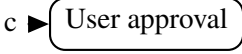
<i>Primitive</i>	<i>Constraint</i>
	$\psi_1 = a \rightarrow (a \wedge (\widehat{a} = d_1 \vee \widehat{a} = d_2 \vee \widehat{a} = d_3 \vee \mathbf{more}(\widehat{a})))$
	$\psi_2 = \mathbf{result} = \widehat{b}$
	$\psi_3 = c \rightarrow (c \wedge \mathbf{UserAppr}(\widehat{c}))$

Table 2: Updated (interactive) constraints of a set of primitives.

The updated constraints in Table 2 illustrate the usage of the extensions to the logic. External constraints can model *on-the-fly* constraint generation. The interpretation of **more** can refer to new external constraints, and this process can be repeated an unbounded number of times. Communication variables provide a mean to communicate the result to the external world, as the constraints of Client show, via the variable **result**. Finally, the external predicate **UserAppr** in  $\psi_3$  illustrates the possibility of asking external primitives if some predicates hold for specific data values.

**Example of the execution of the engine** Recall Example 2, which is based in a set of primitives  $P$ . We partition  $P$  into  $Q$  and  $R$ , where  $Q = \{\text{Client}, \text{CDG}_1, \text{Lossy}_1, \text{Merger}\}$ , and  $R = \{\text{CDG}_2, \text{Lossy}_2\}$ . To provide a better understanding of how the engine evolves with respect to the external interaction, we present a possible trace of the evolution of the constraints. The relation  $\rightarrow^*$  denotes the evolution of the constraints by either applying transformations that preserve the set of possible solutions, or by extending the interpretation  $\mathcal{I}$  based on external interaction. The initial persistent and ephemeral constraints of each primitive  $p \in P$  are denoted by  $\rho_p$  and  $\varepsilon_p$ , respectively. As in Example 2,  $\phi_i$ , where  $i \in 1..6$ , are the constraints for  $\text{CDG}_1$ ,  $\text{CDG}_2$ ,  $\text{Lossy}_1$ ,  $\text{Lossy}_2$ ,  $\text{Merger}$ , and  $\text{Client}$ , respectively.

$$\begin{array}{llll}
\rho_{\text{CDG}_1} = SFA(a) & \varepsilon_{\text{CDG}_1} = \phi_1 & \rho_{\text{Lossy}_1} = \phi_3 \bar{\wedge} SFA(a, b) & \varepsilon_{\text{Lossy}_1} = \text{tt} \\
\rho_{\text{CDG}_2} = SFA(c) & \varepsilon_{\text{CDG}_2} = \phi_2 & \rho_{\text{Lossy}_2} = \phi_4 \bar{\wedge} SFA(c, d) & \varepsilon_{\text{Lossy}_2} = \text{tt} \\
\rho_{\text{Client}} = \phi_6 \bar{\wedge} SFA(e) & \varepsilon_{\text{Client}} = \text{tt} & \rho_{\text{Merger}} = \phi_5 \bar{\wedge} SFA(b, d, e) & \varepsilon_{\text{Merger}} = \text{tt}
\end{array}$$

The initial configuration of the system is given by the set  $\langle \rho_p \wedge \varepsilon_p \rangle_p^{p \in P}$ . We write  $\varepsilon_p^n$  to denote the ephemeral constraint of  $p$  in round  $n$ . During the execution of the engine, both the constraints and the interpretation changes during the solving stage, which we make explicit by using a pair with the interpretation and the constraints. The evolution of a possible trace for our example and its explanation follows:

$$\begin{array}{l}
1 \left\{ \begin{array}{l} \mathcal{I}, \langle \rho_p \wedge \varepsilon_p^1 \rangle_p^{p \in P} \\ \rightarrow^* \mathcal{I}, \langle \phi_1 \wedge \phi_3 \wedge \phi_5 \wedge \phi_6 \wedge SFA(\{a, b, d, e\}) \rangle_Q, \langle \phi_2 \wedge \phi_4 \wedge SFA(\{c, d\}) \rangle_R \\ \rightarrow^* \mathcal{I}, \langle (a \wedge b \wedge e \wedge \mathbf{more}(\widehat{a}) \wedge \widehat{b} = \widehat{a} \wedge \widehat{e} = \widehat{b} \wedge \mathbf{result} = \widehat{e}) \vee \psi_q \rangle_Q, \langle (c \wedge \widehat{c} = d_2 \wedge \neg d) \vee \psi_r \rangle_R \end{array} \right. \\
2 \left\{ \begin{array}{l} \rightarrow^* \mathcal{I}, \langle (a \wedge b \wedge e \wedge \mathbf{more}(\widehat{a}) \wedge \widehat{b} = \widehat{a} \wedge \widehat{e} = \widehat{b} \wedge \mathbf{result} = \widehat{e}) \vee \psi_q \rangle_Q, \langle \rho_r \wedge \varepsilon_r^2 \rangle_r^{r \in R} \end{array} \right. \\
3 \left\{ \begin{array}{l} \rightarrow^* \mathcal{I}', \langle (a \wedge b \wedge e \wedge \widehat{a} = d_4 \wedge \mathbf{evenmore}(\widehat{a}) \wedge \widehat{b} = \widehat{a} \wedge \widehat{e} = \widehat{b} \wedge \mathbf{result} = \widehat{e}) \vee \psi_q \rangle_Q, \\ \langle \rho_r \wedge \varepsilon_r^2 \rangle_r^{r \in R} \\ \rightarrow^* \mathcal{I}', \langle (a \wedge b \wedge e \wedge \widehat{a} = d_4 \wedge \widehat{b} = d_4 \wedge \widehat{e} = d_4 \wedge \mathbf{result} = d_4) \vee \psi'_q \rangle_Q, \langle \rho_r \wedge \varepsilon_r^2 \rangle_r^{r \in R} \end{array} \right. \\
4 \left\{ \begin{array}{l} \rightarrow^* \mathcal{I}, \langle \rho_q \wedge \varepsilon_q^2 \rangle_q^{q \in Q}, \langle \rho_r \wedge \varepsilon_r^2 \rangle_r^{r \in R} \end{array} \right.
\end{array}$$

We now look in more detail into each of the transitions applied above.

1. The blocks of constraints are joined into two blocks based on the partition  $Q$  and  $R$  (of  $P$ ). The persistent and ephemeral constraints are replaced by their definitions. The constraints inside each new block are manipulated following traditional constraint solving techniques until we obtain a disjunction of cases. We then focus on one specific disjunct in each block.
2. The block tagged with  $R$  (in the last step of (1)) has a trivial solution  $\{c \mapsto \text{tt}, d \mapsto \text{ff}, \hat{c} \mapsto d_2\}$  that does not cause any state change. As the boundary conditions hold ( $d \neq \text{tt}$ ), we can perform the update phase on this block. Hereafter the individual blocks for each primitive  $r \in R$  are restored, updating the ephemeral constraints to  $\varepsilon_r^2$ . In this case there is no state change, i.e.,  $\varepsilon_r^2 = \varepsilon_r^1$ .
3. Interaction with the external world is performed to extend the interpretation for **more**, obtaining  $\mathcal{I}' = \mathcal{I} \cup \{\mathbf{more} \mapsto \lambda(v).(v = d_4 \vee \mathbf{evenmore}(v))\}$ . External predicate **more**( $\hat{a}$ ) is replaced by its new interpretation, and the manipulation of the constraints continues as in (1), until we find a new conjunction for the first block which satisfies the trivial solution.
4. In the last step the update phase is performed on the first block. Note that the trivial solution obeys the boundary conditions ( $d \neq \text{tt}$ ). The individual blocks for each primitive  $q \in Q$  are restored, using the corresponding persistent constraints and the new ephemeral constraints for round 2. In this case the ephemeral constraint for the primitive  $\text{CDG}_1$  is updated, while the other primitives in  $Q$  keep the same ephemeral constraints. The update of the constraints of  $\text{CDG}_1$  is performed by querying the external primitive  $\text{CDG}_1$  for its new ephemeral constraints, providing the value of the communication variable of  $\text{CDG}_1$  (**result** =  $d_4$ ). We call this new constraint  $\varepsilon_{\text{CDG}_1}^2$ . After the update, the interpretation “forgets” the value of **more** and is reset to  $\mathcal{I}$ .

We leave for future work the formalisation of the rules that describe the evolution of the constraints and the interpretation during the constraint solving process.

### 7.3 Discussion

Local firings can be discovered concurrently. Furthermore, the explicit connection introduced by the ownership relation, from blocks of constraints and external symbols to external primitives, paves the way for constraint-solving techniques that interact with the external world while searching for solutions for constraints (concurrently). We start by discussing some of our motivation to introduce the local satisfaction relation, and we then explore some more details of our proposed interactive engine.

#### Why locality?

Some of the inspiration for developing a semantic framework that takes into account locality aspects of a model that requires global synchronisation came from experiments undertaken during the development of a distributed implementation of Reo.<sup>4</sup> This implementation is incorporated in the Eclipse Coordination Tools, and its distributed entities roughly correspond to primitives in our constraint approach. There we also make a similar distinction between the two phases of the engine. While developing the distributed engine, we realised the following useful property of the  $\text{FIFO}_1$  channels: in each round it is sufficient to consider the two halves of a  $\text{FIFO}_1$  independently. This property went against the implicit globality assumption in current Reo models, and was never clearly exploited by Reo. This locality property becomes particularly relevant in the extreme case of a Reo connector consisting of several  $\text{FIFO}_1$

<sup>4</sup><http://reo.project.cwi.nl/cgi-bin/trac.cgi/reo/wiki/Tools#DistributedReoEngine>

channels are composed sequentially. In the communication between any two FIFO's from this sequence, traditional Reo models require all the FIFO's to agree, while our distributed implementation requires only the agreement of the two FIFO's involved in the communication.

In more complex Reo connectors, such as the *multiple merger*,<sup>5</sup> it is possible to see that most of the steps involve only the flow on a small part of the connector. It is also possible to find *islands* of synchronous regions, with FIFO channels in the boundaries, where our boundary condition holds for the possible solutions. The approach described in this paper not only justifies the correctness of the locality obtained by the FIFO<sub>1</sub> channels, but it also generalises it to arbitrary solutions where the boundary conditions hold on the boundaries of the synchronous region.

### Interactive engine

We now explore some characteristics of the engine described in § 6.2, using the logic with external symbols introduced in § 7.1. We assume that the interpretation  $\mathcal{I}$  is initially empty regarding external symbols. During the solve stage,  $\mathcal{I}$  is extended every time the external world provides new information about these external symbols. Similarly, the engine can request for the interpretation of specific symbols whenever these are required to find solution. The communication variables play a similar role to state variables. Instead of being directly used in the next round, their value is sent to the primitive that owns the variable, and the engine waits for new (ephemeral) constraints from that primitive. These constraints are then used in the next round.




Figure 4: Interaction with Reo components (left) and with our view of components (right).

The interaction between components and the engine differs in our model with respect to other descriptions of Reo, in that the components play a more active role in the coordination, as depicted in Figure 4. The usual execution of Reo [2] is also divided in two main steps, but the interaction is more restricted in previous models of Reo. In the solve stage the components attempt to write or take a data value. In the update stage the engine requests or sends data values, and restarts the solve stage. In our model we blur the distinction between connectors and components. During the solve stage components can provide constraints with external symbols, that will only be prompted by the engine as required. During the update stage the engine sends the components the values of their communication variables, if

<sup>5</sup><http://homepages.cwi.nl/~proenca/webreo/generated/multimerge/frameset.htm>

defined, and waits for their new constraints for the next round. Our approach therefore offers components the ability to play a more active and dynamic role during the coordination.

## 8 Conclusion and related work

Despite Wegner’s interesting perspective on coordination as constrained interaction [16], little work takes this perspective literally, representing coordination as constraints. Montanari and Rossi express coordination as a constraint satisfaction problem, in a similar and general way [13]. They view networks as graphs, and use the tile model to distinguish between synchronisation and sequential composition of the coordination pieces. In our approach, we explore a more concrete coordination model, which not only captures the semantics of the Reo coordination language [1], but also extends it with a refined notion of locality and a variety of notions of external interaction not found in Montanari and Rossi’s work.

Minsky and Ungureanu took a practical approach and introduced the Law-Governed Interaction (LGI) mechanism [12], implemented by the Moses toolkit. The mechanism targets distributed coordination of heterogeneous agents, enforcing laws that are defined using constraints in a Prolog-like language. The main innovation is the enforcement of laws by certified controllers that are not centralised. Their laws, as opposed to our approach, are not global, allowing them to achieve good performance, while compromising the scope of the constraints. Our approach can express constraints globally, but can solve them locally where possible.

In the context of the Reo coordination language, Lazovik et al. [10] provide a choreography framework for web services based on Reo, where they use constraints to solve a coordination problem. This work is an example of a concrete application of constraints to coordination, using a centralised and non-compositional approach. We formalised and extended their ideas in our work on Deconstructing Reo [5]. The analogy between Reo constraints and constraint solving problems is also pursued by Klüppelholz and Baier [9], who describe a symbolic model checking for Reo, and by Maraïkar et al. [11], who present a service composition platform based on Reo using a mashup’s data-centric approach. The latter can be seen as an scenario where constraint solving techniques are used for executing a Reo-based connector.

One of the main novelties with respect to our previous work [5] is the introduction of a partial semantics to the logic, and techniques for exploiting this semantics. Partiality favours solutions that address only a relevant subset of variables, and can furthermore capture solutions in only part of a network, which cannot be considered independently in a classical setting. Other applications of partial or 3-valued logic exist [4, 8], and model checking and SAT-based algorithms exist for such logics. We do not address verification of partially defined systems, but instead we focus on the specification and execution of these systems. Verification of systems specified by a partial logic would require to assume a fixed interpretation of external symbols, but still presents an interesting challenge, which is out of the scope of this paper. Note that the constraint solving of our partial logic is different from the partial constraint satisfaction problem (PCSP) [7], which consists of finding solutions satisfying a constraint problem  $P$  that are as close as possible to the original problem, although they may be different in most cases.

Faltings et al. [6] explore interactive constraint satisfaction, which bears some similarity to our approach. They present a framework of open constraint satisfaction in a distributed environment, where constraints can be added on-the-fly. They also consider weighted constraints to find optimal solutions. In this paper we do not explore strategies to make the constraint solving process more efficient, such as considering the order in which the rules should be applied. The main differences between our work and theirs are that we focus on the coordination of third parties, making a clear distinction between computation and coordination, we use a partial logic, and we have more modes of interaction.

CRIME (Consistent Reasoning in a Mobile Environment) is an implementation of the Fact Space Model [14], which addresses highly interactive programs in a constantly changing environment. Applications publish their *facts* in a *federated fact space*, which is a tuple space shared by nearby devices. Each fact is defined as a Prolog-like constraint, and the federated fact space evolves as other applications connect or disconnect. The resulting system is a set of reactive objects whose topology is constantly changing. Many of the fact space model ideas are orthogonal to the interaction constraints described in this paper, and its implementation could form a possible base platform for our approach.

## Conclusion

The key contributions of our work are the use of a local logic which does not require all constraints to be satisfied, and the different modes of interaction. Together these enable more concurrency, more flexibility, and more scalability, providing a solid theoretical basis for constraint satisfaction-based coordination models. Furthermore, constraints provide a flexible framework in which it may be possible to combine other constraint based notions, such as service-level agreements. As future work we plan to explore the extension of Reo-based tools, and to implement an interactive and iterative constraint-solving process based on the logic described in this paper. In the process, we will introduce rules describing how to manipulate blocks of constraints that preserve simple solutions, in order to describe in more detail the concurrent search for local firings. Later we plan to explore strategies for the application of these rules, and to understand better the efficiency of our approach.

## References

- [1] F. Arbab (2004): *Reo: a channel-based coordination model for component composition*. *Math. Struct. in Comp. Science* 14(3), pp. 329–366.
- [2] F. Arbab, C. Koehler, Z. Maraïkar, Y. Moon & J. Proença (2008): *Modeling, testing and executing Reo connectors with the Eclipse Coordination Tools*. In: *International Workshop on Formal Aspects of Component Software (FACS)*. Electronic Notes in Theoretical Computer Science (ENTCS), Malaga.
- [3] C. Baier, M. Sirjani, F. Arbab & J. Rutten (2006): *Modeling component connectors in Reo by constraint automata*. *Sci. Comput. Program.* 61(2), pp. 75–113.
- [4] Glenn Bruns & Patrice Godefroid (1999): *Model Checking Partial State Spaces with 3-Valued Temporal Logics*. In: *CAV '99: Proceedings of the 11th International Conference on Computer Aided Verification*. Springer-Verlag, London, UK, pp. 274–287.
- [5] Dave Clarke, Jose Proenca, Alexander Lazovik & Farhad Arbab (2008): *Deconstructing Reo*. *Electr. Notes Theor. Comput. Sci.*, pp. 43–58.
- [6] B. Faltings & S. Macho-Gonzalez (2005): *Open constraint programming*. *Artif. Intell.* 161(1-2), pp. 181–208. Available at <http://dx.doi.org/10.1016/j.artint.2004.10.005>.
- [7] Eugene C. Freuder & Richard J. Wallace (1992): *Partial constraint satisfaction*. *Artif. Intell.* 58(1-3), pp. 21–70.
- [8] Orna Grumberg, Assaf Schuster & Avi Yadgar (2007): *3-Valued Circuit SAT for STE with Automatic Refinement*. In: Kedar S. Namjoshi, Tomohiro Yoneda, Teruo Higashino & Yoshio Okamura, editors: *ATVA, Lecture Notes in Computer Science* 4762. Springer, pp. 457–473. Available at [http://dx.doi.org/10.1007/978-3-540-75596-8\\_32](http://dx.doi.org/10.1007/978-3-540-75596-8_32).
- [9] S. Klueppelholz & C. Baier (2006): *Symbolic Model Checking for Channel-based Component Connectors*. In: *FOCLASA'06*.
- [10] A. Lazovik & F. Arbab (2007): *Using Reo for Service Coordination*. In: *Conf. on Service-Oriented Computing (ICSOC-07)*, Lecture Notes in Computer Sciences 4749. Springer, pp. 398–403.

- [11] Ziyang Maraikar, Alexander Lazovik & Farhad Arbab (2008): *Building Mashups for the Enterprise with SABRE*. In: *ICSOC, Lecture Notes in Computer Science* 5364. pp. 70–83. Available at [http://dx.doi.org/10.1007/978-3-540-89652-4\\_9](http://dx.doi.org/10.1007/978-3-540-89652-4_9).
- [12] Naftaly H. Minsky & Victoria Ungureanu (2000): *Law-governed interaction: a coordination and control mechanism for heterogeneous distributed systems*. *ACM Transactions on Software Engineering and Methodology* 9(3), pp. 273–305.
- [13] Ugo Montanari & Francesca Rossi (1998): *Modeling Process Coordination via tiles, graphs, and constraints*. In: *IDPT'98*.
- [14] Stijn Mostinckx, Christophe Scholliers, Eline Philips, Charlotte Herzeel & Wolfgang De Meuter (2007): *Fact Spaces: Coordination in the Face of Disconnection*. In: Amy L. Murphy & Jan Vitek, editors: *COORDINATION, Lecture Notes in Computer Science* 4467. Springer, pp. 268–285. Available at [http://dx.doi.org/10.1007/978-3-540-72794-1\\_15](http://dx.doi.org/10.1007/978-3-540-72794-1_15).
- [15] George A. Papadopoulos & Farhad Arbab (1998): *Coordination models and languages*. In: M. Zelkowitz (Ed.), *The Engineering of Large Systems, Advances in Computers* 46. Academic Press, pp. 329–400.
- [16] P. Wegner (1996): *Coordination as Constrained Interaction (extended abstract)*. In: *Coordination Languages and Models, Lecture Notes in Computer Sciences* 1061. pp. 28–33.