

ÉCOLE DOCTORALE : SCIENCES MATHÉMATIQUES DE PARIS CENTRE

Équipe-projet Lifeware, Inria

DOCTORAT
Informatique

Thierry Martinez

**Execution models for Constraint Programming:
kernel language design through semantics equivalence.**

Modèles d'exécution pour la programmation par contraintes :
conception d'un langage noyau par le biais d'équivalences sémantiques

Thèse dirigée par François Fages

Soutenue le 17 décembre 2015

JURY

M. Roberto Di Cosmo	Professeur des universités Université Paris 7	Président du jury
M. François Fages	Directeur de recherche Inria	Directeur de thèse
M. Tom Schrijvers	Professeur Katholieke Universiteit Leuven	Rapporteur
M. Maurizio Gabbrielli	Professeur Università di Bologna	Rapporteur
M. Jean Goubault-Larrecq	Professeur École normale supérieure de Cachan	Examineur
M. Frank Valencia	Chargé de recherche CNRS	Examineur

Contents

1	Introduction	11
1.1	Committed-Choice Programming	12
1.2	Constraint Model and Search	13
1.3	Angelic Programming	14
1.4	Thesis	16
1.5	Publications	17
2	Ask-lifting: from LLCC to CHR	19
2.1	Syntax & Semantics of CHR and LLCC	19
2.1.1	Syntax and Semantics of CHR	19
2.1.2	Syntax and Semantics of LLCC	23
2.1.3	Circumscribing non-determinism in CHR and LLCC operational semantics	26
2.2	Translations between sub-languages CSR and flat-LLCC	27
2.2.1	From Constraint Simplification Rules (CSR) to flat-LLCC	27
2.2.2	From flat-LLCC to CSR	29
2.2.3	CHR Linear-Logic and Phase Semantics Revisited	31
2.3	Ask-lifting: encoding LLCC into CSR	33
2.3.1	Preliminary step: labeling LLCC-agents	33
2.3.2	The ask-lifting transformation	34
2.4	Encoding the call-by-value λ -calculus in CHR	35
3	Reified Search: search strategies as constraint solving	39
3.1	Constraint Satisfaction	39
3.1.1	Constraints as Relational Languages	39
3.1.2	Propagators and Consistency	43
3.1.3	Abstractions	45
3.2	The Language ClpZinc	46
3.3	Extending Zinc with CLP Clauses	49
3.3.1	Partial Evaluation of ClpZinc into And/Or Trees	49
3.3.2	Compiling And/Or Trees into Zinc Reified Constraints	50

3.4	Computation Results on Korf’s Square Packing Benchmark	53
3.5	LDS and SBDS as Strategy Transformers in ClpZinc	57
3.5.1	Limited Discrepancy Search	57
3.5.2	Symmetry Breaking During Search	58
3.6	Beyond Tree Search Strategies	59
3.7	ClpZinc Models of First-Fail and Middle-out Strategies	60
3.8	ClpZinc Model of Symmetry Breaking During Search	62
4	Angelic Programming: from compound guards to atomic kernel for LLCC	67
4.1	Angelic Semantics and Derivation nets	67
4.1.1	Process calculi	67
4.1.2	Nets for sets of reductions	68
4.1.3	Sharing strategies	70
4.1.4	Derivation nets in presence of Propagation Rules	72
4.2	Angelic Programming	73
4.2.1	Head Decomposability	73
4.2.2	Controlling the Angelism	74
4.3	Applications	74
4.3.1	Angelic CHR \forall	74
4.3.2	Meta-interpreters	74
4.3.3	User-defined Indexation	75
4.3.4	CHR with Ask and Tell	77
4.3.5	Hierarchical Definition of a Rational Terms Constraint Solver	86
4.4	The SiLCC Programming Language	90
4.4.1	Procedures	91
4.4.2	Sequencing	92
4.4.3	Arithmetic and functional notation	94
4.4.4	Conditional	94
4.4.5	Angelism	96
4.4.6	Modules and the package system	97
4.4.7	Recursion	98
4.4.8	Lists and iterations	99
4.4.9	Pattern-matching and terms	99
4.4.10	References	101
4.4.11	Records	102
4.4.12	Hash-tables and argument indexing	102
4.4.13	Asks on variable number of tokens	103
4.5	Related work	105

5	Conclusion	107
5.1	Committed-Choice Semantics for LLCC through CHR	108
5.2	Constraint Model and Search	108
5.3	Angelic Programming	109
5.4	Language Implementations	110

Summary

Logic programming and constraint programming are two declarative programming paradigms which rely on the identification of programs to theories, and programming to modeling. Execution models result from the operational interpretation of logical provability in logic programming, and of constraint propagation in constraint programming. However, the control of execution is crucial for the practicability of these schemes and extra-logical traits are thus added in those programming systems, with the classical slogans "logic program = logical theory + control", "constraint program = constraint model + search".

This thesis investigates execution models in which control and search can be shifted into the logic or the constraint model, while preserving the semantics. The three parts of the thesis correspond to the three semantics equivalence that are showed: the first between two committed-choice forward-chaining logic languages, the second between constraint logic programs and constraint models, and the third between guard semantics in angelic settings. Each of these equivalence is constructive in the sense that there exists an encoding that enables the compilation from one of the paradigm to the other.

First, we show that simple program transformations exist back and forth between Constraint Handling Rules (CHR) and Linear Logic Concurrent Constraint (LLCC) languages, making them semantically equivalent even if syntactically different, which closes the question of implementing a committed-choice semantics for LLCC by using CHR as kernel language.

Secondly, we show that a wide variety of search procedures can be internalized in the constraint model with a fixed enumeration strategy. Transforming search procedures into constraint satisfaction problems presents several advantages: (1) it makes search strategies declarative and modeled as constraint satisfaction problems; (2) it makes it possible to express search strategies in existing front-end modeling languages without any extension; (3) it opens up constraint propagation algorithms to search constraints and to the implementation of novel search procedures based on constraint propagation. This is illustrated with the design of the ClpZinc modeling language, with an angelic interpretation of Horn clauses which allows compilation to a reification-based constraint kernel.

Finally, in concurrent constraint logic programming, committed-choice semantics create a hierarchy of non-equivalent semantics axed on the expressive power of synchronization mechanism. We show that the hierarchy of guard semantics collapses with angelic semantics, allowing the most primitive synchronization mechanism to encode all the others. The main consequence of this collapse is the identification of a kernel language, with a primitive synchronization mechanism and an elementary constraint system which are sufficient to reconstruct the other forms of synchronizations and other constraint systems as modules, in the software engineering sense.

Résumé

La programmation logique et la programmation par contraintes sont deux paradigmes de programmation déclaratifs qui reposent sur l'identification entre programmes et théories logiques d'une part, et entre programmation et modélisation d'autre part. Les modèles d'exécution résultent de l'interprétation opérationnelle de la prouvabilité logique en programmation logique, et de la propagation des contraintes en programmation par contraintes. Cependant, le contrôle de l'exécution est crucial pour rendre ces paradigmes utilisables en pratique et des traits extra-logiques sont ajoutés à cette fin dans ces systèmes de programmations, avec les slogans classiques « programmation logique = théorie logique + contrôle » et « programmation par contraintes = modèle de contraintes + recherche ».

Cette thèse explore des modèles d'exécution au sein desquels le contrôle et la recherche peuvent être remontés selon le contexte dans la logique ou dans le modèle de contraintes, tout en préservant la sémantique. Les trois parties de cette thèse correspondent aux trois équivalences sémantiques démontrées : la première entre deux langages logiques en chaînage avant avec choix commis, la seconde entre programmes logiques avec contraintes et modèles de contraintes, la troisième entre les sémantiques des gardes dans un cadre angélique. Chacune de ces équivalences est constructive au sens qu'il existe un codage permettant la compilation depuis l'un des paradigmes dans l'autre.

Premièrement, nous montrons que des transformations simples de programmes existent dans les deux sens entre les langages Constraint Handling Rules (CHR) et Linear Logic Concurrent Constraint (LLCC), les rendant sémantiquement équivalents bien que syntaxiquement différents, ce qui ferme la question de l'implémentation d'une sémantique de choix commis pour LLCC, en utilisant CHR comme langage noyau.

Deuxièmement, nous montrons qu'une large variété de procédures de recherche peut être internalisée dans le modèle de contraintes avec une stratégie d'énumération fixe. La transformation de procédures de recherche en problèmes de satisfaction de contraintes présente plusieurs avantages : (1) cela rend les stratégies de recherche déclaratives et modélisables au même titre que les problèmes de satisfaction de contraintes ; (2) cela rend possible l'expression de stratégies de recherche dans les

langages de modélisation existants, sans recourir à des extensions ; (3) cela ouvre la voie à des algorithmes de propagation de contraintes pour les contraintes de recherche et à l'implémentation de nouvelles procédures de recherche basées sur la propagation de contraintes. Ceci est illustré par la conception du langage de modélisation ClpZinc, avec une interprétation angélique des clauses de Horn qui permet la compilation vers un noyau de contraintes basé sur la réification.

Pour finir, en programmation concurrente avec contraintes, la sémantique par choix commis crée une hiérarchie de sémantiques non équivalentes axée sur la puissance d'expression des mécanismes de synchronisation. Nous montrons que la hiérarchie des gardes s'effondre avec la sémantique angélique, permettant au plus primitif des mécanismes de synchronisation de coder tous les autres. La principale conséquence de cet effondrement est l'identification d'un langage noyau, avec un mécanisme de synchronisation primitif et un système de contraintes élémentaires qui sont suffisants pour reconstruire les autres formes de synchronisation et les autres systèmes de contraintes comme modules, au sens de l'ingénierie logicielle.

Chapter 1

Introduction

Program semantics are the mathematical abstractions of program executions. In the settings of logic programming [15, 49], the program semantics results from the operational interpretation of logical provability: the execution is a proof search procedure for the program interpreted as a logical property expressed in a computational fragment of a logic [62]. In the settings of constraint programming, derived from constraint propagation algorithms first introduced by Waltz for computer vision [88], the execution is an exploration procedure of a combinatorial space searching for a variable assignment satisfying all the relations, the “constraints”. In the settings of concurrent programming, the execution is a certain interleaving of computation paths and message transfers. In all these settings, the execution path does not directly follow the program structure: the precise execution relies on the proof method, the search strategy or the scheduler. The precise semantics of the programs is therefore intimately linked with the execution model.

Logic programming and constraint programming have been linked into constraint logic programming [45] which generalizes from the Herbrand domain of terms with equalities to arbitrary domains of partially-known values with constraints. In the constraint programming settings, constraint propagators act concurrently to reduce variable domains. The use of variable domains as communication channel inspired the foundation of the class CC of *Concurrent Constraint* programming languages [53, 73] which rely on a model of concurrent computation, where agents communicate through a shared constraint store, with a synchronization mechanism based on constraint entailment. In classical constraint settings, the store evolves monotonically, similarly to the built-in constraint store of CHR. The LLCC languages [72, 28] introduce linear constraint systems, based on Girard’s intuitionistic linear logic (ILL) [35]. A remarkable kind of linear constraints are linear tokens [28], which can be freely added or consumed, comparably to CHR constraints. Linear logic leads to a natural semantics for classical CC languages as well [28].

Logic programming and constraint programming share the common trait of lifting programming into modeling: as in linear programming [12], or mathematical programming more generally, the paradigm tends to identify the program with a model and the execution with its resolution. However, since the introduction of Prolog [15, 49], logic programming has been living with the dichotomy: “programs = logic + control”. Similarly, constraint programming is traditionally presented as the combination of two components: a constraint model and a search procedure [87]. In the general case, concurrent programs are not necessary confluent and their execution relies upon scheduler choices. While denotational semantics are faithful with respect with the logic of a logical program, the model of a constraint satisfaction problem or the set of all execution paths of a concurrent program, the control suffers from being only described by a procedural interpretation.

This thesis investigates how execution models can internalize the control into the logic.

1.1 Committed-Choice Programming

Constraint Handling Rules (CHR) [30] is a rule-based declarative programming language. CHR and LLCC have been developed independently and with distinct concerns. More recently, a precise declarative semantics for CHR has been described in linear logic [10]. Implementations of CHR follow a committed-choice forward-chaining execution model: the non-determinism of the abstract semantics is partly refined with extra-logical syntactic convention on the program order and possibly notations for weighted semantics (with priorities or probabilities), and partly left unspecified in the underlying compiler.

In CHR, programs are sets of transformation rules on constraint stores. Some constraints are built-ins and can only be accumulated into the store. Other constraints are user-defined and can be added or deleted. Although initial motivations were the definition of constraint solvers and propagators, nowadays applications include typing [17, 85], software testing [66], scheduling [4] and so on.

The linear-logic CC language (or LLCC) provides the non-monotonous traits for imperative programming in addition to traits for logic, concurrent and constraint programming, the whole in a simple and pure semantics. This semantics identifies a fragment of first-order linear logic as a monoparadigm programming language with a rich expressive power.

Two translations from CHR to LLCC and back are proposed, both preserving the semantics. Strong bisimilarity results are formulated. As direct corollary, we obtain a natural encoding of the λ -calculus in CHR. While existence of low-level translations is guaranteed by Turing-completeness *via* a compilation process, there are more fine-grained criteria to compare expressiveness [36]. In particular,

translations presented here are natural and (relatively) agnostic with respect to the constraint theory.

As every logic-based language, the declarativity of Constraint Handling Rules [31, 84] relies on the logical interpretation of the programs. However, this interpretation hides syntactic conventions, like the order of the rules, distinguished abbreviations such as propagation rules, and annotations which control the effective execution of the program. These control features are formally described in a hierarchy of semantics, from the abstract semantics ω_{va} [31] to more fine-grained semantics, describing the handling of propagations (the theoretical semantics ω_t), of the rule ordering (the refined semantics ω_r [25]), or of annotations like priorities [21] or probabilities [81].

This thesis formalizes connections between CHR with naive operational semantics and LLCC.

1.2 Constraint Model and Search

Front-end modeling languages are designed for solving problems independently of the solvers, and solving them with generic solvers using fixed search procedures (e.g. Essence [29]), or contain special features for specifying the search strategy for the constraint solvers (e.g. Zinc [64]).

In this chapter, we show that a completely different approach for specifying search is possible, by internalizing the search procedure in the constraint model with a fixed enumeration strategy. In principle, transforming search procedures into constraint satisfaction problems presents several advantages:

1. it makes search strategies declarative, and modeled as constraint satisfaction problems;
2. it makes it possible to express search strategies in existing front-end modeling languages without any extension;
3. it opens up constraint propagation algorithms to search constraints and to the implementation of novel search procedures based on constraint propagation.

The idea of this transformation is to associate to each choice point a reified constraint with an auxiliary model variable for representing that choice (e.g. value enumeration, domain splitting or any constraint). The search heuristic can then be specified simply by the enumeration strategy for the choice variables. This approach is not limited to static search procedures in which all choice points are precisely known statically, but can accommodate dynamic search strategies, such

as dichotomic or interval splitting search [80] for example. In constraint programming, dynamic search procedures rely on the values of indexicals (domain size, minimum value, etc.). They are expressed in the framework presented here by extending the enumeration strategy with annotations that assign the values of indexicals to auxiliary model variables. Static search procedures do not rely on the values of indexicals and their encoding do not need any specific support on the solver-side. The encoding of dynamic search procedures can be run through simple additions in the solvers for providing the capability to query the values of indexicals.

In this chapter, to make concrete the presentation of the transformation, we consider the Zinc modeling language and introduce ClpZinc¹, a language extending Zinc with the ability to describe new relations by Horn clauses. The choice of CLP as a specification language for search procedures is guided by CLP being the smallest language with the addition of constraint to the store as primitive and closed by conjunction and disjunction (for expressing choices), and with a general form of recursion. Given a constraint system \mathcal{X} (e.g. finite domains) and the Herbrand constraint system \mathcal{H} , we describe a partial evaluation procedure to transform any terminating CLP($\mathcal{X} + \mathcal{H}$) goal to an and/or tree with constraints over \mathcal{X} .

1.3 Angelic Programming

Finally, this thesis proposes an alternative execution model which explores all the possible choices, by opposition to the committed-choice strategy. This execution model is angelic in the sense that if there exists a successful execution strategy (with respect to a given observable), then this strategy will be found. Formally, the set of computed goals is complete with respect to the set of the logical consequences of the interpretation of the initial goal in linear logic. In practice, this chapter introduces a new data representation for sets of goals, the derivation nets. Sharing strategies between computation paths can be defined for derivation nets to make execution algorithmically tractable in some cases where a naive exploration would be exponential. Control for refined execution is recovered with the introduction of user constraints to encode sequencing, fully captured in the linear-logic interpretation. As a consequence of angelic execution, CHR rules become decomposable while preserving accessibility properties. This decomposability makes natural the definition in angelic CHR of meta-interpreters to change the execution strategy. More generally, arbitrary computation can be interleaved during head matching,

¹The Clp2Zinc compiler that transforms ClpZinc models into MiniZinc is available for download, together with patches for Choco, JaCoP, SICStus, Gecode and or-tools: <http://lifeware.inria.fr/~tmartine/clp2zinc/>

for custom user constraint indexation and deep guard definition.

CHR enjoys two logical interpretation: the first to have been introduced historically interprets rules and goals as first-order classical logic formulae; more recently [10], an interpretation as first-order linear logic [35] formulae has been given. The latter provides a finer reading of the dynamics of the rules and will be the logical interpretation considered in this chapter.

All these semantics are correct with respect to the linear-logic interpretation: if a configuration is reachable through any of these operational semantics, then this configuration is indeed a linear-logic consequence of the initial goal. However, only the abstract semantics enjoys completeness: the purpose of all other semantics is to provide syntactic construction to force the execution to choose some particular branches. The downside is that these scheduling choices escape the declarative framework provided by logic. The programmer should ensure that the scheduler can only make the good choice, either by writing a confluent program or by relying on extra-logical traits (order of the rules, priorities, etc.) to drive the scheduler.

Focusing on completeness entails the exploration of all the logical consequences of the interpretation of the initial goal in linear logic. For this purpose, we propose angelic scheduling as an alternative execution model for CHR. Observationally, the scheduler always makes the good choice: more precisely, if a successful (i.e., non-blocking) choice exists, it will be explored. Accessible configurations exactly match the set of logical consequences of the linear interpretation of the initial goal: the operational behavior is fully described by the linear-logic interpretation, including the control. More formally, linear logic is the most faithful logic for CHR [10], since it captures the non-monotonous evolution of configurations. Control structures like sequencing and branching have natural encoding in this logic and their usage for CHR have already been showed through the log-linear encoding of RAM machines [83].

In angelic settings, the atomicity of head consumption is not essential, in opposition to the committed-choice case. Since absence of user constraints cannot be observed, partial head consumptions just lead to silent unsuccessful computation branches. This property allows the interleaving of arbitrary computations between multiple head consumptions. Meta-interpreters for CHR rules can therefore be written by sequencing the consumption of the successive parts of the head. Specific representations can be chosen for some heads to enable user-defined indexation strategy. To reduce the combinatorial explosion among computation branches, the formalism of derivation nets is introduced: this formalism provides a graphical representation for sets of computation paths. Non-determinism during the execution of a CHR program can be intrinsic to the rule dynamics, and all choices should be explored, but the abstract operational semantics suffers from a large part of scheduling non-determinism between independent paths of the com-

putation that should be quotiented for a tractable execution. The derivation nets are a convenient representation to define sharing strategies between computation paths to eliminate scheduling non-determinism. Two decidable sharing strategies are explored in this chapter. The first strategy shows that optimal sharing is decidable but is computationally expensive. The second one is polynomial in the worst case and induces essentially a constant overhead in practice while being optimal relatively to a conservative interpretation of user constraint identity.

Angelic semantics have been identified as the natural semantics for Concurrent Constraint (CC) programming languages [46] since the very beginning of the introduction of this language family: in this forward-chaining framework, the set of accessible computations is more natural to link with a logical interpretation than a particular computation path. However, the CC language and its angelic semantics is considered in [46] as an abstract language to reason about concurrency-related questions that can be captured in this formalism: there is no consideration about implementation. Moreover whenever CC languages have only a monotonous interpretation in classical logic, LLCC and CHR handle non-monotonous traits with consumptions.

Besides these algorithmic results, the angelic execution of programs has numerous semantically good properties. The behavior of programs is by construction precisely captured by the logical interpretation. Committed-choice semantics create a hierarchy of non-equivalent semantics axed on the expressive power of synchronization mechanism: this hierarchy collapses with angelic semantics, allowing the most primitive synchronization mechanism to encode all the others. The main consequence of this collapse is the identification of a kernel language, with this primitive synchronization mechanism and an elementary constraint system, sufficient to reconstruct the other forms of synchronizations and other constraint systems as modules, in the software engineering sense. The kernel language can be seen as a generalization of Warren’s abstract machine for Prolog [89].

1.4 Thesis

This thesis proposes three program transformations that describe three execution models for constraint programming that are guided by lifting programming into modeling. Constraint Handling Rules are used as a compilation target for the committed-choice semantics for LLCC, constraint satisfaction problems are used as a compilation target for search strategies, and logical interpretation of agents, omniscient over execution paths, enables extra-logical control operators to be modeled as blocked paths.

- The second chapter, “Ask-lifting: from LLCC to CHR”, shows that the question of implementing a committed-choice semantics for LLCC can be reduced

to a program transformation to the actively developed field of CHR compilation,

- The third chapter, “Reified Search: search strategies as constraint solving”, shows that the fundamental dichotomy between modeling and search in constraint programming can in fact be eliminated by showing how search strategies can be modeled as constraint satisfaction problems with fixed enumeration strategies,
- The fourth chapter, “Angelic Semantics: from compound guards to atomic kernel for LLCC”, shows that the question of choosing sophisticated semantics for control under the presence of non-elementary asks can be reduced to a simple kernel language with atomic guards.

Through these three semantics equivalences, between LLCC and CHR, between tree-search exploration and constraint satisfaction, and between deep guards and atomic guards in angelic settings, this thesis proposes three novel compilation schemes, with CHR, CSP with fixed enumeration strategy, and atomic asks as kernel languages.

1.5 Publications

The results of the second chapter have been published in PPDP’10 [55]. The results of the third chapter have been published in PPDP’15 [58], which extends a preliminary work on another modeling language published in RAC’11 [54]. Parts of the results of the fourth chapter have been published in CHR’11 [56] and CHRat was presented at CHR’08 [27].

In parallel to the work presented in this thesis, I contributed to the design and the complexity analysis of a constraint model for enumerating minimal siphons in Petri nets [63], and to the complexity analysis and ordering properties of subgraph epimorphism [33]. Among the applications of constraint programming, I studied in particular placement problems, like the Korf’s problem that is presented in the third chapter. I worked on stochastic and continuous solving of placement problems [59], and the modelling of continuous problems in MiniZinc [57].

Chapter 2

Ask-lifting: from LLCC to CHR

Section 2.1 presents CHR and LLCC in full generality and recalls some already published and well-known results. Section 2.2 focuses on distinguished subsets *Constraint Simplification Rules* (CSR) and flat-LLCC, provides translations between these two subsets. Linear logic semantics [10] and phase semantics [39] of CHR are recovered as corollary. Section 2.3 introduces the *ask-lifting* transformation from full LLCC to flat-LLCC. Section 2.4 presents the encoding of the call-by-value λ -calculus in CHR.

2.1 Syntax & Semantics of CHR and LLCC

Let \mathcal{V} be a set of variables, and Σ a signature for constant, function and predicate symbols. The set of free variables of a formula e is denoted $\text{fv}(e)$, a sequence of variables is denoted by \mathbf{x} . $e[\mathbf{t}/\mathbf{x}]$ denotes the formula e in which free occurrences of variables \mathbf{x} are substituted by terms \mathbf{t} (with the usual renaming of bound variables to avoid variable clashes).

For a set S , S^* denotes the set of finite sequences of elements of S and $\mathcal{M}(S)$ denotes the set of finite multi-sets of elements of S . More formally, $(S^*; \cdot; \varepsilon)$ denotes the free monoid and $(\mathcal{M}(S); \cdot; \emptyset)$ the free commutative monoid over S . For relations \mathcal{R} and \mathcal{R}' , $a \mathcal{R} \cdot \mathcal{R}' c$ if there exists b such that $a \mathcal{R} b \mathcal{R}' c$. For a relation \rightarrow , $\overset{*}{\rightarrow}$ is the reflexive and transitive closure of \rightarrow .

2.1.1 Syntax and Semantics of CHR

Notation 1. For any set S , let $\mathcal{M}(S)$ be the set of multisets on S . For every $m \in \mathcal{M}(S)$, the *characteristic application* of m is denoted $\mathbf{1}_m : S \rightarrow \mathbf{N}$. The support of a multiset $m \in \mathcal{M}_f(S)$ is $\{s \in S \mid m(s) \geq 1\}$. The set of finite multisets on S is denoted $\mathcal{M}_f(S)$, this is the set of multisets m with finite support. For

every $s \in S$, $\mathbf{1}_m(s)$ is the *power* of s in m . For any set S' and any application $f : S \rightarrow S'$, we note $f(m)$ the image multiset of m by f , with the characteristic application $\mathbf{1}_{f(m)} : s' \in S' \rightarrow \sum_{s \in f^{-1}(s')} \mathbf{1}_m(s)$.

Syntax.

Let \mathcal{P}_b and \mathcal{P}_c be two disjoint subsets of predicate symbols in Σ . Predicates built from Σ with predicate symbols in \mathcal{P}_b are *atomic built-in constraints*, their set is denoted \mathcal{B}_0 . *Built-in constraints* are conjunctions of atomic built-in constraints, their set is denoted \mathcal{B} . Predicates built from Σ with predicate symbols in \mathcal{P}_c are *atomic CHR constraints*, their set is denoted \mathcal{U}_0 . *CHR constraints* are (finite) multi-sets of atomic CHR constraints, their set is denoted \mathcal{U} . A *goal* is a multi-set of built-in constraints and CHR constraints.

Built-in constraints are supposed to include the syntactic equality $=$. There is a *constraint theory* CT over the built-in constraints: CT is supposed to be a non-empty, consistent and decidable first-order theory. For two multi-sets $H = (H_1, \dots, H_m)$ and $H' = (H'_1, \dots, H'_n)$, $H \doteq H'$ denotes the formula $H_1 = H'_1 \wedge \dots \wedge H_m = H'_n$ if $m = n$, and **false** if $m \neq n$ [3].

Definition 1 (Syntax). A *CHR program* is a sequence of rewriting rules, or *CHR rules*, each rule being denoted $\langle H \setminus H' \Leftrightarrow G \mid B \rangle$ where *heads* H and H' are CHR constraints such that $\langle H, H' \rangle \neq \emptyset$, the *guard* G is a built-in constraint, and the *body* B is a goal.

A *simplification rule* is a rule where H' is empty, and is denoted $H' \Leftrightarrow G \mid B$. A *propagation rule* is a rule where H' is empty, denoted $H \Rightarrow G \mid B$. A *simpagation rule* is a rule where H and H' are both non-empty. H and H' cannot be both empty.

Example 1. The CHR program below, adapted from [9], describes the dining philosophers protocol [23], where N philosophers are sitting around a table and alternate thinking and eating. N forks are dispatched between them. Each philosopher is in competition with her neighbors to take her two adjacent forks and eat.

$$\begin{aligned}
& \text{diner}(N) \Leftrightarrow \text{recphilo}(0, N). \\
& \text{recphilo}(I, N) \Leftrightarrow \\
& \quad J \text{ is } (I + 1) \bmod N, \text{philo}(I, J), \text{fork}(I), \\
& \quad \text{nextphilo}(I, N). \\
& \text{nextphilo}(I, N) \Leftrightarrow I < N - 1 \mid \\
& \quad J \text{ is } I + 1, \text{recphilo}(J, N). \\
& \text{philo}(I, J) \setminus \text{fork}(I), \text{fork}(J) \Leftrightarrow \text{eat}(I, J). \\
& \text{eat}(I, J) \Leftrightarrow \text{fork}(I), \text{fork}(J).
\end{aligned}$$

Logical Semantics.

CHR programs enjoy two declarative semantics: one in classical logic for left-linear programs [30], and one in linear logic without restriction [10]. A rule is *left-linear* if a user defined constraint matches at most one constraint in its head. A program is left-linear when all rules are.

In classical logic, a multiset M of constraints is interpreted by the conjunction $M^\dagger = \langle \bigwedge_{c \in \text{sup}(M)} c \rangle$ of its elements. For each CHR rule:

$$(H \setminus H' \Rightarrow G \mid B_x, B_u)^\dagger = \forall \mathbf{x} (\exists \mathbf{y} (G) \wedge H^\dagger \rightarrow (H'^\dagger \leftrightarrow \exists \mathbf{z} (G \wedge B_x \wedge B_u^\dagger)))$$

where $\mathbf{x} = \text{fv}(H_0, H_1)$, $\mathbf{y} = \text{fv}(G) \setminus \text{fv}(H_0, H_1)$ and $\mathbf{z} = \text{fv}(G, B_x, B_u) \setminus \text{fv}(H_0, H_1)$. The logical interpretation of a program $(P)^\dagger$ is the logical conjunction of the interpretations of the rules of P . The logical semantics of a query q is:

$$\mathcal{L}_P(q) = \{c \mid (P)^\dagger \models_{\mathcal{X}} q \rightarrow c\}$$

$\mathcal{L}_P(q)$ is closed by logical implication in this translation: for any set S of logical facts, we denote $\downarrow S = \{c' \mid \exists c \in S, \models_{\mathcal{X}} c \rightarrow c'\}$, then $\mathcal{L}_P(q) = \downarrow \mathcal{L}_P(q)$.

Let V be the free variables of the initial query (T_0, c_0) . Configurations are interpreted as $(T, c)^\dagger = \langle \exists \mathbf{x} ((T)^\dagger \wedge c) \rangle$ where $\mathbf{x} = \text{fv}(T, c) \setminus \text{fv}(T_0, c_0)$: the query variables appear in the observables and are therefore left free in the interpretation. The notation is extended for sets S of configurations: $S^\dagger = \{(T, c)^\dagger \mid (T, c) \in S\}$.

Theorem 1 ([30]). *For any left-linear CHR(\mathcal{X}) program P and initial query (T_0, c_0) :*

$$\downarrow (\mathcal{O}_P^a(T_0, c_0))^\dagger \subseteq \mathcal{L}_P((T_0, c_0)^\dagger)$$

This formulation gives only a soundness result: $\mathcal{L}_P((T_0, c_0)^\dagger) \subseteq \downarrow (\mathcal{O}_P^a(T_0, c_0))^\dagger$ does not hold in general. For instance, consider the program $\{a \Leftrightarrow b. a \Leftrightarrow c.\}$: $\mathcal{L}_P((a, \top)^\dagger) = \downarrow \{a \wedge b \wedge c\}$ whereas $\downarrow (\mathcal{O}_P^a(a, \top))^\dagger = \{a, b, c, \top\}$. Weaker soundness and completeness hold: $\mathcal{L}_P((T_0, c_0)^\dagger) = \mathcal{L}_P((\mathcal{O}_P^a(T_0, c_0))^\dagger)$, but completeness is just a consequence of the immediate membership $(T_0, c_0)^\dagger \in (\mathcal{O}_P^a(T_0, c_0))^\dagger$. If all rules of P are simplification rules [3] and if $(\mathcal{O}_P^t(T_0, c_0))^\dagger \neq \emptyset$ (that is to say, P has at least one finite computation), then $\mathcal{L}_P((T_0, c_0)^\dagger) \subseteq \mathcal{L}_P((\mathcal{O}_P^t(T_0, c_0))^\dagger)$: in this case, the property is a direct consequence of $(T_0, c_0)^\dagger \in \mathcal{L}_P((\mathcal{O}_P^t(T_0, c_0))^\dagger)$.

The classical logical semantics is not correct if P is not left-linear. For instance, with the single rule $a, a \Leftrightarrow b$ we have $\mathcal{L}_P((a, \top)^\dagger) = \downarrow \{a \wedge b\}$ while $\downarrow (\mathcal{O}_P^a(a, \top))^\dagger = \downarrow \{a\} \neq \downarrow \{a \wedge b\}$. Furthermore, this logical semantics identifies too many programs (for instance, $\{a \Leftrightarrow b. a \Leftrightarrow c.\}$ and $\{a \Leftrightarrow b, c.\}$).

To overcome these limitations, a declarative semantics in Girard's linear logic [35] have been developed in [10] for all CHR programs. Built-in constraints over \mathcal{X}

are translated using Girard's translation of classical logic in linear logic with the bang operator [35]. A multiset M of constraints is interpreted by a linear tensor of the constraints $M^{\dagger\dagger} = \langle \bigotimes_{c \in M} c \rangle$. For each CHR rule:

$$(H \setminus H' \Rightarrow G \mid B)^{\dagger\dagger} = \langle !\forall \mathbf{x} (\exists \mathbf{y} (G^{\dagger\dagger}) \otimes H^{\dagger\dagger} \otimes H'^{\dagger\dagger} \multimap \exists \mathbf{z} (H^{\dagger\dagger} \otimes G^{\dagger\dagger} \otimes B_{\mathcal{X}}^{\dagger\dagger} \otimes B_u^{\dagger\dagger})) \rangle$$

where $\mathbf{x} = \text{fv}(H, H')$, $\mathbf{y} = \text{fv}(G) \setminus \text{fv}(H, H')$ and $\mathbf{z} = \text{fv}(G, B_{\mathcal{X}}, B_u) \setminus \text{fv}(H, H')$.

The linear logical interpretation of a program $(P)^{\dagger}$ is the linear tensor of the interpretations of the rules of P . The linear logical semantics of a query q is:

$$\mathcal{LL}_P(q) = \{c \mid (P)^{\dagger\dagger} \models_{LL, \mathcal{X}} q \multimap c \otimes \top\}$$

$\mathcal{LL}_P(q)$ is closed by linear implication: for any set S of logical facts, we denote $\downarrow S = \{c' \mid \exists c \in S, \models_{LL, \mathcal{X}} c \multimap c' \otimes \top\}$, then $\mathcal{LL}_P(q) = \downarrow \mathcal{LL}_P(q)$.

Let V be the free variables of the initial query (T_0, c_0) . Configurations are interpreted as $(T, c)^{\dagger\dagger} = \langle \exists \mathbf{x} (T^{\dagger\dagger} \otimes c) \rangle$ where $\mathbf{x} = \text{fv}(T, c) \setminus \text{fv}(T_0, c_0)$: the query variables appear in the observables and are therefore left free in the interpretation. The notation is extended to sets S of configurations: $S^{\dagger\dagger} = \{(T, c)^{\dagger\dagger} \mid (T, c) \in S\}$.

Theorem 2 ([10, 28]). *For any CHR(\mathcal{X}) program P and initial query (T_0, c_0) :*

$$\downarrow (\mathcal{O}_P^a(T_0, c_0))^{\dagger\dagger} = \mathcal{LL}_P((T_0, c_0)^{\dagger\dagger})$$

Operational Semantics.

A *state* is a tuple denoted $\langle g; b; c \rangle_V$ where g is a goal, b is a built-in constraint, c is a CHR constraint and V is a set of variables. The relation \equiv_C over states is the smallest equivalence relation such that:

- $\langle g; b; c \rangle_V \equiv_C \langle g; b'; c \rangle_V$ for $CT \models b \leftrightarrow b'$;
- $\langle g; b; c \rangle_V \equiv_C \langle g; b; c \rangle_V[y/x]$ for variables $x \notin V$ and $y \notin V \cup \text{fv}(g, b, c)$.

Let \mathcal{P} be the set of pairs of CHR programs and states.

Definition 2 (Naive Operational Semantics [30]). A CHR program P is executed along a *transition relation* \rightarrow_P over states:

<p style="margin: 0;">FIRING RULE</p> <p style="margin: 0;">APPLY</p> <p style="margin: 0; text-align: center;">$\langle H \setminus H' \Leftrightarrow G \mid B \rangle$ is a fresh variant of a rule in P</p> <p style="margin: 0; text-align: center;">with variables \mathbf{x} $CT \models \forall (b \rightarrow \exists \mathbf{x} (H \doteq h \wedge H' \doteq h' \wedge G))$</p> <hr style="width: 80%; margin: 0 auto;"/> <p style="margin: 0; text-align: center;">$\langle g; b; h, h', c \rangle_{V \rightarrow P}$</p> <p style="margin: 0; text-align: center;">$\langle B, g; H \doteq h \wedge H' \doteq h' \wedge G \wedge b; h, c \rangle_V$</p>

SOLVING RULES

$$\frac{\text{SOLVE} \quad B \in \mathcal{B} \quad CT \models B \wedge b \leftrightarrow b'}{\langle B, g; b; c \rangle_V \rightarrow_P \langle g; b'; c \rangle_V} \quad \frac{\text{INTRODUCE} \quad C \in \mathcal{U}}{\langle C, g; b; c \rangle_V \rightarrow_P \langle g; b; C, c \rangle_V}$$

Let q be an initial goal, the *query*. V is defined as $\text{fv}(q)$ and, from the *initial state* $s_0 = \langle q; \top; \emptyset \rangle_V$, a *derivation* is a sequence $s_0 \rightarrow_P s_1 \rightarrow_P \cdots \rightarrow_P s_n$. Such a state s_n is an *accessible state*.

Definition 3 (Linear Logic Semantics [10]).

For any built-in constraint $B = \langle B_1 \wedge \cdots \wedge B_n \rangle$,
let $B^\dagger = \langle !B_1 \otimes \cdots \otimes !B_n \rangle$.

For any CHR constraint $C = \langle C_1, \dots, C_n \rangle$,
let $C^\dagger = \langle C_1 \otimes \cdots \otimes C_n \rangle$.

For any goal $G = \langle G_1, \dots, G_n \rangle$,
let $G^\dagger = \langle G_1^\dagger \otimes \cdots \otimes G_n^\dagger \rangle$.

For any state $S = \langle g; b; c \rangle_V$, let $S^\dagger = \exists \mathbf{x} (g^\dagger \otimes b^\dagger \otimes c^\dagger)$,
where $\mathbf{x} = \text{fv}(G, B, C) \setminus V$.

The semantics of a rule $r = \langle H \setminus H' \Leftrightarrow G \mid B \rangle$ is
 $r^\dagger = \langle !\forall (G^\dagger \otimes H^\dagger \otimes H'^\dagger \multimap \exists \mathbf{x} (H^\dagger \otimes B^\dagger)) \rangle$

with $\mathbf{x} = \text{fv}(B) \setminus \text{fv}(H, H', G)$.

For a program $P = \{r_1, \dots, r_n\}$, the *linear logic semantics of P* is $P^\dagger = \langle r_1^\dagger \otimes \cdots \otimes r_n^\dagger \rangle$.

Theorem 3 (Soundness & Completeness [10]). *Let CT^\dagger be the Girard translation of CT [35], P a CHR program and q a query.*

- (Sound) *If s is an accessible state from q in P , then $P^\dagger, CT^\dagger \models \forall (q^\dagger \multimap s^\dagger)$.*
- (Complete) *For every formula c such that $P^\dagger, CT^\dagger \models \forall (q^\dagger \multimap c)$, there is an accessible state s from q in P such that $CT^\dagger \models \forall (s^\dagger \multimap c)$.*

2.1.2 Syntax and Semantics of LLCC

Definition 4 (Linear Constraint System [28]). A *linear constraint system* is a pair $(\mathcal{C}, \vdash_{\mathcal{C}})$, where:

- \mathcal{C} is a set of formulas (the *linear constraints*) built from variables \mathcal{V} and the signature Σ , with logical operators: multiplicative conjunction \otimes , its neutral 1, existential \exists , exponential $!$ and constant \top ; \mathcal{C} is assumed to be closed by renaming, multiplicative conjunction and existential quantification;
- $\vdash_{\mathcal{C}}$ is a binary relation over \mathcal{C} , which defines the non-logical axioms.

- $\vdash_{\mathcal{C}}$ is the least subset of $\mathcal{C}^* \times \mathcal{C}$ containing $\Vdash_{\mathcal{C}}$ and closed by the rules of intuitionistic multiplicative exponential linear logic for $1, \top, \otimes, !$ and \exists .

Definition 5 (Syntax with Persistent Asks [41]). The syntax for building LLCC agents follows the grammar: $A ::= \forall \mathcal{V}^*(\mathcal{C} \rightarrow A) \mid \forall \mathcal{V}^*(\mathcal{C} \Rightarrow A) \mid \exists \mathcal{V}.A \mid \mathcal{C} \mid A \parallel A$ where \parallel stands for parallel composition, \exists for variable hiding, \rightarrow for (transient) ask and \Rightarrow for persistent ask. In the particular case where there are no universally quantified variables in an ask, the notation $(c \rightarrow a)$ is preferred to $\forall \varepsilon(c \rightarrow a)$.

Agent $\forall \mathbf{x}(c \rightarrow a)$ suspends until c is entailed then wakes up and does a . Transient asks wake up at most one time. Persistent asks are introduced [41] to replace declarations by agents. The agent $\forall \mathbf{x}(c \Rightarrow a)$ can wake up as many times as c is entailed. This behavior makes sense as entailment consumes resources.

Example 2. Here is the LLCC version for dining philosophers [28, 38]. Compared to the CHR version, the following code is reentrant: the variable K identifies tokens and let several diners to be run in parallel (a banquet [38]) and separation results from the LLCC module theory prove that tables cannot steal cutlery from each other [41].

$$\begin{aligned} & \forall N(\text{diner}(N) \Rightarrow \\ & \quad \exists K(\forall I(\text{recphilo}(K, I) \Rightarrow \\ & \quad \quad \text{fork}(K, I) \parallel \\ & \quad \quad \exists J.(J \text{ is } (I + 1) \bmod N \parallel \\ & \quad \quad \quad (\text{fork}(K, I) \otimes \text{fork}(K, J) \Rightarrow \\ & \quad \quad \quad \text{eat}(K, I) \parallel \\ & \quad \quad \quad (\text{eat}(K, I) \rightarrow \\ & \quad \quad \quad \quad \text{fork}(K, I) \otimes \text{fork}(K, J)) \parallel \\ & \quad \quad \quad (I < N - 1 \rightarrow \text{recphilo}(K, J)) \parallel \\ & \quad \quad \text{recphilo}(K, 0)))))) \end{aligned}$$

This example makes use of non-trivial scopes: variables N, K, I and J are in turn introduced and shared by subsequent asks. The recursive loop $\langle \text{recphilo} \rangle$ installs N forks and composes N agents (the philosophers) in parallel. The philosopher between forks I and J is an agent in LLCC, whereas she is materialized in 1 by the CHR constraint $\text{philo}(I, J)$ in order to carry the environment $\{I, J\}$.

A configuration is a triple $(X; c; \Gamma)$ where c is a constraint (the store), Γ is a multi-set of agents and X is a set of variables (the hidden variables). The relation \equiv_{L} over configurations is the smallest equivalence relation such that:

- $(X; c; a \parallel b, \Gamma) \equiv_{\text{L}} (X; c; a, b, \Gamma)$ for all agents a and b ;
- $(X; c; 1, \Gamma) \equiv_{\text{L}} (X; c; \Gamma)$;

- $(X; c; \Gamma) \equiv_L (X; c'; \Gamma)$ for all constraints c, c' such that $c \dashv\vdash_c c'$;
- $(X; c; \Gamma) \equiv_L (X; c; \Gamma)[y/x]$ for all variables $x \in X$ and $y \notin \text{fv}(X, c, \Gamma)$

Let \mathcal{K} be the set of configurations.

Definition 6 (Operational Semantics [28, 41]). The *transition relation* \rightarrow_L is the least relation on configurations satisfying the following rules:

FIRING RULES			
<p style="margin: 0;">TRANSIENT ASK</p> $\frac{c \vdash_c \exists Y(d \otimes e[\mathbf{t}/\mathbf{x}]) \quad Y \cap \text{fv}(X, c, \Gamma) = \emptyset \quad \forall d'((c \vdash_c \exists Y(d' \otimes e[\mathbf{t}/\mathbf{x}])) \wedge (d' \vdash_c d) \Rightarrow d \dashv\vdash_c d')}{(X; c; \forall \mathbf{x}(e \rightarrow a), \Gamma) \rightarrow_L (X \cup Y; d; a[\mathbf{t}/\mathbf{x}], \Gamma)}$			
<p style="margin: 0;">PERSISTENT ASK</p> $\frac{c \vdash_c \exists Y(d \otimes e[\mathbf{t}/\mathbf{x}]) \quad Y \cap \text{fv}(X, c, \Gamma) = \emptyset \quad \forall d'((c \vdash_c \exists Y(d' \otimes e[\mathbf{t}/\mathbf{x}])) \wedge (d' \vdash_c d) \Rightarrow d \dashv\vdash_c d')}{(X; c; \forall \mathbf{x}(e \Rightarrow a), \Gamma) \rightarrow_L (X \cup Y; d; a[\mathbf{t}/\mathbf{x}], \forall \mathbf{x}(e \Rightarrow a), \Gamma)}$			
SOLVING RULES			
<p style="margin: 0;">HIDING</p> $\frac{y \notin X \cup \text{fv}(c, \Gamma)}{(X; c; \exists x.a, \Gamma) \rightarrow_L (X \cup \{y\}; c \otimes d; a[y/x], \Gamma)}$			
<p style="margin: 0;">TELL</p> $\frac{}{(X; c; d, \Gamma) \rightarrow_L (X; c \otimes d; \Gamma)}$			
<table style="width: 100%; border: none;"> <tr> <td style="padding: 5px; vertical-align: top;"> <p style="margin: 0;">EQUIVALENCE</p> $\frac{\kappa_0 \equiv_L \kappa'_0 \rightarrow_L \kappa'_1 \equiv_L \kappa_1}{\kappa_0 \rightarrow_L \kappa_1}$ </td> <td style="padding: 5px; vertical-align: top;"> <p style="margin: 0;">DECOMPOSITION</p> $\frac{}{(X; c; a \parallel b, \Gamma) \rightarrow_L (X; c; a, b, \Gamma)}$ </td> </tr> </table>		<p style="margin: 0;">EQUIVALENCE</p> $\frac{\kappa_0 \equiv_L \kappa'_0 \rightarrow_L \kappa'_1 \equiv_L \kappa_1}{\kappa_0 \rightarrow_L \kappa_1}$	<p style="margin: 0;">DECOMPOSITION</p> $\frac{}{(X; c; a \parallel b, \Gamma) \rightarrow_L (X; c; a, b, \Gamma)}$
<p style="margin: 0;">EQUIVALENCE</p> $\frac{\kappa_0 \equiv_L \kappa'_0 \rightarrow_L \kappa'_1 \equiv_L \kappa_1}{\kappa_0 \rightarrow_L \kappa_1}$	<p style="margin: 0;">DECOMPOSITION</p> $\frac{}{(X; c; a \parallel b, \Gamma) \rightarrow_L (X; c; a, b, \Gamma)}$		

An agent a is associated with the *initial configuration* $(\emptyset; \top; a)$. *Accessible observables* from a configuration κ are the configurations κ' such that $\kappa \xrightarrow{*}_L \kappa'$.

Definition 7 (Linear Logic Semantics [28, 41]). The translation $(\cdot)^\ddagger$ of LLCC agents into their *linear logic semantics* is defined inductively as follows:

$$\begin{aligned}
(\forall \mathbf{x}(c \rightarrow a))^\ddagger &= \forall \mathbf{x}(c \multimap a^\ddagger) & (\forall \mathbf{x}(c \Rightarrow a))^\ddagger &= !\forall \mathbf{x}(c \multimap a^\ddagger) & (\exists x.a)^\ddagger &= \exists x(a^\ddagger) \\
c^\ddagger &= c & (a \parallel b)^\ddagger &= a^\ddagger \otimes b^\ddagger
\end{aligned}$$

If Γ is a multi-set of agents (a_1, \dots, a_n) , we define $\Gamma^\ddagger = \langle a_1^\ddagger \otimes \dots \otimes a_n^\ddagger \rangle$. Configurations are translated to $(X; c; \Gamma)^\ddagger = \langle \exists X(c \otimes \Gamma^\ddagger) \rangle$.

The following theorem has first been proved by Fages, Ruet & Soliman [28]. The proof was done for a presentation of LLCC with declarations and without persistent asks. Rémy Haemmerlé showed in its Ph. D thesis that declarations can be internalized in the agents by the introduction of persistent asks: the work was presented in [41] and the full proof is in Rémy’s Ph. D, [38].

Theorem 4 (Soundness & Completeness [28, 41, 38]). *For all agents a :*

- (Sound) *If κ is an accessible observable from $(\emptyset; \top; a)$, then $a^\dagger \vdash_C \kappa^\dagger$.*
- (Complete) *If c is such that $a^\dagger \vdash_C c$, then there is an accessible observable $(X; d; \Gamma)$ from $(\emptyset; \top; a)$ with $\exists X(d) \vdash_C c$ and agents in Γ are persistent asks.*

2.1.3 Circumscribing non-determinism in CHR and LLCC operational semantics

Whereas non-determinism in firing rules seems to be inherent to the computation model (and is tackled in CHR by the committed-choice strategy and by the refined semantics), the non-determinism in sequencing solving rules can be completely eliminated. This is a classical result for constraint logic programming [44]. We formalize such a result for CHR and LLCC since the precise bisimulation results presented in next sections rely on it.

Let \rightarrow_P^s and \rightarrow_P^f be the restrictions of \rightarrow_P to solving and firing rules respectively. Let \rightarrow_L^s and \rightarrow_L^f be the similar restrictions for \rightarrow_L .

We define \Rightarrow_P^s such that $s \Rightarrow_P^s s'$ if and only if $s \xrightarrow{*}_P^s s' \not\rightarrow_P^s$. Similarly, \Rightarrow_L^s is such that $\kappa \Rightarrow_L^s \kappa'$ if and only if $\kappa \xrightarrow{*}_L^s \kappa' \not\rightarrow_L^s$.

Lemma 1 (Solving rules terminate and are confluent modulo \equiv). *For every CHR program P , for all state s , there exists s' such that $s \Rightarrow_P^s s'$ and for all s', s'' , if $s \Rightarrow_P^s s'$ and $s \Rightarrow_P^s s''$, then $s' \equiv_C s''$.*

For every configuration κ , there exists κ' such that $\kappa \Rightarrow_L^s \kappa'$ and for all κ', κ'' , if $\kappa \Rightarrow_L^s \kappa'$ and $\kappa \Rightarrow_L^s \kappa''$ then $\kappa' \equiv_L \kappa''$.

Thus, observed configurations can be restricted to be final for \rightarrow^s (or, equivalently, normalized by \Rightarrow^s) without losing derivations. The following lemma is a specialization of the “Andorra” principle [90] to the rule selection strategy:

Lemma 2 (Full solving before firing). *For every CHR program P ,*

$$\left(\xrightarrow{*}_P \cdot \Rightarrow_P^s \right) = \left((\Rightarrow_P^s \cdot \rightarrow_P^f)^* \cdot \Rightarrow_P^s \right)$$

and, similarly,

$$\left(\xrightarrow{*}_L \cdot \Rightarrow_L^s \right) = \left((\Rightarrow_L^s \cdot \rightarrow_L^f)^* \cdot \Rightarrow_L^s \right)$$

The lemma 2 is a corollary of the monotonous selection strategy [38]: intuitively, \rightarrow^s can always be exhausted before applying \rightarrow^f .

Lemma 3 (Solving rules preserve declarative semantics). *For every CHR program P , if $s \rightarrow_P^s s'$, then $s^\dagger \equiv s'^\dagger$. Similarly, if $\kappa \rightarrow_L^s \kappa'$, then $\kappa^\ddagger \equiv \kappa'^\ddagger$.*

Therefore, next sections focus on \Rightarrow -transitions where $\Rightarrow_P = (\Rightarrow_P^s \cdot \rightarrow_P^f \cdot \Rightarrow_P^s)$, and $\Rightarrow_L = (\Rightarrow_L^s \cdot \rightarrow_L^f \cdot \Rightarrow_L^s)$: a \Rightarrow_P -accessible state from s is a state s' such that $s \xrightarrow_P^* s'$ and a \Rightarrow_L -accessible observable from κ is a configuration κ' such that $\kappa \xrightarrow_L^* \kappa'$. It is worth noticing that a firing occurs at each \Rightarrow -transition.

2.2 Translations between sub-languages CSR and flat-LLCC

From now on, we consider the linear constraint system $(\mathcal{C}, \vdash_{\mathcal{C}})$ induced by the constraint theory CT and with atomic CHR constraints as linear tokens. More precisely, \mathcal{C} is the least set of formulas which contains \top and $!B$ for all $B \in \mathcal{B}_0$ and C for all $C \in \mathcal{U}_0$, closed by renaming, multiplicative conjunction and existential quantification. We suppose that $c \Vdash_{\mathcal{C}} d$ if and only if $CT^\dagger \models \forall(c \multimap d)$. The result is a particular form of linear constraint system where non-logical axioms follow from the translation of a classical theory.

Bisimulation is the most popular method for comparing concurrent processes [71], characterizing a notion of strong equivalence between processes. A *transition system* is a tuple (S, \rightarrow) with S a set of states and \rightarrow a binary relation over S . We define the CHR transition system as $(\mathcal{P}, \Rightarrow_{\mathcal{C}})$ where $(P, s) \Rightarrow_{\mathcal{C}} (P', s')$ when $P = P'$ and $s \Rightarrow_P s'$, and the LLCC transition system as $(\mathcal{K}, \Rightarrow_L)$.

Definition 8 (Bisimulation). Let $(S_1, \xrightarrow{1})$ and $(S_2, \xrightarrow{2})$ be two transition systems. A *bisimulation* is a relation $\sim \subseteq S_1 \times S_2$ such that for all s_1, s_2 such that $s_1 \sim s_2$:

- for all s'_1 such that $s_1 \xrightarrow{1} s'_1$, there exists s'_2 such that $s_2 \xrightarrow{2} s'_2$ and $s'_1 \sim s'_2$;
- for all s'_2 such that $s_2 \xrightarrow{2} s'_2$, there exists s'_1 such that $s_1 \xrightarrow{1} s'_1$ and $s'_1 \sim s'_2$.

2.2.1 From Constraint Simplification Rules (CSR) to flat-LLCC

Resulting configurations of LLCC FIRING RULES enjoy a new store where guards have been consumed. This behavior corresponds to simplification rules in CHR.

Definition 9 (CSR programs[3]). A CHR program P is a CSR program when all rules of P are simplifications (*i.e.* rules are of the form $\langle H \Leftrightarrow G \mid B. \rangle$).

As far as naive operational semantics and linear-logic semantics are concerned, expressiveness of CHR and CSR is identical. For a rule $r = \langle H \setminus H' \Leftrightarrow G \mid B \rangle$, let $r^\times = \langle H, H' \Leftrightarrow G \mid H, B. \rangle$ and for $P = \{r_1, \dots, r_n\}$, let $P^\times = \{r_1^\times, \dots, r_n^\times\}$.

Example 3. Here is leq^\times translated from a version of the leq program [78]:

$$\begin{aligned} leq(X, X) &\Leftrightarrow true. \\ leq(X, Y) &\Leftrightarrow number(X), number(Y) \mid X \leq Y. \\ leq(X, Y), leq(Y, X) &\Leftrightarrow X = Y. \\ leq(X, Y), leq(Y, Z) &\Leftrightarrow leq(X, Y), leq(Y, Z), leq(X, Z). \\ leq(X, Y), leq(X, Y) &\Leftrightarrow leq(X, Y). \end{aligned}$$

Proposition 1 (CHR and CSR equivalence). *For every CHR program P , we have $\rightarrow_P = \rightarrow_{P^\times}$ and $P^\dagger \equiv (P^\times)^\dagger$.*

This equivalence only holds for naive CHR semantics. There is probably no natural encoding of the traditional semantics for propagation [25] in LLCC, at least without *ad-hoc* support hard-wired in the constraint system.

Let $r = \langle H' \Leftrightarrow G \mid B. \rangle$ be a simplification rule. $G^\dagger \otimes H'^\dagger$ and B^\dagger are in \mathcal{C} , thus the following agent is well-formed: $r^{-\circ} = \langle \forall \mathbf{y}(G^\dagger \otimes H'^\dagger \Rightarrow \exists \mathbf{x}. B^\dagger) \rangle$, where $\mathbf{x} = \text{fv}(B) \setminus \text{fv}(H', G)$ and $\mathbf{y} = \text{fv}(H', G)$. For every CSR program $P = \{r_1, \dots, r_n\}$, the translation of P in LLCC is: $P^{-\circ} = \langle r_1^{-\circ} \parallel \dots \parallel r_n^{-\circ} \rangle$. States $\langle g; b; c \rangle_V$ are translated in \mathcal{C} as well: $\langle g; b; c \rangle_V^{-\circ} = g^\dagger \otimes b^\dagger \otimes c^\dagger$.

Definition 10 (CSR to LLCC translation). A CSR program P and a query q are translated to the agent $a(P, q) = \langle P^{-\circ} \parallel q^\dagger \rangle$.

Example 4. The leq^\times program (Example 3) is translated to the agent $leq^{-\circ}$:

$$\begin{aligned} leq^{-\circ} = &\forall X (leq(X, X) \Rightarrow 1) \parallel \\ &\forall XY (number(X) \otimes number(Y) \otimes leq(X, Y) \Rightarrow \\ &\quad X \leq Y) \parallel \\ &\forall XY (leq(X, Y) \otimes leq(Y, X) \Rightarrow X = Y) \parallel \\ &\forall XYZ (leq(X, Y) \otimes leq(Y, Z) \Rightarrow \\ &\quad leq(X, Y) \otimes leq(Y, Z) \otimes leq(X, Z)) \parallel \\ &\forall XY (leq(X, Y) \otimes leq(X, Y) \Rightarrow leq(X, Y)) \end{aligned}$$

Since there is no possible confusion between linear tokens and classical constraints, then, by abuse of notations, we omit the $!$ operator on \mathcal{U}_0 constraints.

Main Result 1 (Bisimilarity). *Let $\sim \subseteq \mathcal{P} \times \mathcal{K}$ be the relation where $(P, s) \sim \kappa$ if and only if $\kappa \equiv_{\text{L}} (X; s^{-\circ}; P^{-\circ})$ with $X = \text{fv}(s) \setminus V$. Then, \sim is a bisimulation.*

Proof. Straightforward for \rightarrow_P^s and \Rightarrow_L^s .

Let $(P, s) \sim \kappa$ and $s \rightarrow_P^f s' \Rightarrow_P^s s''$. Let $r = \langle H \Leftrightarrow G \mid B \rangle$ be the fired rule: then there exists a sub-multiset of CHR constraints in s and a substitution such that H is matched and G is entailed. Therefore the ask resulting from the translation of r can be fired and $\kappa \rightarrow_L^f \kappa' \Rightarrow_L^s \kappa''$ with $(P, s'') \sim \kappa''$.

Let $(P, s) \sim \kappa$ and $\kappa \rightarrow_L^f s' \Rightarrow_L^s \kappa''$. Then there exists a rule in P which the fired ask is translated from and this rule can be fired from s . Since the store has not been weakened in κ' , we have $(P, s'') \sim \kappa''$. \square

Corollary 1 (Semantics preservation). *For CSR program P , query q :*

- *if κ is a \Rightarrow_L -accessible observable of $a(P, q)$, then $\kappa \equiv (X; c; P^\circ)$ and there is a \Rightarrow_P -accessible state s from q with $\exists \mathbf{x}(s^\circ) \dashv\vdash_c \exists X(c)$, $\mathbf{x} = \text{fv}(s) \setminus \text{fv}(q)$;*
- *if s is a \Rightarrow_P -accessible state from q , then there is a \Rightarrow_L -accessible observable $(X; c; P^\circ)$ from $a(P, q)$ such that $\exists \mathbf{x}(s^\circ) \dashv\vdash_c \exists X(c)$, where $\mathbf{x} = \text{fv}(s) \setminus \text{fv}(q)$.*

2.2.2 From flat-LLCC to CSR

The translation of CSR into LLCC generates agents of the particular form $p \parallel q$, where the sub-agent p is the translation of a CSR program and is therefore a parallel composition of persistent asks without any nested asks, and the sub-agent q is a translation of a query and is therefore reduced to a constraint. Moreover, every ask guard consumes at least a linear token (since CHR heads are non-empty) and asks are closed term (*i.e.* without free variables). Such agents are characterized by the following definition:

Definition 11 (flat-LLCC). *Flat-LLCC agents are restricted to the grammar: $A^\dagger ::= A^\forall \parallel \mathcal{C}$ where $A^\forall ::= \forall \mathcal{V}^*(\mathcal{C} \Rightarrow \mathcal{C}) \mid A^\forall \parallel A^\forall \mid 1$ with the following side condition for every ask $\forall \mathbf{x}(g \Rightarrow c)$: $g \not\vdash_c g \otimes g$ (consumption) and $\text{fv}(g, c) \subseteq \mathbf{x}$.*

This subsection is dedicated to establishing the reverse translation, from A^\dagger to CSR. It is worth noticing first that, like a CSR program, an A^\dagger -agent essentially transforms constraint stores without introducing new suspensions:

Lemma 4 (Configurations form). *Non-initial \Rightarrow_L -accessible configurations from an A^\dagger -agent a are \equiv_L -equivalent to configurations of the form $(_ ; _ ; a^\forall)$.*

The translation from flat-LLCC to CSR should handle the LLCC existential variables which have no counter part in CSR and the splitting between built-in constraints and CHR constraints. Fresh variables should be introduced to translate constraints such as $a(X, Y) \otimes \exists X(b(X, Y))$ into $\langle a(X, Y), b(K, Y) \rangle$ where K is a new local variable. The function f^c translates every constraint in \mathcal{C} to a tuple

$(X; B; C)$ where B is a built-in constraint, C a CHR constraint and X a set of variables local to B and C :

$$\begin{aligned}
f^c(\top) &= (\emptyset; \mathbf{true}; \emptyset) \\
f^c(!B) &= (\emptyset; B; \emptyset) \\
&\quad \text{for all } B \in \mathcal{B}_0 \\
f^c(C) &= (\emptyset; \mathbf{true}; C) \\
&\quad \text{for all } C \in \mathcal{U}_0 \\
f^c(c \otimes d) &= (\sigma_c(X_c) \cup \sigma_d(X_d); \sigma_c(B_c) \wedge \sigma_d(B_d); \\
&\quad \sigma_c(C_c), \sigma_d(C_d)) \\
&\quad \text{where } f^c(c) = (X_c; B_c; C_c) \\
&\quad \text{and } f^c(d) = (X_d; B_d; C_d) \\
&\quad \text{with } \sigma_c \text{ and } \sigma_d \text{ renaming of } X_c \text{ and} \\
&\quad X_d \text{ respectively such that} \\
&\quad \sigma_c(X_c) \cap \text{fv}(\sigma_d(B_d, C_d)) = \emptyset \text{ and} \\
&\quad \sigma_d(X_d) \cap \text{fv}(\sigma_c(B_c, C_c)) = \emptyset \\
f^c(\exists x(c)) &= (X_c \cup \{x\}; B_c; C_c) \\
&\quad \text{where } f^c(c) = (X_c; B_c; C_c)
\end{aligned}$$

A^\forall -agents are translated to CSR programs through the function f^\forall . Translation of asks should take care of clashes with similar renaming as for \otimes in f^c :

$$\begin{aligned}
f^\forall(\forall \mathbf{x}(g \Rightarrow c)) &= \\
&\quad \{\langle \sigma_g(C_g) \Leftrightarrow \sigma_g(B_g) \mid \sigma_c(B_c), \sigma_c(C_c) \rangle\} \\
&\quad \text{where } f^c(g) = (X_g; B_g; C_g) \\
&\quad \text{and } f^c(c) = (X_c; B_c; C_c) \\
&\quad \text{and } \sigma_g \text{ and } \sigma_c \text{ renaming of } X_g \text{ and} \\
&\quad X_c \text{ respectively such that} \\
&\quad \sigma_g(X_g) \cap \text{fv}(\sigma_c(B_c, C_c)) = \emptyset \text{ and} \\
&\quad \sigma_c(X_c) \cap \text{fv}(\sigma_g(B_g, C_g)) = \emptyset \\
f^\forall(a \parallel b) &= f^\forall(a) \cup f^\forall(b) \\
f^\forall(1) &= \emptyset
\end{aligned}$$

For every ask $\forall \mathbf{x}(g \Rightarrow c)$, $f^\forall(\forall \mathbf{x}(g \Rightarrow c))$ is a well-formed CHR rule. In particular, the side condition on g ensures that $\sigma_g(C_g) \neq \emptyset$.

$f_V^s : c \mapsto \langle \emptyset; b; c \rangle_V$ maps constraints to states with $(_; b; c) = f^c(c)$.

Note that all variables in CSR queries are global. The query should hide existentially quantified variables in the top-level constraint c_0 of the agent. We suppose a fresh symbol $\langle \text{start}/n \rangle \in \mathcal{U}_0$ where $n = \#\text{fv}(c_0)$.

Definition 12 (Flat-LLCC to CSR translation). A flat-LLCC agent $\langle a^\forall \parallel c_0 \rangle$ is translated to the CHR program $P(a^\forall \parallel c_0) = f^\forall(a^\forall) \cup \{\text{start}(\mathbf{v}) \Leftrightarrow B_0, C_0.\}$ and the query $q(a) = (\text{start}(\mathbf{v}))$ where $(_; B_0; C_0) = f^c(c_0)$ and $\mathbf{v} = \text{fv}(c_0)$.

Main Result 2 (Bisimilarity). *Let $\sim \subseteq \mathcal{K} \times \mathcal{P}$ be the relation where $\kappa \sim (P, s)$ if and only if there exists a flat-LLCC agent $\langle a^\forall \parallel c_0 \rangle$ where $\kappa \equiv_L (X; c; a^\forall)$ and $P = P(a)$ and $s \equiv_C f_V^s(c)$, with $V = \text{fv}(c_0)$. Then, \sim is a bisimulation.*

Proof. Let $\kappa \sim (P, s)$ and $\kappa \xrightarrow{f}_L \kappa' \Rightarrow_L^s \kappa''$. Then $\kappa \equiv_L (X; d; a^\forall)$ and there is an ask $\forall x(e \Rightarrow d')$ in a^\forall such that $d \vdash_C e[\mathbf{x}/\mathbf{t}] \otimes d''$ and $\kappa' \equiv_L (X; d''; d', a^\forall)$. Then $\kappa'' = (X; d' \otimes d''; a^\forall)$. The rule $f^\forall(\forall x(e \Rightarrow d'))$ is in P . Let $\langle \emptyset; b; c \rangle_V = s$, $(_; G; H') = f^c(e)$ and $(_; g; h') = f^c(e[\mathbf{x}/\mathbf{t}])$. Since $s = f^c(d)$, $CT \models \forall(b \rightarrow g)$ and there exists c such that $c' = (h', c)$. As (G, H') and (g, h') differ by a substitution on \mathbf{x} , $\exists \mathbf{x}(H' = h' \wedge G)$. Therefore $s \xrightarrow{f}_P s'$ and $s' = \langle b'', c''; H' = h' \wedge G \wedge b; c \rangle_V$ with $(b'', c'') = f^c(d')$ and $s' \Rightarrow_P^s s''$ where $s'' = \langle \emptyset; b'' \wedge H' = h' \wedge G \wedge b; c'' \rangle_V$. Since \mathbf{t} and d'' are the most general choices, the equivalence between s'' and $f^c(d' \otimes d'')$ holds.

Let $\kappa \sim (P, s)$ and $s \xrightarrow{f}_P s' \Rightarrow_P^s s''$. Then there exists an ask $\forall x(e \Rightarrow d')$ in a^\forall which the fired rule is translated from. Side-conditions of CHR firing rules imply the existence of a substitution for e such that the store d of κ is such that $d \vdash_C e[\mathbf{x}/\mathbf{t}] \otimes d''$. Then there exists κ' such that $\kappa \xrightarrow{f}_L \kappa'$ where κ' has the store d'' . Then $\kappa' \Rightarrow_L^s \kappa''$ and $\kappa'' \sim s''$. \square

Corollary 2 (Semantics preservation). *For every flat-LLCC agent $a = \langle a^\forall \parallel c_0 \rangle$, let $s_0 = \langle q(a); \top; \emptyset \rangle_V$, $V = \text{fv}(c_0)$, then:*

- for all \Rightarrow_L -accessible configuration $(X; c; a^\forall)$ from a , there exists a $\Rightarrow_{P(a)}$ -accessible state s from s_0 such that $\exists \mathbf{x}(s^{-\circ}) \dashv\vdash_C \exists X(c)$;
- for all $\Rightarrow_{P(a)}$ -accessible state s from s_0 , if $s \neq s_0$, there exists a \Rightarrow_L -accessible configuration $(X; c; a^\forall)$ from a , such that $\exists \mathbf{x}(s^{-\circ}) \dashv\vdash_C \exists X(c)$;

where, in both cases, $\mathbf{x} = \text{fv}(s) \setminus V$.

2.2.3 CHR Linear-Logic and Phase Semantics Revisited

Lemma 5 (Identical Semantics). *For every CSR program P and query q , $(P^{-\circ})^\ddagger \equiv P^\ddagger$ and we have $P^\ddagger, CT^\ddagger \models \forall(q^\ddagger \multimap c)$ if and only if $(P^{-\circ})^\ddagger, q^\ddagger \vdash_C c$.*

Relating 3 and 4 supposes to prove that accessible constraints are included in provable constraints (correctness), and conversely (completeness). Thus, the correctness and completeness result amounts to equality between sets, which we

make explicit here to prove both ways at the same time.

$$\begin{aligned}
\mathcal{O}_C(P, q) &= \{(P, s) \in \mathcal{P} \mid (q; \top; \emptyset) \xrightarrow{*}_P s\} & \mathcal{O}_L(a) &= \{\kappa \in \mathcal{K} \mid (\emptyset; \top; a) \xrightarrow{*}_L \kappa\} \\
\mathcal{O}_C^{\Rightarrow}(P, q) &= \{(P, s) \in \mathcal{P} \mid (q; \top; \emptyset) \xrightarrow{\Rightarrow}_P s\} & \mathcal{O}_L^{\Rightarrow}(a) &= \{\kappa \in \mathcal{K} \mid (\emptyset; \top; a) \xrightarrow{\Rightarrow}_L \kappa\} \\
\mathcal{LL}_C(P, q) &= \{c \in \mathcal{C} \mid P^\dagger, CT^\dagger \models \forall(q^\dagger \multimap c)\} & \Downarrow_C^s S &= \{s \in \mathcal{P} \mid \exists s' \in S \Rightarrow_C^s s\} \\
\mathcal{LL}_L(a) &= \{c \in \mathcal{C} \mid a^\dagger \vdash_C c\} & \Downarrow_L^s S &= \{\kappa \in \mathcal{K} \mid \exists \kappa' \in S \Rightarrow_L^s \kappa\} \\
\downarrow S &= \{c \in \mathcal{C} \mid \exists c' \in S, c' \vdash_C c\}
\end{aligned}$$

Some results mentioned up to now are summarized in the following table:

For every CSR program P , query q , and flat-LLCC agent a ,	
– 3:	$\downarrow(\mathcal{O}_C(P, q))^\dagger = \mathcal{LL}_C(P, q)$;
– 4:	$\downarrow(\mathcal{O}_L(a))^\ddagger = \mathcal{LL}_L(a)$;
– 2:	$\Downarrow_C^s \mathcal{O}_C(P, q) = \mathcal{O}_C^{\Rightarrow}(P, q)$ and $\Downarrow_L^s \mathcal{O}_L(a) = \mathcal{O}_L^{\Rightarrow}(a)$;
– 3:	$(\Downarrow_C^s \mathcal{O}_C(P, q))^\dagger = (\mathcal{O}_C(P, q))^\dagger$ and $(\Downarrow_L^s \mathcal{O}_L(a))^\ddagger = (\mathcal{O}_L(a))^\ddagger$;
– 1:	$\mathcal{O}_C(P, q) = \mathcal{O}_C(P^\times, q)$ and $\mathcal{LL}_C(P, q) = \mathcal{LL}_C(P^\times, q)$;
– 1:	$(\mathcal{O}_C^{\Rightarrow}(P^\times, q))^\dagger = (\mathcal{O}_L^{\Rightarrow}(a(P^\times, q)))^\ddagger$;
– 5:	$\mathcal{LL}_C(P^\times, q) = \mathcal{LL}_L(a(P^\times, q))$

We are now ready to prove 3 again from the other results.

Proof of 3. For every CHR program P and query q :

$$\begin{aligned}
\downarrow(\mathcal{O}_C(P, q))^\dagger &\stackrel{1}{=} \downarrow(\mathcal{O}_C(P^\times, q))^\dagger &&\stackrel{3}{=} \downarrow(\Downarrow_C^s \mathcal{O}_C(P^\times, q))^\dagger \\
&\stackrel{2}{=} \downarrow(\mathcal{O}_C^{\Rightarrow}(P^\times, q))^\dagger &&\stackrel{1}{=} \downarrow(\mathcal{O}_L^{\Rightarrow}(a(P^\times, q)))^\ddagger \\
&\stackrel{2}{=} \downarrow(\Downarrow_L^s \mathcal{O}_L(a(P^\times, q)))^\ddagger &&\stackrel{3}{=} \downarrow(\mathcal{O}_L(a(P^\times, q)))^\ddagger \\
&\stackrel{4}{=} \mathcal{LL}_L(a(P^\times, q)) &&\stackrel{5}{=} \mathcal{LL}_C(P^\times, q) \\
&\stackrel{1}{=} \mathcal{LL}_C(P, q) && \square
\end{aligned}$$

The following proposition describes a method to prove unreachability property in CHR using phase semantics, adapted from similar result in LLCC [28].

Proposition 2 (Safety through Phase Semantics [39]). *To prove a safety property of the kind $s \not\rightarrow_P s'$ for a given CHR program P , it is enough to prove that for a well-chosen phase space \mathbf{P} and valuation η compatible with CT and P , there exists an element $a \in \eta(s^\dagger)$ such that $a \notin \eta(t^\dagger)$.*

Such a valuation η is compatible with \vdash_C and $(P^\times)^\circ$ (note that it is immediate to see $(P^\times)^\circ$ as declarations in the sense of the original presentation of LLCC [28]). Let κ and κ' be the respective images of s and s' by the transformation. Then the element a is such that $a \in \eta(\kappa^\dagger)$ and $a \notin \eta(\kappa'^\dagger)$. Thus $\kappa \not\rightarrow_L \kappa'$ comes from the phase semantics of LLCC. Therefore the property $s \not\rightarrow_{P^\times} s'$ follows from 1, and is generalizable to P with 1. That proves 2. \square

2.3 Ask-lifting: encoding LLCC into CSR

The main result of this section is a translation from LLCC to flat-LLCC which preserves the semantics. Consequently, thanks to 2, we can deduce a semantics-preserving translation from LLCC to CSR. This section begins with a preliminary step introducing an intermediary language LLCC^ℓ where asks are labeled with linear tokens: these tokens do not change the operational semantics and there is a trivial labeling to transform LLCC programs to LLCC^ℓ programs. These linear tokens are introduced in order to follow asks through the transitions of the operational semantics, which is used to prove the semantics preservation.

2.3.1 Preliminary step: labeling LLCC-agents

Labeled LLCC agents A^ℓ differ from agents A by labels inserted on each ask. In the following definition, labels are arbitrary linear tokens.

Definition 13 (LLCC^ℓ agents). The syntax of LLCC^ℓ agents is given by the following grammar:

$$A^\ell ::= \forall \mathcal{V}^*(\mathcal{C} \xrightarrow{\mathcal{U}_0} A^\ell) \mid \forall \mathcal{V}^*(\mathcal{C} \xRightarrow{\mathcal{U}_0} A^\ell) \mid \\ \exists \mathcal{V}. A^\ell \mid \mathcal{C} \mid A^\ell \parallel A^\ell$$

The transition relation \rightarrow_L is lifted to the transition $\rightarrow_{\text{LLCC}^\ell}$ for LLCC^ℓ .

$$\begin{array}{l} \text{TRANSIENT ASK (WITH LABELING)} \\ \frac{c \vdash_c \exists Y(d \otimes e[\mathbf{t}/\mathbf{x}]) \quad Y \cap \text{fv}(X, c, \Gamma) = \emptyset \\ \forall d'((c \vdash_c \exists Y(d' \otimes e[\mathbf{t}/\mathbf{x}]))) \wedge (d' \vdash_c d) \Rightarrow d \dashv\vdash d'}{(X; c; \forall \mathbf{x}(e \xrightarrow{l} a), \Gamma) \rightarrow_{\text{LLCC}^\ell} (X \cup Y; d; a[\mathbf{t}/\mathbf{x}], \Gamma)} \\ \\ \text{PERSISTENT ASK (WITH LABELING)} \\ \frac{c \vdash_c \exists Y(d \otimes e[\mathbf{t}/\mathbf{x}]) \quad Y \cap \text{fv}(X, c, \Gamma) = \emptyset \\ \forall d'((c \vdash_c \exists Y(d' \otimes e[\mathbf{t}/\mathbf{x}]))) \wedge (d' \vdash_c d) \Rightarrow d \dashv\vdash d'}{(X; c; \forall \mathbf{x}(e \xRightarrow{l} a), \Gamma) \rightarrow_{\text{LLCC}^\ell} (X \cup Y; d; a[\mathbf{t}/\mathbf{x}], \forall \mathbf{x}(e \xRightarrow{l} a), \Gamma)} \end{array}$$

Agents A are translated to a particular family of labeled agents denoted A^{ℓ_0} with the labeling transformation, which ensures the following conditions: each label carries a distinct symbol taken from a set \mathcal{P} of fresh predicate symbols and the arguments enumerate exactly the free variables of the ask. Such a labeling is simple to obtain as soon as \mathcal{P} is large enough to label each ask of a .

Example 5. *The dining philosophers (2) can be labeled as follows:*

$$\begin{aligned}
& \forall N(\text{diner}(N) \xrightarrow{p_1} \\
& \quad \exists K(\forall I(\text{recphilo}(K, I) \xrightarrow{p_2(K, N)} \\
& \quad \quad \text{fork}(K, I) \parallel \\
& \quad \quad \exists J.(J \text{ is } (I + 1) \bmod N \parallel \\
& \quad \quad \quad (\text{fork}(K, I) \otimes \text{fork}(K, J) \xrightarrow{p_3(I, J, K)} \\
& \quad \quad \quad \text{eat}(K, I) \parallel \\
& \quad \quad \quad (\text{eat}(K, I) \xrightarrow{p_4(I, J, K)} \\
& \quad \quad \quad \text{fork}(K, I) \otimes \text{fork}(K, J)) \parallel \\
& \quad \quad \quad (I < N - 1 \xrightarrow{p_5(I, J, K, N)} \\
& \quad \quad \quad \text{recphilo}(K, J)) \parallel \\
& \quad \quad \text{recphilo}(K, 0))))))
\end{aligned}$$

2.3.2 The ask-lifting transformation

The ask-lifting transformation is defined with two helper functions. $\langle a \rangle^c$ transforms the agent a to constraints where asks become linear tokens. $\langle a \rangle^\forall$ puts in parallel every ask occurring in a and the representing token is added to the guard. A persistent ask restores the token, a transient ask consumes it.

The function $\langle \cdot \rangle^c : A^\ell \rightarrow \mathcal{C}$ is defined inductively as follows:

$$\begin{aligned}
\langle \forall \mathbf{x}(c \xrightarrow{f(\mathbf{t})} a) \rangle^c &= f(\mathbf{t}) & \langle \forall \mathbf{x}(c \xrightarrow{f(\mathbf{t})} a) \rangle^c &= f(\mathbf{t}) & \langle \exists x.a \rangle^c &= \exists x.\langle a \rangle^c \\
\langle a \parallel b \rangle^c &= \langle a \rangle^c \otimes \langle b \rangle^c & \langle c \rangle^c &= c
\end{aligned}$$

The function $\langle \cdot \rangle^\forall : A^{\ell_0} \rightarrow A^\forall$ is defined inductively as follows:

$$\begin{aligned}
\langle \forall \mathbf{x}(c \xrightarrow{f(\mathbf{v})} a) \rangle^\forall &= \forall \mathbf{v} \mathbf{x}(f(\mathbf{v}) \otimes c \xrightarrow{f(\mathbf{v})} \langle a \rangle^c) \parallel \langle a \rangle^\forall \\
\langle \forall \mathbf{x}(c \xrightarrow{f(\mathbf{v})} a) \rangle^\forall &= \forall \mathbf{v} \mathbf{x}(f(\mathbf{v}) \otimes c \xrightarrow{f(\mathbf{v})} f(\mathbf{v}) \otimes \langle a \rangle^c) \parallel \langle a \rangle^\forall \\
\langle \exists x.a \rangle^\forall &= \langle a \rangle^\forall & \langle a \parallel b \rangle^\forall &= \langle a \rangle^\forall \parallel \langle b \rangle^\forall & \langle c \rangle^\forall &= 1
\end{aligned}$$

The function $\langle \cdot \rangle^\forall$ is well-defined: every ask satisfies the side-condition for A^\forall .

Definition 14 (Ask-lifting). The *agent ask-lifting function* $\llbracket \cdot \rrbracket^\uparrow : A \rightarrow A^\uparrow$ transforms the agent a to the agent $\llbracket a \rrbracket^\uparrow = \langle a^\ell \rangle^\forall \parallel \langle a^\ell \rangle^c$ where a is translated to a^ℓ by the labeling defined in 2.3.1 with symbol predicates from a subset \mathcal{P} of \mathcal{P}_c whose predicates do not appear in a . $\llbracket \cdot \rrbracket^\uparrow$ is well-defined as soon as the set \mathcal{P} is large enough to label agent a .

Main Result 3 (Bisimilarity). *Let a be a labeled LLCC agent. Let $\sim \subseteq \mathcal{K} \times \mathcal{K}$ be the relation such that $\kappa \sim \kappa'$ if and only if $\kappa \equiv_{\text{L}} (X; c; \Gamma)$ is \Rightarrow -accessible from a and $\kappa' \equiv_{\text{L}} (X; c \otimes \langle \Gamma \rangle^c; \langle a \rangle^\forall)$. Then, \sim is a bisimilarity.*

Corollary 3 (Semantics preservation). *For every LLCC agent a :*

- *for all \Rightarrow_{L} -accessible configuration $(X; c; \Gamma)$ from a , there is a \Rightarrow_{L} -accessible configuration $(X; c'; \langle a \rangle^\forall)$ from $\llbracket a \rrbracket^\uparrow$ such that $\exists X(c \otimes \langle \Gamma \rangle^c) \dashv\vdash_c \exists X'(c')$;*
- *for all \Rightarrow_{L} -accessible configuration $(X; c'; \langle a \rangle^\forall)$ from $\llbracket a \rrbracket^\uparrow$, there exists a \Rightarrow_{L} -accessible configuration $(X; c; \Gamma)$ from a and $\exists X(c \otimes \langle \Gamma \rangle^c) \dashv\vdash_c \exists X'(c')$.*

Example 6. *The labeled diner (5) can be lifted as follows:*

$$\begin{aligned}
& \forall N(\quad p_1 \otimes \text{diner}(N) \Rightarrow \\
& \quad p_1 \otimes p_2(K, N) \otimes \text{recphilo}(K, 0)) \parallel \\
& \forall IK N(\quad p_2(K, N) \otimes \text{recphilo}(K, I) \Rightarrow \\
& \quad p_2(K, N) \otimes \text{fork}(K, I) \otimes \\
& \quad \exists J(J \text{ is } (I + 1) \bmod N \otimes p_3(I, J, K) \otimes \\
& \quad \quad p_5(I, J, K, N))) \parallel \\
& \forall IJK(\quad p_3(I, J, K) \otimes \text{fork}(K, I) \otimes \text{fork}(K, J) \Rightarrow \\
& \quad p_3(I, J, K) \otimes \text{eat}(K, I) \otimes p_4(I, J, K)) \parallel \\
& \forall IJK(\quad p_4(I, J, K) \otimes \text{eat}(K, I) \Rightarrow \\
& \quad \text{fork}(K, I) \otimes \text{fork}(K, J)) \parallel \\
& \forall IJK N(p_5(I, J, K, N) \otimes I < N - 1 \Rightarrow \text{recphilo}(K, J)) \parallel \\
& p_1
\end{aligned}$$

2.4 Encoding the call-by-value λ -calculus in CHR

The following transformation from pure λ -terms to LLCC is proved correct and complete with respect to the call-by-value semantics of the λ -calculus [41]. Every function, *aka* λ -value, is represented by a variable K . The constraint $\text{apply}(K, X, V)$ represents that V should code the result of the application of the function (coded by) K to the λ -term (coded by) X . Therefore, the transformation of a λ -abstraction $\lambda x.e$ coded by K should be a persistent ask which transforms, for all X and V , the constraint $\text{apply}(K, X, V)$ to the equality constraint between V and the evaluation of $e[t/x]$, where t is the λ -term coded by X . The equality constraint is put at the level of λ -variables. The constraint $\text{value}(K)$ indicates that the λ -term K has been reduced to a value so as to encode the particular call-by-value strategy [68].

Definition 15 (Call-by-value λ -calculus in LLCC [41]). For every λ -term e , $\llbracket e \rrbracket$ is a function from variables to LLCC agents. $\llbracket e \rrbracket$ is described inductively on the structure of e :

- $\llbracket X \rrbracket(K) = \langle X = K \otimes \text{value}(K) \rangle$
- $\llbracket \lambda X.e \rrbracket(K) =$
 $\forall XV(\text{apply}(K, X, V) \otimes \text{value}(X) \Rightarrow$
 $\llbracket e \rrbracket(V) \parallel \text{value}(X))$
- $\llbracket f e \rrbracket(K) = \exists XY(\text{apply}(X, Y, K) \parallel \llbracket f \rrbracket(X) \parallel \llbracket e \rrbracket(Y))$

Each ask introduced by this transformation corresponds to a λ -abstraction and this property is preserved by ask-lifting. Therefore, the CSR program obtained by translation has one rule for each λ -abstraction.

We explicit below the direct transformation from λ -terms to CSR. We suppose that the labeling has been prepared directly in λ -terms: λ -abstractions are of the form $\lambda_i X.e$ where i is a unique index.

Definition 16 (Call-by-value λ -calculus in CHR). For every λ -term e , $[e]$ is a function from variables to pairs CHR programs and queries, each component being denoted $[e]^p$ and $[e]^g$. $[e]$ is described inductively on the structure of e as follows.

- $[X](K) = (\emptyset; (X = K, \text{value}(K)))$
- $[\lambda X_i.e](K) = ([e]^p(V) \cup \{r\}; p_i(K, \mathbf{v}))$
where $\mathbf{v} = \text{fv}(\lambda X.e)$ and
 X and V are fresh variables and
 $r = \langle p_i(K, \mathbf{v}), \text{value}(X), \text{apply}(K, X, V) \Leftrightarrow$
 $p_i(K, \mathbf{v}), \text{value}(X), [e]^g(V). \rangle$
- $[f e](K) = ([f]^p(X) \cup [f]^p(Y);$
 $([f]^g(X), [e]^g(Y), \text{apply}(X, Y, K)))$
where X and Y fresh variables

The $p_i(\mathbf{v})$ CHR constraints are supposed to be fresh. Then, the CSR program associated to e is $P[e] = [e]^p(R) \cup \{\langle \text{start}(R, \mathbf{v}) \Leftrightarrow [e]^g(R). \rangle\}$ and the query is $q[e] = \text{start}(R, \mathbf{v})$ with $\mathbf{v} = \text{fv}(e)$.

It is immediate that the program and the goal produced by the transformation above correspond syntactically to the composition of the three transformations: λ -terms to LLCC (15) to flat-LLCC (14) to CSR (12). Therefore, the transformation preserves the semantics as composition of semantics preserving transformations.

In the case of a CHR encoding, the rule associated to each λ -abstraction can be denoted as a simpagation: $\langle p_i(K, \mathbf{v}), \text{value}(X) \setminus \text{apply}(K, X, V) \Leftrightarrow [e]^g(V). \rangle$.

Example 7. The λ -term $(\lambda_1 X. \lambda_2 Y. X) A B$ is transformed to the rules:

$$\begin{aligned}
& \text{start}(R, A, B) \Leftrightarrow \\
& \quad p1(F1), \text{apply}(F1, A0, F2), \text{apply}(F2, B0, R), \\
& \quad A=A0, \text{value}(A0), B=B0, \text{value}(B0). \\
& p1(F1), \text{value}(X) \mid \text{apply}(F1, X, F2) \Leftrightarrow p2(F2, X). \\
& p2(F2, X), \text{value}(Y) \mid \text{apply}(F2, Y, R) \Leftrightarrow X = R, \text{value}(R).
\end{aligned}$$

and the following goal, where the variable R codes the result:

$$\begin{aligned}
& / \text{ ?- start}(R, A, B). \\
& p1(_) \text{value}(_ X) \text{value}(_) \text{value}(_ X) p2(_, _ X) \\
& R = A
\end{aligned}$$

Related work

The adaptations of functional concepts in LLCC languages have been initiated with the embedding of closures and modules, leading to an encoding of λ -calculus in LLCC [41]. This chapter pursues the effort of transposing results in functional languages to concurrent constraint systems.

The translation from full LLCC to CHR relies on *ask-lifting*. This is a transformation comparable to the λ -lifting [47] for functional languages: the common idea is the materialization of the environment in data structures, *i.e.* values in functional languages or tokens in LLCC.

Flattening nested programming structures to CHR programs was suggested in [5] for connecting the Celf system [74] to CHR but, to our knowledge, no formal description of the transformation has been published.

Chapter 3

Reified Search: search strategies as constraint solving

In Section 3.1, we define constraint satisfaction problems as relational languages, that leads to very general definitions for constraints and propagators: in particular, “global constraints”, that is to say constraints that act on a unbounded number of variables, can here be presented as a single constraint, *i.e.* relational language, whereas usual presentations introduce them as families of constraints. Domain representations are introduced in the form of abstractions over the point-wise domain. In Section 3.3, we introduce the ClpZinc language and describe the transformation of search procedures from $\text{CLP}(\mathcal{X} + \mathcal{H})$ to and/or trees over \mathcal{X} with some tree traversal. In Section 3.3.2, we describe the transformation from those and/or trees to their internalization in $\text{CSP}(\mathcal{X})$, *i.e.*, into Zinc models with a back-end solver for \mathcal{X} . In the subsequent sections, we evaluate this approach on benchmarks of models with specific search strategies, namely: Korf’s Square Packing problem in Section 3.4, limited discrepancy search in Section 3.5.1 and symmetry breaking during search in Section 3.5.2. In Section 3.6, we show how it is possible to go beyond tree search procedures by using a simple mechanism of annotations for global store, and specify optimization procedures such as Branch-and-Bound.

3.1 Constraint Satisfaction

3.1.1 Constraints as Relational Languages

Let \mathcal{U} be a (possibly infinite) set, the *universe of values*. The set of finite sequences of values is denoted $\mathcal{U}^* = \bigcup_{n \in \mathbb{N}} \mathcal{U}^n$.

Definition 17. A *valuation* is an element of \mathcal{U}^* .

Definition 18. A *constraint* is a subset of \mathcal{U}^* .

A constraint is a set of admissible valuations. The set of constraints is therefore $\mathfrak{P}(\mathcal{U}^*)$. Given a constraint $C \subseteq \mathcal{U}^*$, the valuations $\nu = (\nu_1, \dots, \nu_n)$ that are in C are called the solutions of C . The proposition $(\nu_1, \dots, \nu_n) \in C$ is usually written $C(\nu_1, \dots, \nu_n)$. Constraints generalize n -ary relations over \mathcal{U} .

Definition 19. A constraint C is n -ary if for all $\nu \in C$, we have $\#\nu = n$ (i.e., $C \subseteq U^n$).

Example 8. The following constraints are present in most constraint solvers.

- Equality $=$ and difference \neq are binary constraints defined regardless of the universe \mathcal{U} . These constraints are usually written $\nu_1 = \nu_2$ and $\nu_1 \neq \nu_2$ for $=(\nu_1, \nu_2)$ and $\neq(\nu_1, \nu_2)$ respectively.
- Every n -ary function $f : \mathcal{U}^n \rightarrow \mathcal{U}$ is an $(n + 1)$ -ary functional relation and therefore induces an $(n + 1)$ -ary constraint: $f = \{(\nu_1, \dots, \nu_n, \nu_{n+1}) \in \mathcal{U}^{n+1} \mid f(\nu_1, \dots, \nu_n) = \nu_{n+1}\}$. The constraint is usually written $f(\nu_1, \dots, \nu_n) = \nu_{n+1}$ instead of $f(\nu_1, \dots, \nu_n, \nu_{n+1})$.
- Suppose that a subset $A \subseteq \mathcal{U}$ is ordered (typically, A is a set of numbers). The usual comparison operators $<, >, \leq, \geq$ over A are binary constraints.
- If C only contains a finite number m of solutions, these solutions can be explicitly given in a matrix $M \in \mathcal{U}^{m \times n}$. In this case, C is usually called the TABLE_M constraint or the RELATION_M constraint: $\text{TABLE}_M = \{(\nu_1, \dots, \nu_n) \in \mathcal{U}^n \mid \exists i, \forall j, M_{i,j} = \nu_j\}$.
- Subsets of U are 1-ary constraints, usually called domain constraints. The notation $\nu_1 \in C$ is often preferred over $C(\nu_1)$ in this case.

Suppose $\mathbf{Z} \subseteq \mathcal{U}$.

- The usual arithmetic operators $+, -, \times, \div$, as binary functions, are ternary constraints.
- The usual congruence is a ternary constraint: $\text{CONGRUENT} = \{(a, b, c) \in U^3 \mid a \equiv b \pmod{c}\}$.

A constraint satisfaction problem describe, for a set \mathcal{X} of variables, the constraints that the instantiations of these variables should satisfy to be solutions to the problem.

Definition 20. A constraint satisfaction problem is given by

- a finite set \mathcal{X} of variables;

- an initial family of *domains* $D : \mathcal{X} \rightarrow \mathfrak{P}(\mathcal{U})$ for all variables;
- a finite set $\mathcal{C} \subseteq \mathfrak{P}(\mathcal{U}^*) \times \mathcal{X}^*$ of (parameterized) constraints.

The elements (C, \mathbf{x}) of \mathcal{C} are written $C(\mathbf{x})$: $C(x_1, \dots, x_n)$ denotes the formal application of the constraint C to the variables (x_1, \dots, x_n) .

Definition 21. An *instantiation* is an element of $\mathcal{U}^{\mathcal{X}}$.

Let ν be an instantiation. ν is extended to formal applications: $\nu(C(x_1, \dots, x_n))$ denotes the proposition $C(\nu(x_1), \dots, \nu(x_n))$, *i.e.* $(\nu(x_1), \dots, \nu(x_n)) \in C$.

Definition 22. An instantiation ν is a *solution* when, for all $C(\mathbf{x}) \in \mathcal{C}$, we have $\nu(C(\mathbf{x}))$.

A constraint satisfaction problem realizes the conjunction of the constraints \mathcal{C} over \mathcal{X} and is usually denoted $C_1(\mathbf{x}_1) \wedge \dots \wedge C(\mathbf{x}_n)$, where $\mathcal{C} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. \mathcal{X} and D are often left implicit: \mathcal{X} is the set of variables that occur in \mathcal{C} and D can usually be inferred from \mathcal{C} through a simple propagation of domain constraints.

Example 9. Let $\mathcal{G} = (V, E)$ be a graph and \mathcal{U} a set of colors. The following constraint satisfaction problem describes the graph coloring of \mathcal{G} with the colors in \mathcal{U} .

$$\bigwedge_{(u,v) \in E} x_u \neq x_v$$

Example 10. Given two numbers $a, b \in \mathbf{Z}$, the following constraint satisfaction problem over $\mathcal{U} = \mathbf{Z}$ describes the numbers D that are common divisors of a and b . Supplementary variables A, B, U and V are introduced for the computation.

$$A \in \{a\} \wedge B \in \{b\} \wedge D \times U = A \wedge D \times V = B$$

The domain constraints $A \in \{a\}$ and $B \in \{b\}$ serve here to pass constants to the last parameters of the \times constraints.

Example 11. The classical n -queens problem can be described by the following constraint satisfaction problem over $\mathcal{U} = \{1, \dots, n\}$, where each variable x_i represents the column number of the queen at the i th row (or, dually, the row number of the queen at the i th column).

$$\begin{aligned} & \bigwedge_{1 \leq i < n} \bigwedge_{i < j \leq n} x_i \neq x_j \\ \wedge & \bigwedge_{1 \leq i < n} \bigwedge_{i < j \leq n} x_i + i \neq x_j + j \\ \wedge & \bigwedge_{1 \leq i < n} \bigwedge_{i < j \leq n} x_i - i \neq x_j - j \end{aligned}$$

where the constraints $x_i + i \neq x_j + j$ and $x_i - i \neq x_j - j$ are defined as expected.

- The negation of every constraint C can be expressed as a reified constraint.

$$y \in \{\mathbf{0}\} \wedge (y = \mathbf{1} \Leftrightarrow C(x_1, \dots, x_n))$$

- The disjunction between every couple of constraints (C, D) can be expressed as a reified constraint.

$$y \in \{\mathbf{0}, \mathbf{1}\} \wedge z \in \{\mathbf{0}, \mathbf{1}\} \wedge y \neq z \wedge (y = \mathbf{1} \Leftrightarrow C(x_1, \dots, x_n)) \wedge (z = \mathbf{1} \Leftrightarrow D(x'_1, \dots, x'_{n'}))$$

This construction is often referred as constructive disjunction, by opposition to disjunctions occurring in the search phase.

- The equivalence between every couple of constraints (C, D) can be expressed as a reified constraint.

$$y \in \{\mathbf{0}, \mathbf{1}\} \wedge (y = \mathbf{1} \Leftrightarrow C(x_1, \dots, x_n)) \wedge (y = \mathbf{1} \Leftrightarrow D(x'_1, \dots, x'_{n'}))$$

This construction is useful for expressing tunnelling between several representations of the same problem. For example, given two dual instances of the n -queens problem (example 11) x_1, \dots, x_n for the rows and x'_1, \dots, x'_n for the columns respectively, then the tunnelling between the two representations can be expressed as follows.

$$\bigwedge_{1 \leq i, j \leq n} (y_{i,j} \in \{\mathbf{0}, \mathbf{1}\} \wedge (y_{i,j} = \mathbf{1} \Leftrightarrow x_i = j) \wedge (y_{i,j} = \mathbf{1} \Leftrightarrow x'_j = i))$$

It is worth noticing that the variables $y_{i,j}$ induce themselves a third representation: the chess-board matrix where $y_{i,j} = \mathbf{1}$ if there is a queen at the i th row and j th column, otherwise $y_{i,j} = \mathbf{0}$.

3.1.2 Propagators and Consistency

We define in this section propagators as closure operators in the sense of the abstract interpretation theory [19]. Such a presentation is not new: it can be found for example in [14].

Definition 26. Given a set S , a relation $\sqsubseteq \subseteq S \times S$ is a (partial) *order* when \sqsubseteq is

- *reflexive*: for all $x \in S$, $x \sqsubseteq x$;
- *antisymmetric*: for all $x, y \in S$, if $x \sqsubseteq y$ and $y \sqsubseteq x$, then $x = y$;
- *transitive*: for all $x, y, z \in S$, if $x \sqsubseteq y$ and $y \sqsubseteq z$, then $x \sqsubseteq z$.

(S, \sqsubseteq) is called a (partially) ordered set. The *dual* order, \sqsupseteq , is the order such that, for all $x, y \in S$, $x \sqsupseteq y$ when $y \sqsubseteq x$.

Example 15. Let A be a set. The set $\mathfrak{P}(A)$ of all the subsets of A is ordered by inclusion \sqsubseteq .

The following definition relies on the usual vocabulary of the abstract interpretation theory [19].

Definition 27. Let (S, \sqsubseteq) be a (partially) ordered set and f an *operator* over S (i.e., $f : S \rightarrow S$).

- f is *extensive* (or *continuous*) when for all $x \in S$, $x \sqsubseteq f(x)$;
- f is *monotonic* (or *increasing*) when for all $x, y \in S$, if $x \sqsubseteq y$, then $f(x) \sqsubseteq f(y)$.
- f is *idempotent* when $f \circ f = f$;
- f is a *closure operator* when f is extensive, monotonic and idempotent.

Definition 28. Given a closure operator f , an element $x \in S$ is *closed* by f if x is a fix-point of f (i.e., $f(x) = x$).

Since f is idempotent, x is closed by f if and only if $x \in \mathfrak{S}(f)$.

Definition 29. Let (S, \sqsubseteq) be a (partially) ordered set and A be another set. The set of all maps $f : A \rightarrow S$ form a partially ordered set (S^A, \sqsubseteq^*) , the *point-wise order*, such that, for all $f, g : A \rightarrow S$, $f \sqsubseteq^* g$ when $f(a) \sqsubseteq g(a)$ for all $a \in A$.

Example 16. For all $n \in \mathbf{N}$, since S^n is isomorphic to S^A with $A = \{1, \dots, n\}$, the point-wise order induces the partially ordered set (S^n, \sqsubseteq^*) . Considering the disjoint union $S^* = \bigcup_{n \in \mathbf{N}} S^n$, the point-wise order also induces the partially ordered set (S^*, \sqsubseteq^*) such that $(\mu_1, \dots, \mu_m) \sqsubseteq^* (\nu_1, \dots, \nu_n)$ when $m = n$ and $\mu_i \sqsubseteq \nu_i$ for all $1 \leq i \leq n$.

Definition 30. A *sequence of domains* is an element of $\mathfrak{P}(\mathcal{U})^*$.

Definition 31. A *propagator* is an extensive operator over sequences of domains ordered by \sqsupseteq^* .

It is worth noticing that the extensivity ensures that for all $D \in \mathfrak{P}(\mathcal{U})^*$, $\#\pi(D) = \#D$.

Definition 32. Given a sequence of domains $D = (D_1, \dots, D_n)$, the *flattening* of D is the set of valuations $\downarrow D = \{(\nu_1, \dots, \nu_n) \in \mathcal{U}^* \mid \forall i, \nu_i \in D_i\}$.

Definition 33. Given a constraint C and a sequence of domains D , the *set of solutions of C in D* is the set $\text{sol}_C(D) = C \cap \downarrow D$.

sol_C is monotonic: for all sequences of domains D and D' , if $D \subseteq^* D'$, then $\text{sol}_C(D) \subseteq \text{sol}_C(D')$.

Definition 34. A propagator π filters a constraint C when it preserves the set of solutions of C (i.e., for all sequence of domains D , $\text{sol}_C(\pi(D)) = \text{sol}_C(D)$).

Definition 35. Given a constraint C , the *canonical propagator for C* is

$$\begin{aligned} \pi_C : \mathfrak{P}(\mathcal{U})^* &\rightarrow \mathfrak{P}(\mathcal{U})^* \\ D = (D_1, \dots, D_n) &\mapsto D' = (D'_1, \dots, D'_n) \end{aligned}$$

where $D'_i = \{\mu_i \in D_i \mid \exists \nu \in \text{sol}_C(D), \mu_i = \nu_i\}$.

π_C is extensive by definition: $D'_i \subseteq D_i$.

3.1.3 Abstractions

Definition 36. An *abstraction* is a closure operator over sequences of domains ordered by \subseteq^* .

Definition 37. Given an abstraction α , a sequence of domains $D \in \mathfrak{P}(\mathcal{U})^*$ is α -consistent with respect to a constraint C when it is the smallest sequence of domains with respect to \subseteq^* among the sequences of domains D' such that

- D' is closed by α ,
- $\text{sol}_C(D) \subseteq \text{sol}_C(D')$.

Property 1. A sequence of domains $D \in \mathfrak{P}(\mathcal{U})^*$ is α -consistent with respect to a constraint C if and only if for all propagators π that filter C , D is a fix-point for $\alpha \circ \pi$.

Proof. Suppose that D is α -consistent with respect to C . Then, for all propagators π that filter C , let $D' = \alpha \circ \pi(D)$. We will show that $D = D'$.

- $D \subseteq D'$: We show that D' satisfies the hypotheses of Definition 37 so that the minimality of D concludes. Indeed, D' is closed by α since α is idempotent. Moreover, $\text{sol}_C(D) \subseteq \text{sol}_C(D')$ since π preserves the set of solutions and α is extensive with respect to \subseteq^* .
- $D' \subseteq D$: Indeed, $D \supseteq \pi(D)$ since π is extensive, therefore $\alpha \circ \pi(D) \subseteq \alpha(D)$ since α is monotonic, and $\alpha(D) = D$ since D is closed.

Conversely, suppose that, for all propagators π that filter C , $D = \alpha \circ \pi(D)$.

□

3.2 The Language ClpZinc

We propose to use Constraint Logic Programming (CLP) clauses to specify search strategies in Zinc. More precisely, given a constraint system \mathcal{X} (e.g. finite domain constraints) and a CSP model with constraints in \mathcal{X} , we consider search procedures that are expressible as the traversal of an and/or tree with constraints over \mathcal{X} , *i.e.* an and/or tree where every leaf is either labeled by a constraint in \mathcal{X} or, for dynamic search procedures, labeled by a query to indexicals. In addition, we consider the Prolog primitive constraint system, \mathcal{H} , *i.e.* Herbrand terms with unification. The choice of Herbrand terms for representing Zinc data structures makes the language look familiar to Prolog users. Similarly, we fix the CLP strategy as depth-first and left-to-right.

The modeling language Zinc, and its implementation MiniZinc¹, succeeded in becoming a common input format across many solvers in the Constraint Programming community. In Zinc, the search procedure is specified through special annotations that are dedicated to the constraint solver [69] and ignored by the other solvers.

The language ClpZinc is an extension of Zinc where the item `solve satisfy`; in models is replaced by a CLP goal of the form “`:- goal.`”, and where user-defined predicates are defined by CLP clauses of the form “`p(t1, . . . , tn) :- goal.`”.

Example 17 (Labeling). *The following ClpZinc model implements the search strategy that enumerates all possible values for a given variable in ascending order.*

```
var 0..5: x;
constraint x * x = x + x;

labeling (X, Min, Max) :-
    Min <= Max, (X = Min ; labeling(X, Min + 1, Max)).

:- labeling(x, 0, 5).
output [show(x)];
```

As shown in the following section, this ClpZinc model for the given goal of labeling x between 0 and 5, can be expanded to the following MiniZinc model:

```
var 0..5: x;
constraint x * x = x + x;
var 0..5: X1;
constraint X1 = 0 -> x = 0;
constraint X1 = 1 -> x = 1;
constraint X1 = 2 -> x = 2;
constraint X1 = 3 -> x = 3;
constraint X1 = 4 -> x = 4;
constraint X1 = 5 -> x = 5;
solve :: seq_search([
    int_search([X1], input_order, indomain_min, complete)
]) satisfy;
output [show(x)];
```

¹<http://www.minizinc.org/>

Definition 38. A *ClpZinc goal* is either

- a constraint,
- a MiniZinc search annotation,
- a call to a user-defined predicate,
- the conjunction (A;B) or the disjunction (A;|B) of two goals.

A *ClpZinc clause* is an item of the form $p(t_1, \dots, t_n) :- goal$. where t_1 and t_n are terms and *goal* is a ClpZinc goal. The goal part can be omitted: “ $p(t_1, \dots, t_n)$.” is a shorthand for “ $p(t_1, \dots, t_n) :- true$.”.

The search annotations of MiniZinc are accessible in goals in order to allow the composition of user-defined strategies with built-in ones. Terms are either logical variables (X, Y, Max, \dots), numbers, or compound terms of the form $p(t_1, \dots, t_n)$ where t_1, \dots , and t_n are terms. Model variables are a special case of compound terms, either atomic (a, b, \dots) or array accessors ($x[I, J]$). Zinc arrays have been unified with Prolog-like lists to ease their enumeration in search strategies.

In $CLP(\mathcal{X} + \mathcal{H})$, arithmetic differs from Prolog. Indeed, in accordance with the theory of CLP and unlike most Prolog systems, arithmetic is supposed to be contained in \mathcal{X} and is distinguished from \mathcal{H} terms, *e.g.*, “ $1 + 1$ ” is undistinguishable from “ 2 ” and is not a \mathcal{H} term. In ClpZinc, the different forms of unification, equality, and evaluation predicates that are encountered in Prolog systems ($=, \neq, is, \dots$) are thus all unified in a unique notion of equality, which is accessible either explicitly with the predicate $=$, or implicitly when predicate arguments in either \mathcal{X} or \mathcal{H} are unified.

Arithmetic expressions are also extended for accessing the indexicals of the model variables. For instance, the goal $M = \min(X)$ assumes that X is a model variable and unifies M with the currently known lower-bound of X . We consider the indexicals \min , \max , card and dom_nth (for retrieving the n th value in a variable domain). Concretely, an intermediary variable is introduced to receive the value of the indexical and search annotations are emitted for getting them with:

```

annotation indexical_min(var int: target, var int: x);
annotation indexical_max(var int: target, var int: x);
annotation indexical_card(var int: target, var int: x);
annotation indexical_dom_nth(var int: target, var int: x, var int: n);

```

These annotations require to extend the solvers to communicate the indexicals. That is the only change made to the interface of the solvers.

Example 18 (Dichotomic search). *The Zinc `indomain_split` value selection strategy can be implemented in ClpZinc using indexicals. The predicate `dichotomy/3` below expresses the bisection of a variable X that has the initial domain $Min \dots Max$. The bisection defined in the auxiliary predicate `dichotomy/2` is iterated $Depth = \lceil \log_2 \#X \rceil$ times to ensure that the domain is reduced to a value on every leaf.*

```
dichotomy(X, Min, Max) :-
    dichotomy(X, ceil(log(2, Max - Min + 1))).
```

```
dichotomy(X, Depth) :-
    Depth > 0,
    Middle = (min(X) + max(X)) div 2,
    (X <= Middle ; X > Middle),
    dichotomy(X, Depth - 1).
dichotomy(X, 0).
```

```
var 0..5: x;
:- dichotomy(x, 0, 5).
```

The MiniZinc model generated for the given goal is

```
var 0..5: x;
var 0..5: X3; var 0..5: X5; var 0..1: X7;
var 0..5: X4; var 0..5: X6; var 0..5: X2;
var 0..1: X8; var 0..5: X1; var 0..1: X9;
constraint X7 = 0 <-> x <= (X1 + X2) div 2;
constraint X8 = 0 <-> x <= (X3 + X4) div 2;
constraint X9 = 0 <-> x <= (X5 + X6) div 2;
solve :: seq_search([
    indexical_min(X1, x),
    indexical_max(X2, x),
    int_search([X7], input_order, indomain_min, complete),
    indexical_min(X3, x),
    indexical_max(X4, x),
    int_search([X8], input_order, indomain_min, complete),
    indexical_min(X5, x),
    indexical_max(X6, x),
    int_search([X9], input_order, indomain_min, complete)
]) satisfy;
```

The next example shows a partial search strategy that is not available using the usual MiniZinc search annotations. This is the interval slitting strategy introduced in [80] for solving Korf's packing problem [48] (explained in Section 3.4) by making a preliminary coarse-grained filtering of the variable domains.

Example 19 (Interval splitting). *The predicate `interval_splitting/4` defined below expresses the splitting of the domain of X into intervals of width $Step$. X is supposed to have the initial domain $Min \dots Max$.*

```
interval_splitting(X, Step, Min, Max) :-
    Min + Step <= Max, NextX = min(X) + Step,
    (
        X < NextX
```



```

;
  X >= NextX,
  interval_splitting (X, Step, Min + Step, Max)
).
interval_splitting (X, Step, Min, Max) :-
  Min + Step > Max.
var 0..5: x;
:- interval_splitting (x, 2, 0, 5).

```

The corresponding MiniZinc model for the given goal is

```

var 0..5: x;
var 0..1: I3; var 0..1: I4; var 0..5: I2;
var 0..5: I1;
constraint I3 = 0 <-> x < I1 + 2;
constraint I3 = 1 -> (I4 = 0 <-> x < I2 + 2);
constraint I3 = 0 -> I4 = 0;
solve :: seq_search([
  indexical_min(I1, x),
  int_search([I3], input_order, indomain_min, complete),
  indexical_min(I2, x),
  int_search([I4], input_order, indomain_min, complete)
]) satisfy;

```

3.3 Extending Zinc with CLP Clauses

3.3.1 Partial Evaluation of ClpZinc into And/Or Trees

From now on, let us assume that the initial ClpZinc goals provided in the items “:- goal.” of the ClpZinc models that we consider, always terminate. That hypothesis should hold even if \mathcal{X} only resolves fully instantiated constraints, as is the case of the static partial evaluator. Verifying termination of logic programs is a classical topic for which many results have been obtained using type systems or abstract interpretation techniques [13]. The description of these techniques is however beyond the scope of this thesis.

Given a constraint system \mathcal{X} , the partial evaluation of a $\text{CLP}(\mathcal{X} + \mathcal{H})$ goal will lead to an and/or tree with constraints over \mathcal{X} . The partial evaluator resolves predicate calls, Herbrand constraints and fully instantiated arithmetic constraints, *i.e.*, arithmetic tests. Since, without loss of generality, we settled for a DFS left-to-right CLP evaluation, the and/or trees will be traversed in a similar DFS left-to-right fashion in our examples, but any other traversal order can be treated similarly.

As shown in Figure 3.1 for Example 17, or-nodes are flattened so that nested choices become a single large disjunction. And-nodes are similarly flattened into conjunctions. In the general case, the partial evaluation of the continuation may duplicate constraints with different partial instantiations. For instance, Figure 3.2 shows a simple example of duplication with partial instantiation of the bounding constraint $\text{Min} \leq x, x \leq \text{Max}$.




Figure 3.1: And/Or tree derived from Example 17 (Labeling)




Figure 3.2: And/Or tree for the ClpZinc goal

```
var 0..5: x;
:- (Min=3, Max=5; Min=0, Max=2), Min <= x, x <= Max.
```

However, when the partial evaluation store is left unchanged by a choice (typically, when only constraints in \mathcal{X} are involved), the continuation will remain undeveloped, as shown in Figure 3.3 for Example 18. The and/or tree is in logarithmic size with respect to the size of the domain whereas the fully expanded search tree would be in linear size.

We assume that the partial evaluation terminates and that the resulting and/or tree meets the following conditions:

1. all the variables that appear in constraints are finite-domain variables;
2. all lists are well-formed, in particular the tail of every non-empty list is a list and cannot be a variable since such a variable would be a finite-domain variable according to the previous condition (this ensures that lists can be expanded in Zinc array literals);
3. all annotations except indexicals do not appear below a choice point (*i.e.* their execution is unconditional).

3.3.2 Compiling And/Or Trees into Zinc Reified Constraints

Given a CSP(\mathcal{X}) model \mathcal{M} and a tree search strategy represented by the traversal of an and/or tree t , the generation of Zinc code proceeds by assigning an additional model variable to every or-node in the tree t , and by emitting search annotations

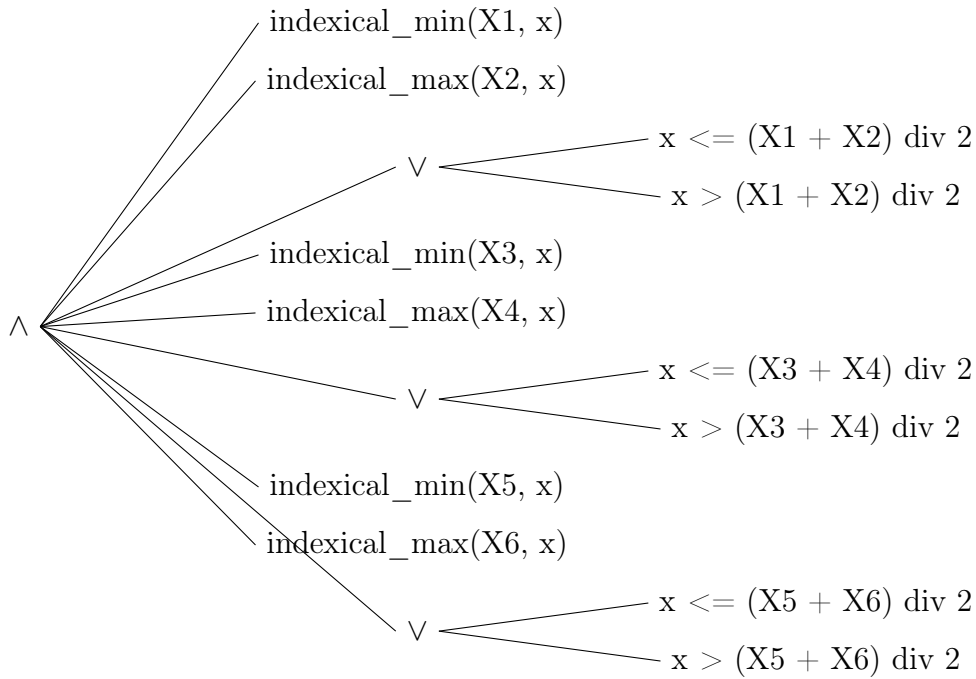


Figure 3.3: And/Or tree for Example 18 (Dichotomic search)




Figure 3.4: And/Or tree for Example 19 (Interval splitting)

that fix the enumeration strategy for these additional variables in a way compatible with the traversal ordering.

In Figure 3.1, the variable $X1$ is assigned to the root node, with the domain $0..5$ corresponding to the arity of the node. As shown in the Zinc model generated for Example 17, each constraint labeling the leaves under this or-node appears in the

model guarded by an implication checking for a particular value of X1. Therefore, when the search annotation `int_search` enumerates the possible values of X1, these guarded constraints are successively enabled for exploring the different branches of the tree.

More generally, the transformation presented in this chapter can be seen as a constructive proof for the following theorem. We call *fixed enumeration strategies* the search strategies that are reduced to a sequence of variables selected in a fixed order and enumerated with the increasing value selection (`indomain_min`). For dynamic search strategies, this sequence is possibly interleaved with accesses to indexicals.

Theorem 5. *For every pair (\mathcal{M}, t) where \mathcal{M} is a CSP model and t a tree search strategy described as the traversal of an and/or tree, there exists a model \mathcal{M}_t and a fixed enumeration strategy t' such that the resolution of $(\mathcal{M} + \mathcal{M}_t, t')$ explores the same search tree as (\mathcal{M}, t) .*

Proof. First, let us remark that in $\mathcal{M} + \mathcal{M}_t$, the variables and the constraints of \mathcal{M} are left unchanged; only additional model variables accompanied with additional constraints are introduced in \mathcal{M}_t .

Let us assume that we have a function ℓ that maps each or-node n of t to a model variable $\ell(n) \in V(\mathcal{M}_t)$, such that for every pair n_1, n_2 of nodes of t , if $\ell(n_1) = \ell(n_2)$, either $n_1 = n_2$ or the lowest common ancestor of n_1 and n_2 is an or-node.

Each constraint c that appears as a leaf of t is translated as a constraint in \mathcal{M}_t . Let n_1, \dots, n_k denote the or-nodes that are traversed by the path π from the root of t to the leaf c and, for every $1 \leq i \leq k$, let p_i be the rank of the branch taken by π at node n_i . We adopt the convention that branches are numbered from left to right and that the left-most branch has rank 0. Then the following constraint is posted in the MiniZinc model, for translating the leaf c :

$$\text{constraint } \ell(n_1)=p_1 \wedge \dots \wedge \ell(n_k)=p_k \rightarrow c;$$

Let $X \in V(\mathcal{M}_t)$ be one of the variables that label or-nodes. The domain of X will be $0.. \max\{w(n) - 1 \mid \ell(n) = X\}$ where $w(n)$ denotes the width of the or-node n (i.e., the number of branches issued from n). For every or-node n_k such that $\ell(n_k) = X$ that does not reach this maximum, the following additional constraint is posted, where $(n_i, p_i)_i$ denotes the or-path to n_k as above:

$$\text{constraint } \ell(n_1)=p_1 \wedge \dots \wedge \ell(n_{k-1})=p_{k-1} \rightarrow \ell(n_k) < w(k);$$

We should now establish the connection between the enumeration of the variables that label the or-nodes and the exploration of the and/or tree.

Search annotations have to be emitted to fix the enumeration strategy for the variables $V(\mathcal{M}_t)$. MiniZinc search annotations have a depth-first semantics. To reproduce the semantics of (\mathcal{M}, t) , it is thus sufficient in t' to emit the annotations that select the variables in the order where their corresponding nodes are encountered following the traversal of t . The value selection strategy fixes the order in which the sub-branches are explored. For t' , it can thus be reduced to the value selection `indomain_min` that selects the left-most branch, by switching the values if necessary. We thus have that, by construction, $(\mathcal{M} + \mathcal{M}_t, t')$ explores the same search tree as (\mathcal{M}, t) . \square \square

The two following optimizations are not mandatory but have been measured to give significant performance improvements:

1. to prevent enumerating on X in branches where X does not occur, the following constraint imposes a fixed value to X on these branches.

$$\text{constraint } \left(\bigwedge_{\substack{(n_i, p_i)_i \\ \ell(n_k)=X}} \ell(n_i) \neq p_i \vee \dots \vee \ell(n_k) \neq p_k \right) \rightarrow X=0;$$

2. in the particular case where a constraint c occurs under an or-node n_k (possibly separated with some and-nodes) and when $\neg c$ occurs in every other branches of n_k , then the following constraint is posted instead (and the constraints corresponding to the leaves $\neg c$ are not posted).

$$\text{constraint } \ell(n_1)=p_1 \wedge \dots \wedge \ell(n_{k-1})=p_{k-1} \rightarrow (\ell(n_k)=p_k \leftrightarrow c);$$

These simplifications can be seen in the Zinc code generated for Examples 18 and 19.

3.4 Computation Results on Korf's Square Packing Benchmark

In this section, we consider Korf's Optimal Rectangle Square Packing problem [48], *i.e.*, given an integer $n \geq 1$, find an enclosing rectangle of smallest area containing n squares from sizes 1×1 , 2×2 , up to $n \times n$, without overlap. Helmut Simonis and Barry O'Sullivan proposed a complex dynamic search strategy for that problem in [80], which is interesting to specify and evaluate in ClpZinc. In Fig. 3.5, an optimal solution is drawn for $n = 24$ squares.

First, the model they consider for packing the n consecutive squares in a rectangle of size $w \times h$ can be written in MiniZinc as follows. Since the 1×1 square




Figure 3.5: Optimal 88×56 solution for Korf's problem with $n = 24$

can always be placed afterward if the area $w \times h$ is big enough, the model only considers the remaining $n - 1$ other squares. Two redundant cumulative constraints are introduced. The two last constraints break some symmetries by forcing the largest square to be in the bottom-left quadrant.

```

int : n;

constraint diffn (
  x,y,[i+1|i in 1..n-1],[i+1|i in 1..n-1]
);
constraint cumulative(
  x,[i+1|i in 1..n-1],[i+1|i in 1..n-1],h
);
constraint cumulative(
  y,[i+1|i in 1..n-1],[i+1|i in 1..n-1],w
);

constraint forall (i in 1..n-1)
  (x[i] <= w - i & y[i] <= h - i);
constraint x[n-1] <= (w - n + 2) div 2;
constraint y[n-1] <= (h + 1) div 2;

```

Second, the optimization procedure used in [80] enumerates all the possible

sizes $w \times h$ for the enclosing rectangle by increasing area. This strategy can be internalized in the model by successively considering all the rectangles up to $\text{max_size} \times \text{max_size}$, from the minimal area covered by the squares themselves ($\sum_{1 \leq i \leq n} i^2$) and with bounds on w and h that are described in [80]:

```

int : max_size;
array [1..n-1] of var 1..max_size: x;
array [1..n-1] of var 1..max_size: y;
var 0..max_size: w; var 0..max_size: h;
var 0..max_size * max_size: area;

constraint w * h = area /\ w <= h;
constraint sum([i*i | i in 1..n]) <= area;
constraint w >= 2 * n - 1
  \/ h >= (n * n + n -
    ((w + 1) div 2 - 1) * ((w + 1) div 2 - 1)
    - ((w + 1) div 2 - 1)) div 2;

```

Now, the search strategy of [80] firsts enumerates in x and then in y , considering in each dimension a preliminary interval splitting on the origins of the squares from sizes $n \times n$ to 7×7 , and then a dichotomic search on the origins, still by considering the biggest square first. This search strategy is implemented in ClpZinc by enumerating first on area and w to find the rectangle of smallest area first. It is worth noticing that we can combine the user-defined interval splitting strategy defined in Example 19 with the built-in dichotomic search (`indomain_split`).

```

interval_splitting_list (L, S, Stop) :-
  (S <= Stop ; S > Stop, L = []).
interval_splitting_list ([H | T], S, Stop) :-
  S > Stop,
  interval_splitting (
    H, max(1, (S * 3) div 10) + 1, 0, max_size
  ),
  interval_splitting_list (T, S - 1, Stop).

:- int_search(
  [area, w], input_order, indomain_min, complete
),
reverse(x, RXs), interval_splitting_list (RXs, n, 6),
int_search(
  RXs, input_order, indomain_split, complete
),
reverse(y, RYs), interval_splitting_list (RYs, n, 0),

```

```

int_search(
  RYs, input_order, indomain_split, complete
).

```

This strategy can be compared to the use of the dichotomic search only, on each dimension, from the biggest to the smallest square, relying on the native `indomain_split` of MiniZinc. This is indeed a good candidate for the best strategy that can be easily written in MiniZinc without the help of ClpZinc.

```

solve :: seq_search([
  int_search(
    [area, w], input_order, indomain_min, complete
  ),
  int_search(
    [x[n-i] | i in 1..n-1] ++ [y[n-i] | i in 1..n-1],
    input_order, indomain_split, complete
  )
]) satisfy ;

```

To measure the overhead of ClpZinc, we also include the version of dichotomic search relying on the user-defined predicate of Example 18.

```

:- int_search(
  [area, w], input_order, indomain_min, complete
),
reverse(x, Rx), dichotomy_list(Rx, 0, max_size),
reverse(y, Ry), dichotomy_list(Ry, 0, max_size).

```

Table 3.1 shows the results of the native dichotomic search procedure in MiniZinc, of the user-defined dichotomic and interval-splitting search procedure in ClpZinc, solved either Choco or SICStus solvers, and of the original SICStus-Prolog program of [80], all of them running on Intel®Xeon®CPU E5-1620 0 @ 3.60GHz machines. As shown in Table 3.1, the overhead introduced by the reification of the search procedure is quite reasonable, averaging a two-fold slowdown of the program. On the other hand, the reified search enables the encoding of the interval splitting strategy that induces a crucial increase in performance comparable to the results obtained in [80].

That table also shows that the specification in ClpZinc of the dichotomic and interval-splitting search strategy makes it readily available in a variety of solvers for which its implementation was not trivial. The implementation in Choco is the most efficient, followed by SICStus-Prolog, probably due to differences in the implementation of reified constraints.

n	Choco 3			SICStus	
	dichotomic indomain_split	dichotomic ClpZinc	interval split then dichotomic ClpZinc	interval split then dichotomic ClpZinc	interval split then dichotomic Original
16	9232	14402	853	710	340
17	16321	21643	982	450	250
18	422116	570407	7978	9400	4850
19	785080	1051418	6984	11710	4310
20			12572	17330	8970
21			42892	88310	32370
22			208632	303810	153860
23			1340816	2104020	999020
24			2312933	3433410	1481910
25			29201522		10662860
26			142702128		62179600

Table 3.1: Solving times in ms for Korf’s problem for strategies implemented in ClpZinc with Choco 3 and SICStus as solvers, compared to the original SICStus program.

3.5 LDS and SBDS as Strategy Transformers in ClpZinc

Since and/or trees are first-class terms in ClpZinc, they can be arguments of ClpZinc predicates to define search strategy transformers. In this section, we illustrate this possibility with the modeling of Limited Discrepancy Search (LDS) [42] and Symmetry Breaking During Search (SBDS) [34] as strategy transformers for labeling or dichotomic search for instance. This technique is closely related to the monadic approach of strategy transformers presented in [77]. The main difference, outside of purely syntactic choices, is that the monadic transformations described in [77] heavily rely on laziness to not expand the trees, whereas in ClpZinc, in order to finally compile towards a CSP, we fully meta-interpret, and therefore expand, the search trees, with some possible benefits thanks to the propagation of search constraints.

3.5.1 Limited Discrepancy Search

LDS can be modeled very simply in ClpZinc using meta-interpretation. Basically the and/or tree is developed but the right turns are counted at the same time, by increment when going in the right branch of an *or* and by addition of the two

branches when going through an *and*:

```

lds(true, L).
lds((A ; B), L) :-
    domain(L0, 0, 1024), domain(D, 0, 1),
    ( D = 0, lds(A, L0)
    ; D = 1, lds(B, L0)),
    L = D + L0.
lds((A, B), L) :-
    domain(L0, 0, 1024), domain(L1, 0, 1024),
    lds(A, L0), lds(B, L1),
    L = L0 + L1.
lds(B, L) :- builtin(B), B, L = 0.
lds(H, L) :- clause(H, B), lds(B, L).

```

Interestingly, since right turns are counted at the constraint level, the propagation of search constraints may actively reduce the search space, whereas a classical procedural implementation of LDS limits the number of right turns by generate-and-test. The following example demonstrates an *exponential speed-up* thanks to this propagation with respect to a procedural implementation of LDS.

```

var 0..1: x;
var 0..1: y;
array[0..n] of var 0..1: a;

:- int_search(a, input_order, indomain_min, complete),
   lds(((x = 0; x = 1), (y = 0; y = 1)), 0), x != y.

```

Whereas a procedural implementation would explore the 2^n possible assignments for *a* before detecting that the model is unsatisfiable within the reduced search space, the inconsistency is immediately detected in the MiniZinc model generated by ClpZinc.

```

var 0..1: x;
var 0..1: y;
array[0..n] of var 0..1: a;
constraint x = 0;
constraint y = 0;
constraint x != y;
solve :: seq_search([
    int_search(a, input_order, indomain_min, complete)
]) satisfy;
n = 1000;

```

3.5.2 Symmetry Breaking During Search

Symmetry Breaking During Search [7, 34] is a general method that transforms a search tree so as to remove symmetric branches from enumeration. Each time the search backtracks from enumerating solutions with given a search constraint *c*, the

other search branch considers $\neg c$ and also all the symmetric constraints $\sigma(\neg c)$ for symmetries σ compatible with search constraints already posted. This schema is implemented in the predicate below, supposing a predicate `cut_symmetry` that adds the symmetric negations for a given constraint.

```

sbds(top, _).
sbds(or(A, B), Path) :-
  ( A = constraint(C, A0),
    ( C, sbds(A, [C | Path])
      ; cut_symmetry(C, Path), sbds(B, Path))
    ; A \= constraint(_, _),
      (sbds(A, Path) ; sbds(B, Path))).
sbds(constraint(C, T), Path) :- C, sbds(T, [C | Path]).
:- search_tree(labeling_list(queens, 1, n), T),
   sbds(T, []).

```

The predicate `search_tree` constructs the search tree associated with the and/or tree of a CLP goal by meta-interpretation (full code in Appendix 3.8).

3.6 Beyond Tree Search Strategies

Some search strategies require to iterate a search tree several times with a memory passed from one branch to another. That is typically the case for optimization methods like branch-and-bound where the best score reached up to now is remembered from one iteration to another of the underlying search strategy, or for shaving, where one step of propagation is performed and undone in order to select the best one. In languages like Prolog, such methods are implemented with the help of a global state, most commonly stored within the fact database (with `assert` and `retract`). We propose two additional annotations for search in MiniZinc to handle global state.

```

annotation store(var bool: c, string: id, array[int] of var int: src);
annotation retrieve(string: id, array[int] of var int: target);

```

The semantics of `store(cond, id, source)` is to remember, if `cond` is true, the current values of the sequence of variables `source` into the global state identified as `id`. The store annotation does nothing if `cond` is false, such that the assignation to `id` is skipped outside the computation branch that involves this assignation. The parameter `cond` does not appear in ClpZinc: it is implicitly fixed to the guard associated to the path leading to the node where the annotation appears in the and/or tree. The semantics of `retrieve(id, target)` is to assign the values previously remembered into the global state identified as `id` into the sequence of variables `target`.

As shown below, these two simple annotations allow the specification of branch-and-bound optimization in ClpZinc. Once again, for such strategy one might also use the native maximize annotation of MiniZinc, but as far as we know, more complex iterative procedures like shaving or enumerating solutions, using previously found ones in the search (whether to guide it or to limit it), cannot be natively written in MiniZinc.

```

maximize(G, S, Min, Max) :-
  domain(I, Min, Max + 1), domain(Best, Min, Max),
  domain(Fail, 0, 1),
  domain(A, 0, 1), domain(B, 0, 1), domain(C, 0, 1),
  (Fail = 0 -> A != B /\ B != C /\ A != C),
  store("bb_best", [Min, 0]),
  labeling(I, Min, Max + 1),
  retrieve("bb_best", [Best, Fail]),
  ( Fail = 0, store("bb_best", [Best, 1]),
    S > Best, G, store("bb_best", [S, 0]),
    labeling(A, 0, 1), labeling(B, 0, 1)
  ; Fail = 1, I = Max + 1, S = Best, G).

```

```

minimize(G, S, Min, Max) :-
  domain(Dual, Min, Max), Dual = Max - S + Min,
  maximize(G, Dual, Min, Max).

```

Note that in order to make this branch-and-bound procedure possible, the gap between failures at the search and at the constraint level has to be bridged. Using the incompleteness of arc-consistency, the reified constraint imposing that A, B and C are all different allows us to fail at will in the success branches (Fail = 0) by labelling A and B. There is also an optimization in the above code where the upper bound on the score is used in the Fail = 1 branch as some kind of *cut*: all attempts after the first failure will be immediately discarded, except the last one where appropriate values for variables will be rebuilt by running the goal G again.

3.7 ClpZinc Models of First-Fail and Middle-out Strategies

The full code for the first-fail variable selection strategy is as follows. Basically, the `first_fail` predicate selects the variable with smallest domain of the list of variables given as first argument.

```

first_fail ([H | T], X) :-

```

```

(
  check_card_greater_than(T, card(H)),
  X = H
;
  first_fail_tail (T, X)
).

first_fail_tail (Vars, X) :-
  select (X, Vars, Other_vars),
  CardX = card(X),
  CardX > 1,
  check_card_greater_than(Other_vars, CardX).

```

```
check_card_greater_than([], _).
```

```
check_card_greater_than([H | T], Min) :-
  CardH = card(H),
  (CardH < Min -> CardH = 1),
  check_card_greater_than(T, Min).

```

The full code for the middle-out value selection strategy is as follows: The `middle_out` predicate selects the value in the middle of the `X` variable's domain, removes it from its domain and iterates upper bound times over it.

```

var 0..10 : a;
var 0..10 : b;
var 0..10 : c;

constraint a != b;

middle_out(X, N) :-
  N > 0,
  Middle = dom_nth(X, card(X) div 2),
  (
    X = Middle
  ;
    X != Middle,
    middle_out(X, N - 1)
  ).

middle_out(X, 0).

```

```

:- middle_out(a, 10).

output [ show(a) ];

```

3.8 ClpZinc Model of Symmetry Breaking During Search

The full code for SBDS on the N-queens example, including the transformation of the and/or tree to a search tree during meta-interpretation.

```

int : n;
array [1..n] of var 1..n: queens;

include "globals.mzn";
include "prolog.plz";

constraint all_different (queens);
constraint all_different ([queens[i] + i | i in 1..n]);
constraint all_different ([queens[i] - i | i in 1..n]);

labeling(X, Min, Max) :-
    Min <= Max,
    (
        X = Min
    ;
        labeling(X, Min + 1, Max)
    ).

labeling_list ([], Min, Max).
labeling_list ([H | T], Min, Max) :-
    labeling(H, Min, Max),
    labeling_list (T, Min, Max).

builtin (true).
builtin (false).
builtin (_ = _).
builtin (_ <= _).
builtin (_ > _).

sequence_tree(top, B, T) :-

```

```

search_tree(B, T).

sequence_tree(bot, _T, bot).

sequence_tree(or(A0, B0), T, or(A1, B1)) :-
    sequence_tree(A0, T, A1),
    sequence_tree(B0, T, B1).

sequence_tree(constraint(C, T0), T, constraint(C, T1)) :-
    sequence_tree(T0, T, T1).

search_tree((A, B), T) :-
    search_tree(A, TA),
    sequence_tree(TA, B, T).

search_tree((A ; B), T) :-
    search_tree(A, TA),
    search_tree(B, TB),
    (
        TA == false,
        TB == false,
        T = bot
    ;
        TA == false,
        TB \= false,
        T = TB
    ;
        TA \= false,
        TB == false,
        T = TA
    ;
        TA \= false,
        TB \= false,
        T = or(TA, TB)
    ).

search_tree(B, T) :-
    builtin(B),

```

```

(
  B == true,
  T = top
;
  B == false,
  T = bot
;
  B \= true,
  B \= false,
  T = constraint(B, true)
).

search_tree(H, T) :-
  findall (B, clause(H, B), L),
  L \= [],
  disjunction(L, D),
  search_tree(D, T).

disjunction ([], false).
disjunction ([X], X).
disjunction ([A, B | T], (A; C)) :-
  disjunction ([B | T], C).

sbds(top, _).

sbds(or(A, B), Path) :-
  (
    A = constraint(C, A0),
    (
      C,
      sbds(A, [C | Path])
    ;
      cut_symmetry(C, Path),
      sbds(B, Path)
    )
  ;
  A \= constraint(_, _),
  (
    sbds(A, Path)
  )

```



```

        ;
        sbds(B, Path)
    )
).

sbds(constraint(C, T), Path) :-
    C,
    sbds(T, [C | Path]).

:- search_tree(labeling_list(queens, 1, n), T), sbds(T, []).

cut_symmetry(C, Path) :-
    (
        C = (queens[I] = J),
        (
            Path = [],
            queens[I] != n - J + 1
        ;
            Path \= []
        )
    ;
        C \= (queens[I] = J),
        true
    ).

output [
    show(queens) ++ "\n"
] ++ [
    if j = 1 then "\n" else "" endif ++
    if fix(queens[i]) = j then
        show_int(2,j)
    else
        "  "
    endif
| i in 1..n, j in 1..n
] ++ ["\n"];

n = 3;

```

Related work.

The presentation of constraint satisfaction problems as relational languages is new to our knowledge, but pursues the tradition of defining a constraint satisfaction problem as a bipartite graph relating variables with constraints [61]. The definition of constraint propagators as closure operators and the link between domain representations and abstractions is borrowed from [14].

Language frameworks for modeling search have been proposed as monadic constructs [77], ad-hoc solver extensions with combinators [79], or extensions of MiniZinc with meta-search operators [70]. Our proposal is singular as it does not need to extend the underlying solver (except for indexical support). In particular, this approach demonstrates that reified constraint gives to constraint satisfaction problems enough expressiveness to describe any tree search procedures.

Chapter 4

Angelic Programming: from compound guards to atomic kernel for LLCC

The first section introduces derivation nets, that give a compact representation of all the consequences of an agent. The observation of all the consequences of an agent is the angelic semantics, and the main consequences of angelic programming, in particular the collapse of the guard semantics hierarchy, are given in the second section. The third section gives some applications to angelic semantics. It is in particular remarkable that some particular forms of decomposition, with *ask* and *tell*, can even be used in the committed-choice settings of CHR, thanks to CHR propagation rules. The fourth section presents the design of the SiLCC programming language.

4.1 Angelic Semantics and Derivation nets

4.1.1 Process calculi

Definition 39. A *process algebra* $(\mathcal{L}, \otimes, \mathbf{1})$ is a free abelian monoid where

- the set \mathcal{L} is the *language*, its elements a_i are the *agents*,
- the composition law \otimes is the *parallel composition*.

The main point of this definition is introducing the parallel composition as operator for building processes as multisets of *atomic agents* in the style of the *chemical abstract machine* framework CHAM [8].

Definition 40. An *atomic agent* is an agent which is neither $\mathbf{1}$, nor of the form $a_1 \otimes a_2$, with $a_1, a_2 \in \mathcal{L}$.

Let \mathcal{L}_0 be the set of atomic agents. The fact that a process algebra is a free abelian monoid is characterised by the following proposition.

Proposition 3. *The application $\otimes : \mathcal{M}_f(\mathcal{L}) \rightarrow \mathcal{L}$ is well-defined and its restriction on $\mathcal{M}_f(\mathcal{L}_0)$ is one-to-one.*

$$\{m_1 \otimes \cdots \otimes m_n\} \mapsto m_1 \otimes \cdots \otimes m_n$$

Proof. Associativity and commutativity justify the definition and $\mathcal{M}_f(\mathcal{L}_0)$ is one-to-one because the process algebra is free. \square

Hence, every agent is a “*soup*” of atomic agents.

Definition 41. A *process calculus* $(\mathcal{L}, \otimes, \mathbf{1}, \longrightarrow)$ is given by

- a process algebra $(\mathcal{L}, \otimes, \mathbf{1})$,
- a *reduction relation* $\longrightarrow \subset \mathcal{L} \otimes \mathcal{L}$ which is compatible for parallel composition: for all $a, a', b \in \mathcal{L}$, if $a \longrightarrow a'$ then $a \otimes b \longrightarrow a' \otimes b$.

The compatibility of \longrightarrow with respect to \otimes captures the monotonic property of the calculus with respect to the constraint store in concurrent constraint programming [28] and, more generally, the *chemical law* of the chemical abstract machine framework.

Example 20. *Let (C, \Vdash_C) be a linear constraint system. Let \mathcal{A} be the set of $\text{LLCC}(C)$ agents and let \equiv be the structural equivalence between agents. $(\mathcal{A}/\equiv, \otimes, \mathbf{1}, \longrightarrow)$ is a process calculus.*

4.1.2 Nets for sets of reductions

Definition 42. An *oriented hypermultigraph* (V, E, \mathbf{i}) is given by two disjoint sets V and E , and an application $\mathbf{i} : V \times E \cup E \times V \rightarrow \mathbf{N}$ where

- the elements v_i of V are the *vertices*,
- the elements e_j of E are the *edges*,
- the application \mathbf{i} gives the *incidences* between vertices and edges.

Remark 1. Equivalently, an oriented hypermultigraph can be defined as an oriented bipartite graph (V', E') where the edges are weighted by natural numbers $w : E' \rightarrow \mathbf{N}$. The correspondence is given by $V' = V \cup E$ and $E' = \{(a, b) \in V \times E \cup E \times V \mid \mathbf{i}(a, b) \geq 1\}$, where the weight of each edge is $w = \mathbf{i}|_{E'}$.

In an oriented graph (V', E') , every vertice $v \in V'$ has a set of predecessors $\bullet v = \{v' \mid (v', v) \in E'\}$ and a set of successors $v^\bullet = \{v' \mid (v, v') \in E'\}$. When the graph is weighted by $w : E' \rightarrow \mathbf{N}$, predecessors and successors are naturally considered as multisets, the power of each vertice is given by the weight of the edge linking it. In other words, the characteristic applications of $\bullet v$ and v^\bullet are given by

$$\mathbf{1}_{\bullet v} : v' \in V' \mapsto \begin{cases} w(v', v) & \text{si } (v', v) \in E' \\ 0 & \text{sinon} \end{cases}$$

$$\mathbf{1}_{v^\bullet} : v' \in V' \mapsto \begin{cases} w(v, v') & \text{si } (v, v') \in E' \\ 0 & \text{sinon} \end{cases}$$

It is worth noticing that in the case of bipartite graphs where vertices are partitioned into two classes without edges between vertices from the same class, $\bullet v$ and v^\bullet associates to every vertice a (multi)set of vertices from the other class. We use the same notations for multihypergraphs.

Notation 2. Let $\mathcal{H} = (V, E, \mathbf{i})$ be an oriented hypermultigraph. For every vertice $v \in V$, let $\bullet v$ and v^\bullet be the multiset of edges which are predecessors and successors respectively of v in \mathcal{H} . Symmetrically, for every edge $e \in E$, let $\bullet e$ and e^\bullet be the multiset of vertices which are predecessors and successors respectively of e in \mathcal{H} .

Let $(\mathcal{L}, \otimes, \mathbf{1}, \longrightarrow)$ be a process calculus.

Definition 43. A *derivation net* $(V, E, \mathbf{i}, \ell_V)$ is given by an oriented hypermultigraph (V, E, \mathbf{i}) and an application $\ell : V \rightarrow \mathcal{L}$ where

- every vertice $v \in V$ is labelled by an agent $\ell(v) \in \mathcal{L}$
- every edge $e \in E$ is such that $\otimes \ell(\bullet e) \longrightarrow \otimes \ell(e^\bullet)$.

Let $(V, E, \mathbf{i}, \ell_V)$ be a derivation net.

Definition 44. A *marking* m is a multiset of vertices.

Definition 45. The *reduction relation* \longrightarrow is the binary relation $\longrightarrow \subseteq \mathcal{M}_f(V) \times \mathcal{M}_f(V)$ such that $m \longrightarrow m'$ when there exists an edge $e \in E$ satisfying

- $\bullet e \subseteq m$, that is to say for every $v \in V$, $\mathbf{1}_{\bullet e}(v) \leq \mathbf{1}_m(v)$,
- $m' = m \ominus \bullet e \oplus e^\bullet$, that is to say for every $v \in V$, $\mathbf{1}_{m'}(v) = \mathbf{1}_m(v) - \mathbf{1}_{\bullet e}(v) + \mathbf{1}_{e^\bullet}(v)$.

Proposition 4. For all markings m and m' , if $m \longrightarrow m'$ then $\otimes \ell(m) \longrightarrow \otimes \ell(m')$.

Proof. If $m \longrightarrow m'$, then by definition of the reduction relation, there exists an edge $e \in E$ such that $\bullet e \subseteq m$ and $m' = m \ominus \bullet e \oplus e^\bullet$. By definition of the edges of a derivation net, we have $\otimes \ell(\bullet e) \longrightarrow \otimes \ell(e^\bullet)$. Let $m_0 = m \ominus \bullet e$. By compatibility of \longrightarrow with \otimes , we have $\otimes \ell(\bullet e) \otimes \otimes m_0 \longrightarrow \otimes \ell(e^\bullet) \otimes \otimes m_0$, that is to say $\otimes \ell(m) \longrightarrow \otimes \ell(m')$. \square

Angelic execution consists therefore in the iterative construction of such a complete derivation net, keeping only the hyperarcs which are involved in a reachable marking from the initial configuration. Since the exploration is potentially infinite, the exploration should be done in breadth first to give an equal chance of execution to every computation path.

4.1.3 Sharing strategies

Derivation nets do not structurally force any sharing to reduce scheduling non-determinism. This is typically the case for *simpagation* rules (we adopt CHR syntax for conveniency without loss of generality): a simpagation rule is of the form $n@H_1 \setminus H_2 \Leftrightarrow B$ where H_1 is a *persistent head*. Such a rule has the same logical interpretation as $n@H_1, H_2 \Leftrightarrow H_1, B$. Two firings of simpagation rules with a common persistent head can lead to two computation paths whether which rule is fired before the other.

Consider the following rules describing the calculus of a two-dimensional scalar product with a concurrent product.

Example 21 (Concurrent two-dimensional scalar product).

$$\begin{aligned}
& \mathit{init} \ @ \ \mathit{scalar}(X1, Y1, X2, Y2, P) \Leftrightarrow \\
& \quad \mathit{product}(X1, X2, X), \mathit{product}(Y1, Y2, Y), \mathit{sum}(X, Y, P) . \\
& \mathit{product} \ @ \ \mathit{product}(A, B, C) \Leftrightarrow \\
& \quad V \ \mathbf{is} \ A * B, \mathit{value}(C, V). \\
& \mathit{sum} \ @ \ \mathit{sum}(A, B, C), \mathit{value}(A, VA), \mathit{value}(B, VB) \Leftrightarrow \\
& \quad V \ \mathbf{is} \ VA + VB, \mathit{value}(C, V) .
\end{aligned}$$

There are essentially two possible derivations in the abstract semantics from a query $\mathit{scalar}(X1, Y1, X2, Y2, P)$, revealing scheduling non-determinism, depending upon which of the two products is evaluated first: $\mathit{product}(X1, X2, X)$ or $\mathit{product}(Y1, Y2, Y)$.

The following derivation net shows the quotiented derivation path for the scalar product.

This non-determinism can be reduced considering the derivation net where each simpagation rule is a hyperarc such that the vertex of the persistent head is the same both for input and output.

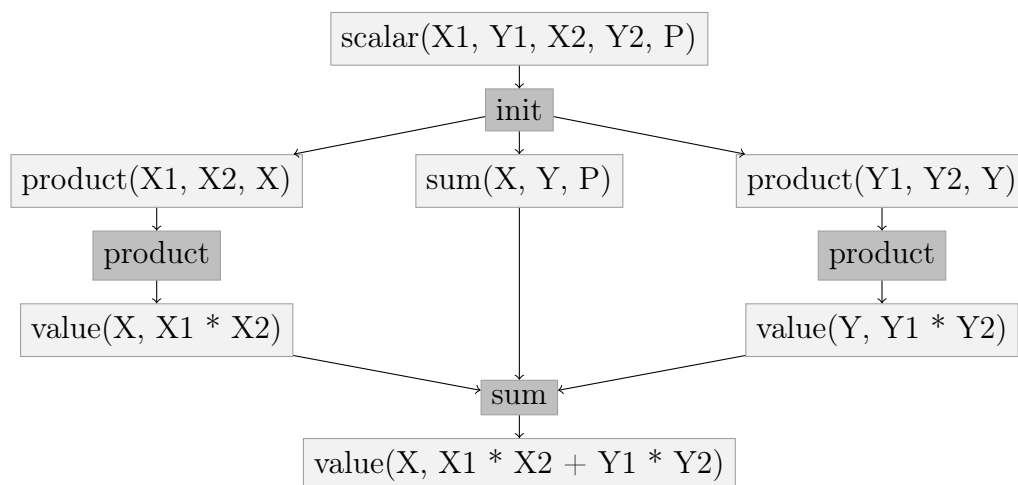


Figure 4.1: Derivation net for the query $\text{scalar}(X1, Y1, X2, Y2, P)$ with the scalar product program (Example 21)

Proposition 5. *The complete derivation net where all hyperarcs are reachable and all equal user constraints are shared to the same vertex can be iteratively constructed by solving an EXPSPACE-complete problem for each new hyperarc.*

Proof. The iterative construction of such a derivation net can be done by sharing all equal user constraints to the same vertex: interpreting the derivation net as a Petri net, testing the reachability of a hyperarc reduces to testing the reachability in a Petri net, which is decidable [60] but computationally expensive [52]. \square

However, in the context of derivation nets where all cycles are trivial, that is to say that every cycle consist in only one hyperarc having some vertices both as input and output, then there exists a log-linear algorithm to decide the reachability. The case of trivial cycles is particularly important as those cycles appear naturally when considering simpagation rules.

- Preparation: at each new vertex creation v_0 , computes a table $t(v_0)$ which associates each ancestor vertice v (outside trivial cycles) to its potential immediate successor hyperarc e (there is at most one!): $t(v_0) : v \mapsto e$ With balanced binary trees, logarithmic time cost for each vertex.
- To check if a binary hyperarc e_0 between v_0 and v_1 introduces a conflict:
 1. choose one of the predecessor vertex, say v_0 , (preferably the one with least ancestors)




Figure 4.2: Derivation net without sharing for the query $H1, H2, H2'$ with two simpagation rules $n @ H_1 H_2 \Leftrightarrow B$ and $n' @ H_1 H_2' \Leftrightarrow B'$




Figure 4.3: Derivation net with sharing for the query $H1, H2, H2'$ with two simpagation rules $n @ H_1 H_2 \Leftrightarrow B$ and $n' @ H_1 H_2' \Leftrightarrow B'$

2. let $t \leftarrow t(v_1) + (v_1 \mapsto e_0)$
3. begin with $v \leftarrow v_0$,
4. for each predecessor vertex v' of each predecessor hyperarc e of v ,
5. if $t(v')$ is defined, succeeds if $t(v') = e$ or v' in trivial cycle, else fails,
6. if not, let $t \leftarrow t + (v' \mapsto e)$ and recursively go to 4 for $v \mapsto v'$.

In worst case, logarithmic cost (table search) for each ancestor.

In practice, either hyperarcs between neighbours ($t(v_0)$ is often defined) or hyperarcs between a vertex and a top-level ask (with few ancestors).

This algorithm gives a polynomial construction for completing the derivation net at each execution step.

4.1.4 Derivation nets in presence of Propagation Rules

As far as the abstract semantics is concerned, propagation rules $n@H \Rightarrow G|B_b, B_u$ are shortcuts for $n@H \Leftrightarrow G|H, B_b, B_u$: the head is restored after firing the rule.

This non-consumption leads to trivially infinite computation paths that can be avoided with the ω_1 semantics [11]. In terms of derivation nets, distinguishing the set of vertices between linear and persistent constraints make construction rules of derivation nets being refined to prevent this trivial non-termination. The strict output vertices of a hyperarc are marked as persistent if and only if all the input vertices are persistent or if it is a propagation (that is to say, if all input vertices are in a trivial cycle).

4.2 Angelic Programming

4.2.1 Head Decomposability

In the abstract semantics, a CHR rule can only observe the presence of a user constraint, not the absence: firing is monotonic relatively to the store. As a consequence, if observations are restricted to the side effects of the firing of some rules, silent partial consumption of the head of these rules cannot be observed. Such an observable is motivated by the fact that the body of fired rules is the place where side effects can happen. In particular, multiple headed rules can be rewritten in 2-headed rules by the introduction of fresh intermediary user constraints (carrying the context variables if any).

The rule

$$a, b, c, d \Leftrightarrow \text{print}(\text{"side effect"})$$

and the set of rules

$$\begin{aligned} a &\Leftrightarrow f1 \\ f1, b &\Leftrightarrow f2 \\ f2, c &\Leftrightarrow f3 \\ f3, d &\Leftrightarrow \text{print}(\text{"side effect"}) \end{aligned}$$

are equivalent provided that $f1$, $f2$, $f3$ are fresh user constraints that do not appear elsewhere neither in the program nor in the initial goal.

This equivalence does not hold in general with a committed-choice scheduler since premature consumptions of a and b can prevent other rules to be fired even if c and d never appear. On the contrary, premature consumptions in angelic settings will only lead to blocking computation branches that will not prevent other branches to be explored. This property can benefit to the implementation: only 2-headed rules have to be considered. More precisely, all CHR rules can be translated to rules with two heads where one of them is an intermediary user constraint.

Such a translation makes trivial cycles of simpagations become non-trivial: the algorithm presented above can nevertheless be adapted in this case, since

all hyperarcs involved in the cycle only introduce intermediary user constraints. Therefore, a hyperarc cannot be unreachable due to the consumption of such constraints by another rules.

4.2.2 Controlling the Angelism

User constraints can be introduced to explicitly sequence the execution of rules. The following program produces as side-effect a unspecified permutation of a, b, c when launched with the goal start.

```
start ⇔ a, b, c.
a ⇔ print("a").
b ⇔ print("b").
c ⇔ print("c").
```

The order can be fixed by the introduction of fresh intermediary constraints to mark the step of the sequence (carrying the context variables if any).

```
start ⇔ s0, a, b, c.
a, s0 ⇔ s1, print("a").
b, s1 ⇔ s2, print("b").
c, s2 ⇔ print("c").
```

Operationally, such an explicit sequencing forces the derivation net to have a linear path instead of branching hyperarcs. Formally, the sequence is coded declaratively in the logic instead of being left to the conventional implementation of the comma sequence operator.

4.3 Applications

4.3.1 Angelic CHR \forall

Angelic execution can be seen as a search among scheduling. Therefore, CHR \forall rules where there can be multiple bodies, leading to a search for a successful one, can be encoded as multiple rules with the same head.

The rule $N @ H \Leftrightarrow G \mid B_1 ; \dots ; B_n$. is encoded as the set of n rules $N_1 @ H \Leftrightarrow G \mid B_1$., ..., $N_n @ H \Leftrightarrow G \mid B_n$. Angelic execution ensures that consequences of B_1 , ..., B_n are explored.

4.3.2 Meta-interpreters

The decomposability of heads allow CHR meta-interpreters to be conveniently written. Suppose that a rule is coded with a user constraint rule $(N @ H \Leftrightarrow G \mid B)$

where H is a list of heads and with a proper encoding for the body B where user constraints are marked with the functor $ucstr$, then the following rules code a meta-interpreter.

```

first_head @ rule(_N @ H  $\Leftrightarrow$  G, B), ucstr(H0)  $\Leftrightarrow$ 
  copy_term((H, G, B), ([H0 | T0], G0, B0)) |
  match(T0, G0, B0).
matching_end @ match([], G, B)  $\Leftrightarrow$  call(G) |
  call(B).
matching_cont @ match([H | T], G, B), ucstr(H)  $\Leftrightarrow$ 
  match(T, G, B).

```

This meta-interpreter uses the decomposability of rules with $match$ as intermediary user constraint.

Example 22. *The scalar product example (Example 21) is encoded into the following constraints. A derivation net for the meta-interpretation of this program is given in Fig. 4.4.*

```

rule(init @ [scalar(X1, Y1, X2, Y2, P)]  $\Leftrightarrow$  true |
  ucstr(product(X1, X2, X)),
  ucstr(product(Y1, Y2, Y)),
  ucstr(sum(X, Y, P))).
rule(product @ [product(A, B, C)]  $\Leftrightarrow$  true |
  V is A * B,
  ucstr(value(C, V))).
rule(sum @ [sum(A, B, C), value(A, VA), value(B, VB)]  $\Leftrightarrow$  true |
  V is VA + VB,
  ucstr(value(C, V))).

```

4.3.3 User-defined Indexation

Thanks to the dependency book-keeping operated by derivation nets, rules can reformulate user constraints to another form while keeping the dependency relation between the original user constraint and the reformulation. For instance, suppose that the underlying implementation only makes indexing on the principal functor of the user constraint arguments. The following rule decomposes a nested term in a user constraint into another user constraint with this term as root argument.

$$c(f(X)) \Leftrightarrow cf(X) .$$

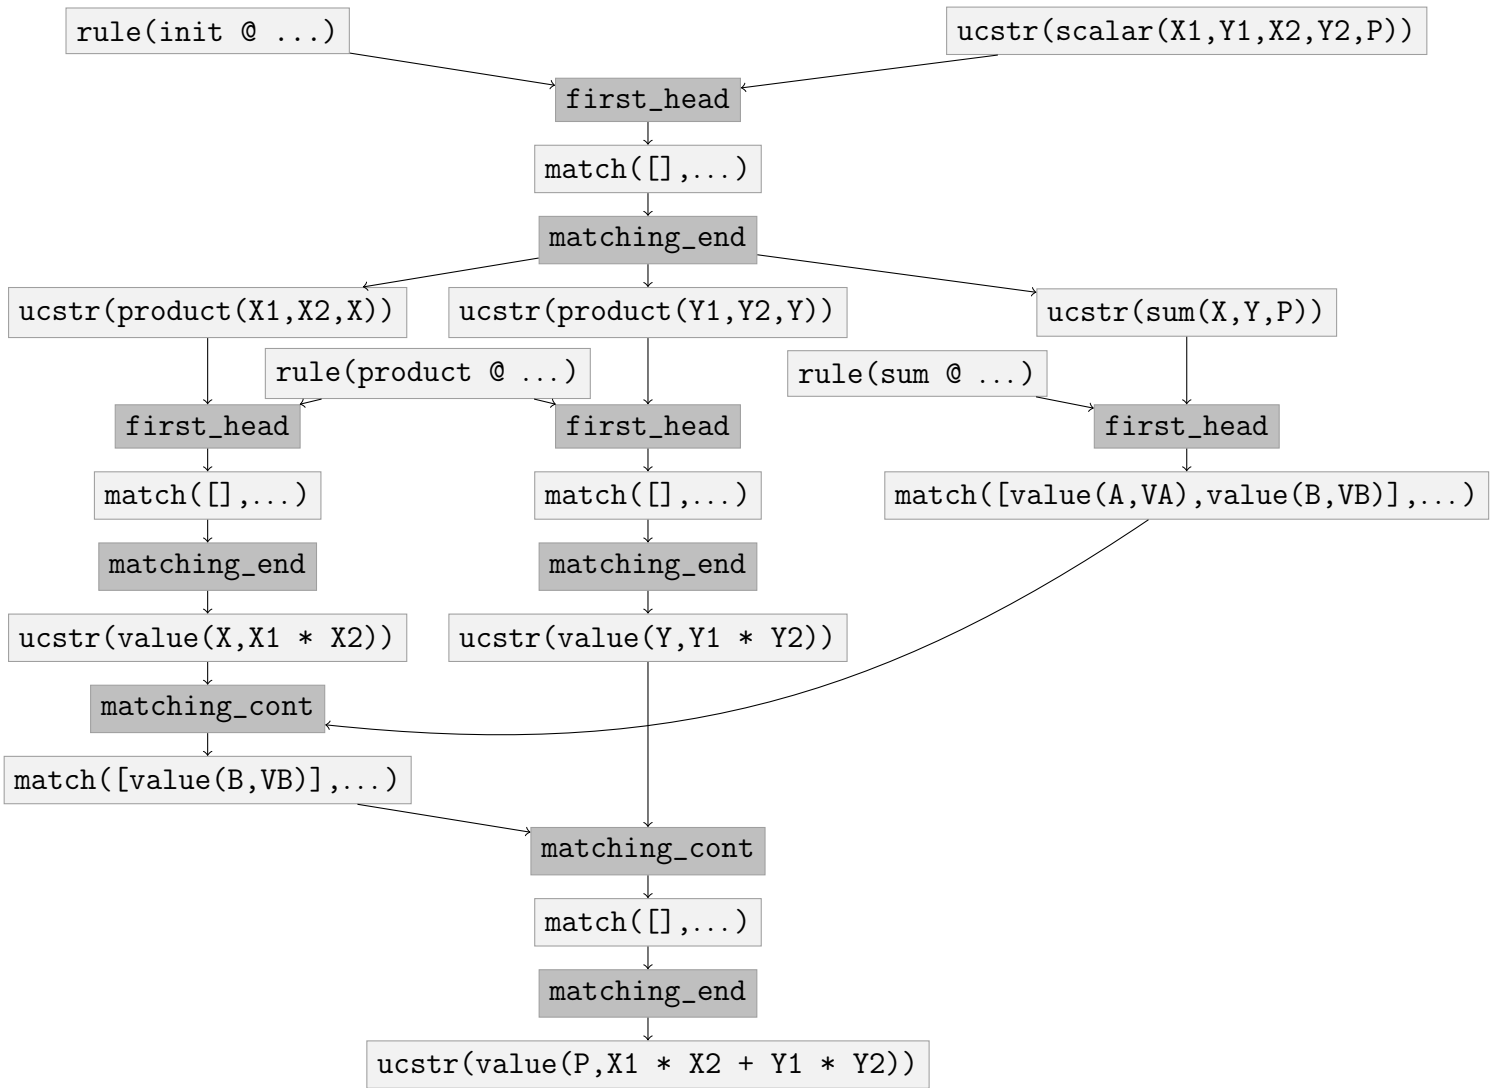


Figure 4.4: Meta-interpretation of the query `ucstr(scalar(X1,Y1,X2,Y2,P))` with the meta-interpreted program given in Example 22

Then consuming $cf(X)$ in another rule is equivalent to consuming $c(f(X))$: therefore, rules having heads matching on $c(X)$ in general and rules having heads matching on $cf(X)$ in particular can coexist and consume observationnaly the same user constraints.

4.3.4 CHR with Ask and Tell

In this section, we apply the decomposability of guards for CHR. Indeed, one main drawback of CHR is the absence of *modularity*. Once a constraint system is defined in CHR with some built-in constraints, this constraint system cannot be reused in another CHR program taking the defined constraints as new built-in constraints. The reason for this difficulty is that a CHR program defines a satisfiability check but not the *constraint entailment* check that is required in guards. We make use of a particular case of decomposability to propose a *modular* version of CHR, called CHR with *ask* and *tell*. Guards are decomposed into an *ask* part, that triggers the computation of an entailment check that should lead to the *tell* part, that fires the goal. By using the particular semantics of CHR propagation, we prove that this decomposition can be used in the committed-choice settings of CHR.

Previous approaches to this problem have studied conditions under which one can derive automatically an entailment check from a satisfiability check. In [75] such conditions are given based on the logical equivalence:

$$\mathcal{X} \models c \rightarrow d \text{ if and only if } \mathcal{X} \models (c \wedge d) \leftrightarrow c$$

The programming discipline proposed in CHRat for programming modular constraint solvers is to enforce, for each user-introduced constraint, the definition of simplification and propagation rules to rewrite *ask* control tokens into *entailed* control tokens. Solvers for *asks* and *tells* are required for the implementation of built-in constraint system [26]; the discipline we propose consists in the internalization of this requirement in the CHR solver itself. A guard constraint $\text{genericconstraint}(\text{argumentvector})$ in a rule instance R is *operationally entailed* in a constraint store containing the control token $\text{ask}(K, \text{genericconstraint}(\text{argumentvector}))$ when its solved form contains the token $\text{entailed}(K, \text{genericconstraint}(\text{argumentvector}))$, where K is a new variable used to associate these control tokens to the rule R . This allows us to program arbitrarily complex entailment checks with rules, instead of event-driven imperative programs [24]. With this programming discipline, CHRat constraints can be reused both in rules and guards in other components to define new solvers.

CHRat Components for `leq/2` and `min/3`

We first consider the classical CHR program defining an ordering constraint.

Example 23. *The satisfiability solver is defined by the following four rules. The first three rules translate the axioms for ordering relations, and the last rule eliminates duplicates in leq constraints.*

$$\begin{aligned} \text{leq}(X, X) &\Leftrightarrow \text{true}. \\ \text{leq}(X, Y), \text{leq}(Y, X) &\Leftrightarrow X = Y. \\ \text{leq}(X, Y), \text{leq}(Y, Z) &\Rightarrow \text{leq}(X, Z). \\ \text{leq}(X, Y) \mid \text{leq}(X, Y) &\Leftrightarrow \text{true}. \end{aligned}$$

In CHRat, a constraint solver must define rules for checking entailment. The entailment checking rules for $\text{leq}(X, Y)$ rewrite the token $\text{ask}(K, \text{leq}(X, Y))$ into the token $\text{entailed}(K, \text{leq}(X, Y))$. K is a variable which associates the control tokens to the rule instances that generated them in order to avoid token capture by other rules with the same guards. In this example, since the store is transitively closed, checking $\text{leq}(X, Y)$ is directly observable in the store with a rule if $X \neq Y$. The reflexive case is handled with an additional rule.

$$\begin{aligned} \text{leq}(X, Y) \mid \text{ask}(K, \text{leq}(X, Y)) &\Leftrightarrow \text{entailed}(K, \text{leq}(X, Y)). \\ \text{ask}(K, \text{leq}(X, X)) &\Leftrightarrow \text{entailed}(K, \text{leq}(X, X)). \end{aligned}$$

The satisfiability solver and the entailment solver together define a CHRat component for the constraint $\text{leq}(X, Y)$. The modularization of CHRat relies on a simple atom-based component name separation mechanism: exported CHR-constraints are prefixed with the component name and internal CHR-constraints are prefixed so as to avoid collisions [40]. Such a component can be used to define new solvers using the $\text{leq}(X, Y)$ constraint both in rules and guards.

Example 24. *For instance, a component for the minimum constraint $\text{min}(X, Y, Z)$, stating that Z is the minimum value among X and Y , is definable as follows:*

$$\begin{aligned} &\text{component } \text{min_solver}. \\ &\text{import } \text{leq}/2 \text{ from } \text{leq_solver}. \\ &\text{export } \text{min}/3. \\ &\text{min}(X, Y, Z) \Leftrightarrow \text{leq}(X, Y) \mid Z=X. \\ &\text{min}(X, Y, Z) \Leftrightarrow \text{leq}(Y, X) \mid Z=Y. \\ &\text{min}(X, Y, Z) \Rightarrow \text{leq}(Z, X), \text{leq}(Z, Y). \\ \\ &\text{min}(X, Y, Z) \mid \text{ask}(K, \text{min}(X, Y, Z)) \Leftrightarrow \text{entailed}(K, \text{min}(X, Y, Z)). \\ &\text{ask}(K, \text{min}(X, Y, X)) \Leftrightarrow \text{leq}(X, Y) \mid \text{entailed}(K, \text{min}(X, Y, Z)). \\ &\text{ask}(K, \text{min}(X, Y, Y)) \Leftrightarrow \text{leq}(Y, X) \mid \text{entailed}(K, \text{min}(X, Y, Z)). \end{aligned}$$

Transformation to a Flat CHR Program

In CHR, the guards are restricted to built-in constraints [30]. CHRat programs are transformed into CHR programs with a program transformation which:

- removes all the user defined constraints from guards;
- renames the predicates $\text{ask}(c(\dots))$ in $\text{ask_c}(K, \dots)$ and $\text{entailed}(c(\dots))$ in $\text{entailed_c}(K, \dots)$ with an extra argument for the association variable K ;
- introduces for each CHRat rule an id_n CHR constraint which links the association variable with the instances of the variables of the head.

For example, the min CHRat solver is transformed to the following CHR program:

```

min(X,Y,Z)\ask_min(K,X,Y,Z) $\Rightarrow$ entailed_min(K,X,Y,Z).
min(X,Y,Z) $\Rightarrow$ ask_leq(K,X,Y),id1(K,X,Y,Z).
id1(X,Y,Z,K),min(X,Y,Z),entailed_leq(K,X,Y) $\Leftrightarrow$ Z=X.
min(X,Y,Z) $\Rightarrow$ ask_leq(K,Y,X),id2(K,X,Y,Z).
id2(X,Y,Z,K),min(X,Y,Z),entailed_leq(K,Y,X) $\Leftrightarrow$ Z=Y.
min(X,Y,Z) $\Rightarrow$ leq(Z,X),leq(Z,Y).
ask_min(X,Y,X,K) $\Rightarrow$ ask_leq(K0,X,Y),id3(K0,X,Y,K).
id3(K0,X,Y,K),ask_min(K,X,Y,X),entailed_leq(K0,X,Y) $\Leftrightarrow$ 
    entailed_min(K,X,Y,X).
ask_min(X,Y,Y,K) $\Rightarrow$ ask_leq(K0,Y,X),id4(K0,X,Y,K).
id4(K0,X,Y,K),ask_min(X,Y,Y,K),entailed_leq(K0,Y,X) $\Leftrightarrow$ 
    entailed_min(K,X,Y,Y).

```

Existentially Quantified Variables in Guards

In CHRat, as in CHR, variables that appear in a guard without appearing in the head of the rule, need a special treatment. For instance, consider the following CHRat rule for eliminating non minimal elements:

$$\text{number}(A), \text{number}(B) \Leftrightarrow \min(A, B, C) \mid \text{number}(C).$$

The ask-solver defined above will never detect the entailment of the guard.

In CHRat, the support for existentially quantified variables in guards is performed with control tokens $\text{exists}(K, V_1, \dots, V_n)$ where the V_i 's are existentially quantified variables in the guard and K is the rule association variable. The rules for $\text{ask}(\min)$ of example 24 must thus be completed with extra rules for instantiating the existentially quantified variables as follows:

```

ask(K,min(X,Y,Z)),exists(K,Z)  $\Leftrightarrow$  leq(X,Y) |
    Z=X, entailed(K, min(X,Y,Z)).
ask(K,min(X,Y,Z)),exists(K,Z)  $\Leftrightarrow$  leq(Y,X) |
    Z=Y, entailed(K, min(X,Y,Z)).
min(X,Y,M) \ask(K, min(X,Y,Z)), exists(K,Z)  $\Leftrightarrow$ leq(X,Y) |
    Z=M, entailed(K, min(X,Y,Z)).

```

Each of these rules adds the entailed constraint to the store. This generalizes the solution proposed in [2] for the operational semantics of CHR for built-in constraints, to user-defined constraints.

Syntax and Semantics of $\text{CHRat}(\mathcal{X})$ Components

Syntax.

In CHRat , control tokens are built with extra symbols with the signature $\Pi_t = \{\text{ask}/2, \text{exists}/2, \text{entailed}/2\}$ disjoint from Π and $\Pi_{\mathcal{X}}$.

Definition 46. A $\text{CHRat}(\mathcal{X})$ program P is a sequence of rules $H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u$ where H is a multiset of elements of \mathcal{A} ; H' is a multiset of elements of $\mathcal{A} \cup \text{ask}(V, \mathcal{A}) \cup \text{exists}(V, V)$; G is a formula of \mathcal{L}' where \mathcal{L}' is the extension of \mathcal{L} containing atomic propositions \mathcal{A} and closed by conjunction and existential quantification; $B_{\mathcal{X}}$ is a formula of \mathcal{L} ; and B_u is a multiset of elements of $\mathcal{A} \cup \text{entailed}(V, \mathcal{A})$. H and H' cannot be both empty.

As formula of \mathcal{L}' , a guard G is of the form $G_u \wedge G_{\mathcal{X}}$ where G_u is a conjunction (or multiset) of tokens and $G_{\mathcal{X}}$ is a formula of \mathcal{L} . The set of variables that only appear in G_u is denoted $V_G = \text{fv}(G_u) \setminus \text{fv}(H, H', G_{\mathcal{X}})$.

Operational Semantics.

Configurations are tuples (T, c, I) :

- the built-in store is a constraint $c \in \mathcal{C}$;
- the user store T is a multiset of $\mathcal{A} \cup \text{ask}(V, \mathcal{A}) \cup \text{exists}(V, V) \cup \text{entailed}(V, \mathcal{A})$;
- pending instances table I is a finite support map from variables to rule instances $(\mathbf{r}, \sigma_{\pi})$ where \mathbf{r} is a (renamed) rule and σ_{π} a permutation of \mathcal{A} .

For any pending instances table I , $\text{sup}(I)$ denotes the support of I . If $K \in V$ is such that $K \notin \text{sup}(I)$, then $I \uplus (\mathbf{r}, \sigma_{\pi})_K$ denotes the map I' such that $\text{sup}(I') = \text{sup}(I) \cup \{K\}$ with $I'|_{\text{sup}(I)} = I$ and $I'(K) = (\mathbf{r}, \sigma_{\pi})$.

Two kind of transitions can occur from (T, c, I) : suspend $\xrightarrow{\text{s}}$ and wake $\xrightarrow{\text{w}}$.

$$(T, c, I) \xrightarrow{\text{s}} ((T \oplus \text{ask}(K, G_u) \oplus \text{exists}(K, V_G)), \langle c \wedge \exists(c_m) \rangle, \langle I \uplus (\mathbf{r}, \sigma_{\pi})_K \rangle)$$

where $\mathbf{r} = \langle H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u \rangle$ is a rule in P renamed with fresh variables with respect to (T, c, I) and π an injection from $T' = H \oplus H'$ to T such that the matching constraint $c_m = \langle (\bigwedge_{t \in \text{sup}(T')} t = \sigma_{\pi}(t)) \wedge G_{\mathcal{X}} \rangle$ is entailed, i.e. $\mathcal{X} \models \forall \mathbf{x}(c \rightarrow \exists \mathbf{y}(c_m))$

where $\mathbf{x} = \text{fv}(c)$ and $\mathbf{y} = \text{fv}(c_m) \setminus \text{fv}(c)$. Then the reduction adds the new pending instance (\mathbf{r}, σ_π) to I , indexed by a fresh variable K .

$$(T, c, \langle I \uplus (\mathbf{r}, \sigma_\pi)_K \rangle) \xrightarrow{\text{w}} (\langle T \ominus \pi'(H') \ominus \pi'(\text{entailed}(K, G_u)) \oplus B_u \rangle, \langle c \wedge c_m \wedge B_{\mathcal{X}} \rangle, I)$$

where (\mathbf{r}, σ_π) is a pending instance of the rule $\mathbf{r} = \langle H \setminus H' \Rightarrow G \mid B_{\mathcal{X}} \wedge B_u \rangle$ and σ_π represents an injection π from $H \oplus H'$ to T , such that there exists an injection π' from $T' = H \oplus H' \oplus \text{entailed}(K, G_u)$ to T , which extends π such that the matching constraint: $c_m = \langle (\bigwedge_{t \in \text{sup}(T')} t = \sigma_{\pi'}(t)) \wedge G_{\mathcal{X}} \rangle$ is entailed, i.e. $\mathcal{X} \models \forall \mathbf{x}(c \rightarrow \exists \mathbf{y}(c_m))$ where $\mathbf{x} = \text{fv}(c)$ and $\mathbf{y} = \text{fv}(c_m) \setminus \text{fv}(c)$.

We define $\rightarrow = \xrightarrow{\text{s}} \cup \xrightarrow{\text{w}}$.

Definition 47. The *operational semantics* for the observation of accessible and terminal configurations for a query $q = T_0 \wedge c_0 \in \mathcal{L}'$ are, with $\mathbf{x} = \text{fv}(T_1, c) \setminus \text{fv}(q)$:

$$\begin{aligned} \mathcal{O}_P^a[q] &= \{\exists \mathbf{x}(T_1 \wedge c) \mid (T_0, c_0, \emptyset) \xrightarrow{*} (T, c, _) , T_1 = T \cap \mathcal{A}\} \\ \mathcal{O}_P^t[q] &= \{\exists \mathbf{x}(T_1 \wedge c) \mid (T_0, c_0, \emptyset) \xrightarrow{*} (T, c, _) \not\rightarrow, T_1 = T \cap \mathcal{A}\} \end{aligned}$$

The accessible store semantics $\mathcal{O}_P^a[q]$ is related to the logical semantics in the next section. The terminal stores semantics $\mathcal{O}_P^t[q]$ captures the resulting stores once no more rules can be applied. Observed stores are restricted to \mathcal{A} . Let min be the solver defined in example 24, then $\mathcal{O}_{\text{min}}^t[\text{min}(X, Y, Z), \text{leq}(X, Y)] = \{\text{leq}(X, Y) \wedge Z = X\}$ and the token $\text{ask}(\text{leq}(Y, X))$ introduced in the unsuccessful guard checking (for the second rule of the leq component) is not exposed.

Logical Semantics.

We define a transformation $(\cdot)^*$ from control tokens and suspended instances to atomic propositions:

$$\begin{aligned} (\text{ask}(K, p(\vec{X})))^* &= \text{ask}_p(K, \vec{X}) \\ (\text{entailed}(K, p(\vec{X})))^* &= \text{entailed}_p(K, \vec{X}) \\ (\text{exists}(K, V))^* &= \text{exists}(K, V) \\ (((H \setminus H' \Rightarrow G \mid B), \sigma_\pi)_K)^* &= \text{id}_{\mathbf{r}}(K, \vec{V}(\sigma_\pi(H \oplus H'))) \end{aligned}$$

In classical logic, a multiset M of constraints is interpreted by the conjunction $M^\ddagger = \langle \bigwedge_{c \in \text{sup}(M)} c^* \rangle$ of its elements. A rule is logically interpreted as:

$$(H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u)^\ddagger = (H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u)_s^\ddagger \wedge (H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u)_w^\ddagger$$

where sub-formulae suspend and wake translate their operational counterpart:

$$\begin{aligned} (H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u)_{\text{s}}^{\dagger} &= \left\langle \forall \mathbf{x} \left(\exists \mathbf{y} (G) \wedge H^{\dagger} \wedge H'^{\dagger} \rightarrow \right. \right. \\ &\quad \left. \left. \exists K (\text{id}(K, \mathbf{z}) \wedge \text{exists}(K, V_G) \wedge (\bigwedge_{p(\mathbf{u}) \in \text{sup}(G_u)} \text{ask}_p(K, \mathbf{u}))) \right) \right\rangle \\ (H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, \wedge B_u)_{\text{w}}^{\dagger} &= \left\langle \forall \mathbf{x} \forall K (\exists \mathbf{y} (G) \wedge H^{\dagger} \rightarrow (H'^{\dagger} \wedge \text{id}(K, \mathbf{z}) \wedge \right. \\ &\quad \left. (\bigwedge_{p(\mathbf{u}) \in G_u} \text{entailed}_p(K, \mathbf{u})) \leftrightarrow \exists \mathbf{z} (G \wedge B_{\mathcal{X}} \wedge B_u^{\dagger}))) \right\rangle \end{aligned}$$

and $\mathbf{x} = \text{fv}(H, H')$, $\mathbf{y} = \text{fv}(G) \setminus \text{fv}(H, H')$ and $\mathbf{z} = \text{fv}(G, B_{\mathcal{X}}, B_u) \setminus \text{fv}(H, H')$. The logical interpretation of a program $(P)^{\dagger}$ is the logical conjunction of the interpretations of the rules of P .

Definition 48. The *classical logical semantics* of a query $q \in \mathcal{L}'$ is:

$$\mathcal{L}_P[q] = \{c \in \mathcal{L}' \mid (P)^{\dagger} \models_{\mathcal{X}} q \rightarrow c\}$$

Let V be the free variables of the initial query $q = T_0 \wedge c_0 \in \mathcal{L}'$. Configurations are interpreted as $(T, c, I)^{\dagger} = \langle \exists \mathbf{x} (T^{\dagger} \wedge c \wedge (\bigwedge_{K \in \text{sup}(I)} I(K)^*)) \rangle$ where $\mathbf{x} = \text{fv}(T, c) \setminus \text{fv}(T_0, c_0)$: the query variables appear in the observables and are therefore left free in the interpretation. This definition is extended for a set of configurations S : $S^{\dagger} = \{(T, c, I)^{\dagger} \mid (T, c, I) \in S\}$.

With the same restriction as for CHR, CHRat operational semantics is sound w.r.t. the classical logical semantics. This result follows from the following lemma:

Lemma 6. *If $c_0 \rightarrow c_1$, then $(P)^{\dagger} \models_{\mathcal{X}} (c_0)^{\dagger} \rightarrow (c_1)^{\dagger}$.*

Proof. If $c_0 \xrightarrow{\text{s}} c_1$, then, by definition of $\xrightarrow{\text{s}}$, there is a rule $\mathbf{r} = \langle H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u \rangle$ in P renamed with fresh variables with respect to $c_0 = (T, c, I)$. Moreover, $c_1 = \langle (T \oplus \text{ask}(K, G_u) \oplus \text{exists}(K, V_G)), \langle c \wedge \exists(c_m) \rangle, \langle I \uplus (\mathbf{r}, \sigma_{\pi})_K \rangle \rangle$ where π is an injection from $T' = H \oplus H'$ to T such that the matching constraint $c_m = \langle (\bigwedge_{t \in \text{sup}(T')} t = \sigma_{\pi}(t)) \wedge G_{\mathcal{X}} \rangle$ is entailed. By definition of $(\cdot)^{\dagger}$ for configurations, $(c_0)^{\dagger} = \langle \exists \mathbf{x} (T^{\dagger} \wedge c \wedge (\bigwedge_{K \in \text{sup}(I)} I(K)^*)) \rangle$ and $(c_1)^{\dagger} = \langle \exists \mathbf{x} \exists K (T^{\dagger} \wedge \text{ask}(K, G_u)^{\dagger} \wedge \text{exists}(K, V_G)^{\dagger} \wedge c \wedge (\bigwedge_{K \in \text{sup}(I)} I(K)^*) \wedge (\mathbf{r}, \sigma_{\pi})_K^*) \rangle$. If $(c_0)^{\dagger}$, let \mathbf{x} such that $T^{\dagger} \wedge c \wedge (\bigwedge_{K \in \text{sup}(I)} I(K)^*)$. Admitting $(\mathbf{r})_{\text{s}}^{\dagger}$, we have $c_m \rightarrow (\mathbf{r}, \sigma_{\pi})_K^* \wedge \text{exists}(K, V_G)^{\dagger} \wedge \text{ask}(K, G_u)^{\dagger}$. Since $(c_0)^{\dagger} \rightarrow c_m$, we have $(c_1)^{\dagger}$.

If $c_0 \xrightarrow{\text{w}} c_1$, then, by definition of $\xrightarrow{\text{w}}$, there is a pending instance of the rule $\mathbf{r} = \langle H \setminus H' \Rightarrow G \mid B_{\mathcal{X}} \wedge B_u \rangle$ with fresh variables with respect to $c_0 = (T, c, I \uplus (\mathbf{r}, \sigma_{\pi})_K)$. Moreover $c_1 = \langle (T_1 \oplus B_u), \langle c \wedge c_m \wedge B_{\mathcal{X}} \rangle, I \rangle$ with $T_1 = T \ominus \pi'(H') \ominus \pi'(\text{entailed}(K, G_u))$ and σ_{π} represents an injection π from $H \oplus H'$ to T , such that there exists an injection π' from $T' = H \oplus H' \oplus \text{entailed}(K, G_u)$ to T , which extends π such that the matching constraint: $c_m = \langle (\bigwedge_{t \in \text{sup}(T')} t = \sigma_{\pi'}(t)) \wedge G_{\mathcal{X}} \rangle$ is entailed. By definition of $(\cdot)^{\dagger}$ for configurations, $(c_0)^{\dagger} = \langle \exists \mathbf{x} (T_1^{\dagger} \wedge \pi'(H')^{\dagger} \wedge$

$\pi'(\text{entailed}(K, G_u))^{\ddagger} \wedge c \wedge (\bigwedge_{K \in \text{sup}(I)} I(K)^* \wedge (\mathbf{r}, \sigma_\pi)_K^*)$ and $(c_1)^{\ddagger} = \langle \exists \mathbf{x}(T_1^{\ddagger} \wedge B_u^{\ddagger} \wedge c \wedge c_m \wedge B_{\mathcal{X}}^{\ddagger}(\bigwedge_{K \in \text{sup}(I)} I(K)^*)) \rangle$. If $(c_0)^{\ddagger}$, let \mathbf{x} such that $T_1^{\ddagger} \wedge \pi'(H')^{\ddagger} \wedge \pi'(\text{entailed}(K, G_u))^{\ddagger} \wedge c \wedge (\bigwedge_{K \in \text{sup}(I)} I(K)^* \wedge (\mathbf{r}, \sigma_\pi)_K^*)$. Admitting $(\mathbf{r})_w^{\ddagger}$, we have $c_m \rightarrow B_u^{\ddagger} \wedge B_{\mathcal{X}}^{\ddagger}$. Since $(c_0)^{\ddagger} \rightarrow c_m$, we have $(c_1)^{\ddagger}$. \square

Theorem 6. For any left-linear $\text{CHRat}(\mathcal{X})$ program P and initial query $q \in \mathcal{L}'$:

$$\downarrow \mathcal{O}_P^a[q] \subseteq \mathcal{L}_P[q]$$

Proof. Let $q = T_0 \wedge c_0 \in \mathcal{L}'$. Let $\langle \exists \mathbf{x}(T_1 \wedge c) \rangle \in \mathcal{O}_P^a[q]$. Then there exists I such that $(T_0, c_0, \emptyset) \xrightarrow{*} (T, c, I)$. Logical implication being transitive, lemma 6 ensures that $(P)^{\ddagger} \models_{\mathcal{X}} (T_0, c_0, \emptyset)^{\ddagger} \rightarrow (T, c, I)^{\ddagger}$. Therefore, by definition of $(T_0, c_0, \emptyset)^{\ddagger}$, $(P)^{\ddagger} \models_{\mathcal{X}} q \leftrightarrow (T, c, I)^{\ddagger}$. Therefore, by definition of $\mathcal{L}_P[q]$, $(T, c, I)^{\ddagger} \in \mathcal{L}_P[q]$. Since $\mathcal{L}_P[q] = \downarrow \mathcal{L}_P[q]$, we have $\downarrow \mathcal{O}_P^a[q] \subseteq \mathcal{L}_P[(q)^{\ddagger}]$. \square

Linear logic semantics use similar atomic propositions to encode control tokens and suspended instances. A multiset M of constraints is interpreted by a linear tensor of the constraints $M^{\ddagger\ddagger} = \langle \bigotimes_{c \in M} c^* \rangle$. A rule is interpreted in linear logic as:

$$\begin{aligned} (H \setminus H' \Rightarrow G \mid B)^{\ddagger\ddagger} &= (H \setminus H' \Rightarrow G \mid B)_s^{\ddagger\ddagger} \otimes (H \setminus H' \Rightarrow G \mid B)_w^{\ddagger\ddagger} \\ (H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u)_s^{\ddagger\ddagger} &= \left\langle !\forall \mathbf{x} \left(\exists \mathbf{y}(G^{\ddagger\ddagger}) \otimes H^{\ddagger\ddagger} \otimes H'^{\ddagger\ddagger} \multimap \right. \right. \\ &\quad \left. \left. \exists K(H^{\ddagger\ddagger} \otimes H'^{\ddagger\ddagger} \otimes \text{id}(K, \mathbf{z}) \otimes \text{exists}(K, V_G) \otimes \right. \right. \\ &\quad \left. \left. \left(\bigotimes_{p(\mathbf{u}) \in G_u} \text{ask}_p(K, \mathbf{u}) \right) \right) \right\rangle \\ (H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u)_w^{\ddagger\ddagger} &= \left\langle !\forall \mathbf{x} \forall K \left(\exists \mathbf{y}(G^{\ddagger\ddagger}) \otimes H^{\ddagger\ddagger} \otimes H'^{\ddagger\ddagger} \otimes \right. \right. \\ &\quad \left. \left. \text{id}(K, \mathbf{z}) \otimes \left(\bigotimes_{p(\mathbf{u}) \in G_u} \text{entailed}_p(K, \mathbf{u}) \right) \multimap \right. \right. \\ &\quad \left. \left. \exists \mathbf{z}(H^{\ddagger\ddagger} G^{\ddagger\ddagger} \otimes B_{\mathcal{X}} \otimes B_u^{\ddagger\ddagger}) \right) \right\rangle \end{aligned}$$

and $\mathbf{x} = \text{fv}(H, H')$, $\mathbf{y} = \text{fv}(G) \setminus \text{fv}(H, H')$ and $\mathbf{z} = \text{fv}(G, B_{\mathcal{X}}, B_u) \setminus \text{fv}(H, H')$. The logical interpretation of a program $(P)^{\ddagger\ddagger}$ is the linear tensor of the interpretations of the rules of P .

Definition 49. The *linear logical semantics* of a query $q \in \mathcal{L}'$ is:

$$\mathcal{LL}_P[q] = \{c \in \mathcal{L}' \mid (P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} q \multimap c \otimes \top\}$$

Let V be the free variables of the initial query $q = T_0 \wedge c_0 \in \mathcal{L}'$. Configurations are interpreted as $(T, c, I)^{\ddagger\ddagger} = \langle \exists \mathbf{x}(T^{\ddagger\ddagger} \otimes c \otimes (\bigotimes_{K \in I} I(K)^*)) \rangle$ where $\mathbf{x} = \text{fv}(T, c) \setminus$

$\text{fv}(T_0, c_0)$: the query variables appear in the observables and are therefore left free in the interpretation. This definition is extended for a set of configurations S : $S^{\ddagger\ddagger} = \{(T, c, I)^{\ddagger\ddagger} \mid (T, c, I) \in S\}$.

CHRat operational semantics is sound and complete w.r.t. the linear logical semantics. This result follows from the following lemma:

Lemma 7. *If $c_0 \rightarrow c_1$, then $(P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} (c_0)^{\ddagger\ddagger} \multimap (c_1)^{\ddagger\ddagger}$.*

Proof. If $c_0 \xrightarrow{s} c_1$, then there is a rule $\mathbf{r} = \langle H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u \rangle$ in P renamed with fresh variables with respect to $c_0 = (T, c, I)$, and $c_1 = (\langle T \oplus \text{ask}(K, G_u) \oplus \text{exists}(K, V_G) \rangle, \langle c \wedge \exists(c_m) \rangle, \langle I \uplus (\mathbf{r}, \sigma_{\pi})_K \rangle)$. Then, according to $(\mathbf{r})_s^{\ddagger\ddagger}$, $(P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} ((T, c, I)^{\ddagger\ddagger} \multimap (c_1)^{\ddagger\ddagger})$.

If $c_0 \xrightarrow{w} c_1$, then there is a rule $\mathbf{r} = \langle H \setminus H' \Rightarrow G \mid B_{\mathcal{X}}, B_u \rangle$ in P renamed with fresh variables with respect to $c_0 = (T, c, I)$, and $c_1 = (\langle T \oplus \text{ask}(K, G_u) \oplus \text{exists}(K, V_G) \rangle, \langle c \wedge \exists(c_m) \rangle, \langle I \uplus (\mathbf{r}, \sigma_{\pi})_K \rangle)$. Then, according to $(\mathbf{r})_w^{\ddagger\ddagger}$, $(P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} ((T, c, I)^{\ddagger\ddagger} \multimap (c_1)^{\ddagger\ddagger})$. \square

Soundness and completeness of the operational semantics w.r.t. the linear logical semantics is summarized below:

Theorem 7. *For any CHRat(\mathcal{X}) program P and initial query $q \in \mathcal{C}'$:*

$$\downarrow \mathcal{O}_P^a[q] = \mathcal{LL}_P[q]$$

Proof. Let $q = T_0 \wedge c_0 \in \mathcal{L}'$.

Soundness Let $\langle \exists \mathbf{x}(T_1 \wedge c) \rangle \in \mathcal{O}_P^a[q]$. Then there exists I such that $(T_0, c_0, \emptyset) \xrightarrow{*} (T, c, I)$. Logical implication being transitive, lemma 7 ensures that $(P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} (T_0, c_0, \emptyset)^{\ddagger\ddagger} \multimap (T, c, I)^{\ddagger\ddagger}$. Therefore, by definition of $(T_0, c_0, \emptyset)^{\ddagger\ddagger}$, $(P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} q \multimap (T, c, I)^{\ddagger\ddagger}$. Therefore, by definition of $\mathcal{LL}_P[q]$, $(T, c, I)^{\ddagger\ddagger} \in \mathcal{LL}_P[q]$. Since $\mathcal{LL}_P[q] = \downarrow \mathcal{LL}_P[q]$, $\downarrow \mathcal{O}_P^a[q] \subseteq \mathcal{LL}_P[(q)^{\ddagger\ddagger}]$.

Completeness (*sketch*) Let $c \in \mathcal{L}_P[(q)^{\ddagger\ddagger}]$. Then $(P)^{\ddagger\ddagger} \models_{LL, \mathcal{X}} q \multimap c$. Main steps of the proof tree are *modus ponens* on linear implications introduced with $(\mathbf{r})_s^{\ddagger\ddagger}$ and $(\mathbf{r})_w^{\ddagger\ddagger}$ which mirror their operational counterpart. \square

Program Transformation from CHRat to CHR

Definition 50. Let $\llbracket \cdot \rrbracket : \text{CHRat}(\mathcal{X}) \rightarrow \text{CHR}(\mathcal{X})$ be the following morphism, where $G_u = \{t_1(\vec{v}_1), \dots, t_m(\vec{v}_m)\}$ and $V_G = \{y_1, \dots, y_p\}$ and $\text{fv}(H, H') = \{x_1, \dots, x_n\}$:

$$\llbracket H \setminus H' \Leftrightarrow G | B_{\mathcal{X}}, B_u \rrbracket = \begin{cases} H^*, H'^* \Rightarrow G_{\mathcal{X}} | \text{id}_i(K, x_1, \dots, x_n), \\ \text{exists}(K, y_1), \dots, \text{exists}(K, y_p), \\ \text{ask}_{t_1}(K, \vec{v}_1), \dots, \text{ask}_{t_m}(K, \vec{v}_m). \\ H^* \setminus H'^*, \text{id}_i(K, x_1, \dots, x_n), \\ \text{entailed}_{t_1}(K, \vec{v}_1), \dots, \text{entailed}_{t_m}(K, \vec{v}_m) \\ \Rightarrow G_{\mathcal{X}} | B_{\mathcal{X}}, B_u^*. \end{cases}$$

User constraints ask_{t_j} , exists , tokentailed_{t_j} and id_i corresponds to atomic propositions introduced in the logical semantics. i is a unique identifier associated to the rule. The linear logical semantics makes then link between the operational semantics and the transformation defined above: this leads to the following result which proves the soundness of the transformation in the logical semantics:

Theorem 8. For all $\text{CHRat}(\mathcal{X})$ program P and initial query $T_0 \wedge c_0 \in \mathcal{L}'$:

$$\mathcal{LL}_P[T_0 \wedge c_0] = \mathcal{LL}_{\llbracket P \rrbracket}(T_0, C_0)$$

Proof. Let $H \setminus H' \Leftrightarrow G | B$ be a CHRat rule. $(\llbracket H \setminus H' \Leftrightarrow G | B_{\mathcal{X}}, B_u \rrbracket)^{\dagger\dagger} = S \otimes W$ where:

$$\begin{aligned} S &= (H^*, H'^* \Rightarrow G_{\mathcal{X}} | \text{id}(K, x_1, \dots, x_n), \text{exists}(K, y_1), \dots, \text{exists}(K, y_p), \\ &\quad \text{ask}_{t_1}(K, \vec{v}_1), \dots, \text{ask}_{t_m}(K, \vec{v}_m))^{\dagger\dagger} \\ W &= (H^* \setminus H'^*, \text{id}(K, x_1, \dots, x_n), \\ &\quad \text{entailed}_{t_1}(K, \vec{v}_1), \dots, \text{entailed}_{t_m}(K, \vec{v}_m) \Rightarrow G_{\mathcal{X}} | B_{\mathcal{X}}, B_u^*)^{\dagger\dagger} \end{aligned}$$

By definition, $S = \langle !\forall \mathbf{x} (\exists \mathbf{y} (G_{\mathcal{X}}^{\dagger\dagger}) \otimes (H^* \oplus H'^*)^{\dagger\dagger} \multimap \exists K ((H^* \oplus H'^*)^{\dagger\dagger} \otimes \text{id}(K, \mathbf{z}) \otimes \text{exists}(K, V_G) \otimes (\bigotimes_{i=1}^m \text{ask}_{t_i}(K, \vec{v}_i)))) \rangle$ where $\mathbf{x} = \text{fv}(H, H')$, $\mathbf{y} = \text{fv}(G_{\mathcal{X}})$, $\mathbf{z} = \{x_1, \dots, x_n\}$ and $V_G = \{y_1, \dots, y_p\}$. Since $(H^* \oplus H'^*)^{\dagger\dagger} = H^{\ddagger\dagger} \otimes H'^{\ddagger\dagger}$, $S \equiv (\llbracket H \setminus H' \Leftrightarrow G | B_{\mathcal{X}}, B_u \rrbracket_s)^{\ddagger\dagger}$. By definition, $W = \langle !\forall \mathbf{x} (\exists \mathbf{y} (G_{\mathcal{X}}^{\dagger\dagger}) \otimes H^{\ddagger} \otimes H'^{\ddagger} \otimes \text{id}(K, \mathbf{z}) \otimes (\bigotimes_{i=1}^m \text{entailed}_{t_i}(K, \vec{v}_i)) \multimap \exists \mathbf{u} (H^{\ddagger} \otimes B_{\mathcal{X}}^{\dagger\dagger} \otimes B_u^*)) \rangle$ where $\mathbf{u} = \text{fv}(G, B_{\mathcal{X}}, B_u) \setminus \text{fv}(H, H')$. Therefore $W \equiv (\llbracket H \setminus H' \Leftrightarrow G | B_{\mathcal{X}}, B_u \rrbracket_w)^{\ddagger\dagger}$. \square

4.3.5 Hierarchical Definition of a Rational Terms Constraint Solver

Union-Find Constraint Component

The classical union-find (or disjoint set union) algorithm [86] has been implemented in CHR in [76] with its best-known quasi-linear algorithmic complexity. This positive result is remarkable because declarative languages such as logic programming do not achieve this efficiency [32]. The union-find algorithm maintains a partition of a universe, such that each equivalence class has a *representative* element. Three operations define this data structure:

- `make(X)` adds the element X to the universe, initially in an equivalence class reduced to the singleton $\{X\}$.
- `find(X)` returns the representative of the equivalence class of X .
- `union(X,Y)` joins the equivalence classes of X and Y (possibly changing the representative).

As a constraint solver in classical logic, such an algorithm solves the equivalence constraint $A =^{\sim} B$. The union and find constraint tokens reflect the imperative interpretation of the disjoint-set data structure, which makes representative elements explicit. However, being a representative element is a non-monotonic property which cannot be captured by constraints in classical logic.

Naive Implementation in CHRat.

A first implementation relies on the classical representation of equivalence classes by rooted trees [76]. Roots are representative elements, they are marked as such with the CHR-constraint `root(X)`. Tree branches are marked with $A \sim > B$, where A is the child and B the parent node.

```
component naive_union_find_solver.  
export make/1, =~/2.  
make(A) ⇔ root(A).  
union(A, B) ⇔ find(A, X), find(B, Y), link(X, Y).  
A ~> B \ find(A, X) ⇔ find(B, X).  
root(A) \ find(A, X) ⇔ X = A.  
link(A, A) ⇔ true.  
link(A, B), root(A), root(B) ⇔  
  B ~> A, root(A).
```

This implementation supposes that the entry-points make and union are used with constant arguments only, and that the first argument of find is always a constant. Telling this constraint yields to the union of both equivalence classes:

$$A =^{\sim} B \Rightarrow \text{union}(A, B).$$

A way to provide a naive implementation for the entailment solver of $=^{\sim}$ is to follow the branches until possibly finding a common ancestor for A and B.

$$\begin{aligned} \text{ask}(K, A =^{\sim} A) &\Leftrightarrow \text{entailed}(K, A =^{\sim} A). \\ A \rightsquigarrow C \setminus \text{ask}(K, A =^{\sim} B) &\Leftrightarrow C =^{\sim} B \mid \text{entailed}(K, A =^{\sim} B). \\ B \rightsquigarrow C \setminus \text{ask}(K, A =^{\sim} B) &\Leftrightarrow A =^{\sim} C \mid \text{entailed}(K, A =^{\sim} B). \end{aligned}$$

The computation required to check the constraint entailment is done by using recursively guards $C =^{\sim} B$ and $A =^{\sim} C$. Let S be a set of ground elements of \mathcal{X} .

Definition 51. A CHR store c has a valid union-find data-structure if $=^{\sim}/2$ and $\text{root}/1$ form a forest with elements in S .

Proposition 6. The property of having a valid union-find data-structure is preserved by all the rules of the solver.

Proposition 7. For every CHR store c with a valid union-find data-structure and for all elements A and B , A and B share the same equivalence class if and only if, for any fresh variable K , $\text{entailed}(K, A =^{\sim} B) \in \mathcal{O}_a(\text{ask}(K, A =^{\sim} B) \wedge c)$.

Proof. Suppose that A and B share the same equivalence class and R be the root of the tree representing this class. Let d_1 and d_2 be the length of the \rightsquigarrow -path between A and R , and between B and R respectively. Then, the property can be proved by recurrence over $d_1 + d_2$. If $d_1 + d_2 = 0$ then $A = B = R$ and askEq concludes. If $d_1 + d_2 \geq 1$ then at least one head of the rules askLeft and askRight matches the store. Suppose without loss of generality that askLeft is applicable. Then there is a \rightsquigarrow -path from C to R of length $d_1 - 1$. Therefore the hypothesis induction is applicable and the entailment test of $C =^{\sim} B$ succeeds. Then askLeft is fired and $\text{entailed}(K)$ is reachable.

Conversely, we first remark that if askEq is never fired, $\text{entailed}(K)$ is not reachable. Then, if A and B are not in the same equivalence class, there is no common element between \rightsquigarrow -paths starting from A and \rightsquigarrow -paths starting from B . By induction on the length of the derivation, every token $\text{ask}(K, C =^{\sim} D) \in \mathcal{O}_a(\text{ask}(K, A =^{\sim} B) \wedge c)$ is such that there is a path between A and C and a path between B and D . Therefore, the rule askEq is never fired. \square

Optimized Implementation in CHRat.

The second implementation proposed in [76] implements both *path-compression* and *union-by-rank* optimizations to reach the quasi-linear complexity in $O(n\alpha(n))$.

```

component union_find.
export make/1, =~/2.
make(A) ⇔ root(A, 0).
union(A, B) ⇔ find(A, X), find(B, Y), link(X, Y).
A ~> B, find(A, X) ⇔ find(B, X), A ~> X.
root(A, _) \ find(A, X) ⇔ X = A.
link(A, A) ⇔ true.
link(A, B), root(A, N), root(B, M) ⇔ N >= M |
  B ~> A, N1 is max(M+1, N), root(A, N1).
link(B, A), root(A, N), root(B, M) ⇔ N >= M |
  B ~> A, N1 is max(M+1, N), root(A, N1).

```

An optimized check for common equivalence class can rely on find to efficiently get the representatives and then compare them. `check(K, A, B, X, Y)` represents the knowledge that the equivalence class representatives of A and B are the roots X and Y respectively. When X and Y are known to be equal, `entailed(K)` is put to the store (`checkEq`).

```

ask(K, A =~ B) ⇔
  find(A, X), find(B, Y),
  check(K, A, B, X, Y).
root(X) \ check(K, A, B, X, X) ⇔ entailed(K).

```

These two rules are not sufficient to define a complete entailment checker because of the changes applied to the tree structure. Indeed, roots found for A and B can be invalidated by subsequent calls to `union`, which may transform these roots into child nodes. When a former root becomes a child node, the following two rules put `find` once again to get the new root.

```

X ~> C \ check(K, A, B, X, Y) ⇔
  find(A, Z), check(K, A, B, Z, Y).
Y ~> C \ check(K, A, B, X, Y) ⇔
  find(B, Z), check(K, A, B, X, Z).

```

Rational Tree Equality Constraint Component

Let us now consider rational terms, i.e. rooted, ordered, unranked, labeled, possibly infinite trees, with a finite number of structurally distinct sub-trees [18]. The nodes are supposed to belong to the universe considered by the union-find solver. Two nodes belonging to the same equivalence class are supposed to be structurally equal. Each node X has a signature F/N, where F is the label of X and N its arity: the associated constraint is denoted `fun(X, F, N)`. For each I between 1 and N, the constraint `arg(X, I, Y)` states that the Ith sub-tree of X is (structurally equal

to) Y. These constraints have just to be compatible between elements of the same equivalence class:

```

component rational_tree_solver.
import =~/2 from union_find_solver.
export fun/3, arg/3, =~/2.
fun(X0, F0, N0) \ fun(X1, F1, N1) ⇔ X0 =~ X1 |
    F0 = F1, N0 = N1.
arg(X0, N, Y0) \ arg(X1, N, Y1) ⇔ X0 =~ X1 | Y0 =~ Y1.

```

Telling that two trees are structurally equal, denoted $X \approx \approx Y$, can be reduced to the union of the two equivalence classes.

$$X \approx \approx Y \Leftrightarrow X \approx Y.$$

The computation associated to asking $A \approx \approx B$ requires a co-inductive derivation of structural comparisons to break infinite loops. This is done by memoization: tokens checking(K, A, B) signal that A can be assumed to be equal to B since this check is already in progress. checkTreeAux checks that signatures of A and B are equal and compares arguments.

```

ask(K, A ≈≈ B) ⇔ checkTree(K, A, B).
checkTree(K, A, B) ⇔ eqTree(K, A, B) | entailed(K, A ≈≈ B).
ask(eqTree(K, A, B)) ⇔
    checking(K, A, B), fun(A, FA, NA), fun(B, FB, NB),
    checkTreeAux(K, A, B, FA, NA, FB, NB).
checkTreeAux(K, A, B, F, N, F, N) ⇔
    askArgs(K, A, B, 1, N), collectArgs(K, A, B, 1, N).

```

askArgs adds every askArg token corresponding to each pair of point-wise subtrees of A and B. askArg answers entailedArg if they match. collectArg ensures every entailedArg token have been put before concluding about the entailment of eqTree(K, A, B). It is very close to the definition of an ask solver, but the guard deals with a variable number of tokens equals to the arity of A and B.

```

askArgs(K, A, B, I, N) ⇔ I ≤ N |
    arg(A, I, AI), arg(B, I, BI),
    askArg(K, A, B, I, AI, BI),
    J is I + 1, askArgs(K, A, B, J, N).
askArgs(K, A, B, I, N) ⇔ true.
collectArgs(K, A, B, I, N), entailedArg(K, A, B, I) ⇔
    J is I + 1, collectArgs(K, A, B, J, N).
collectArgs(K, A, B, I, N) ⇔ I > N |
    entailed(eqTree(K, A, B)).

```

askArg firstly checks if the equality is memoized, otherwise asks for its check.

$$\begin{aligned} \text{checking}(K, AI, BI) \setminus \text{askArg}(K, A, B, I, AI, BI) &\Leftrightarrow \\ &\text{entailedArg}(K, A, B, I). \\ \text{askArg}(K, A, B, I, AI, BI) &\Leftrightarrow \text{eqTree}(K, AI, BI) \mid \\ &\text{entailedArg}(K, A, B, I). \end{aligned}$$

For the sake of simplicity, this program does not perform garbage collection. In particular memoized tokens $\text{checking}(K,A,B)$ are never removed from the store, and disentailed constraints are not eliminated.

4.4 The SiLCC Programming Language

A SiLCC program defines an *LCC agent*. The following agent writes “Hello world!” to standard output.

```
io.std.writeln("Hello world!", _)
```

LCC agents are written with the following grammar. All the other constructions (except the foreign function interface) are just convenient notations for usual programming idioms and can be rewritten with these fundamental constructions.

Terms are constructed with variables and domain symbols: the domain includes integers, floating-point values, character strings (double-quoted), Herbrand terms (with simple quoted functors), and other structures that we will detail later on.

```
exists m (
  m.a('f'(1, 2.5, "example"))
  forall x (m.a(x) -> exists k (io.std.writeln(x, k)))
)
```

Variables are bound with the quantifiers `forall` and `exists` with the obvious scope. Other variables are free and denote values defined in other files.

Remark 2. When an existentially quantified variable is only needed once, we can simply note it “`_`”. The previous example is therefore equivalent to the code below.

```
exists m (
  m.a('f'(1, 2.5, "example"))
  forall x (m.a(x) -> io.std.writeln(x, _))
)
```

A theoretical description of the semantics for these fundamental constructions is given in [41]. Intuitively, `exists` adds tokens to the store. `forall` asks that some tokens are simultaneously present in the store for removing them and executing the `ask` body. Transient asks are fired only once, persistent asks are fired as many times as there are matching tokens to remove. The term which precedes the dot before a token identifier is called the module of the token.

4.4.1 Procedures

Since they can be fired several times, persistent asks are typically useful to define procedures, as in the following example (note that string concatenation is just obtained by putting strings one after the others).

```
exists m (
  forall s (
    m.hello(s)
  =>
    io.std.writeln("Hello " s "!", _)
  )
  m.hello("You")
  m.hello("Me")
)
```

Asks without universally quantified variables are written by omitting the `forall` keyword.

```
exists m (
  (
    m.hello()
  =>
    io.std.writeln("Hello word!", _)
  )
  m.hello()
)
```

There is a convenient notation for persistent asks on single head when all the variables appearing in the arguments of the head are universally quantified (typically, procedure definitions): just write the head followed by the body of the agent surrounded by curly brackets.

```
exists m (
  m.hello(s) {
    io.std.writeln("Hello " s "!", _)
  }
  m.hello("You")
  m.hello("Me")
)
```

Remark 3. When a variable is existentially quantified up to the next closing delimiter (if any, the end of the file otherwise), the quantification may be followed just by a comma and parentheses omitted. The previous example is therefore equivalent to the code below.

```

exists m,
m.hello(s) {
  io.std.writeln("Hello " s "!", _)
}

```

```

m.hello("You")
m.hello("Me")

```

Consequently, successive introduction of variables correspond to nested existential quantification and reintroducing a variable masks its previous definition.

```

exists m,
m.hello(s) {
  io.std.writeln("Hello " s "!", _)
}

```

```

m.hello("You")

```

```

(
  exists m,
  m.hello(s) {
    io.std.writeln("Hello " s ".", _)
  }
  m.hello("Me")
)

```

```

m.hello("everybody")

```

```

exists m,
m.hello("(will not appear)")

```

4.4.2 Sequencing

When agents are put in parallel, as `hello("You")` and `hello("Me")` in the previous example, there is no guarantee that the first is executed before or after the second or, worse, that the two executions are not interleaved. To sequence printings, we should use a variable, say `k`, passed to the continuation argument of `writeln ("...", k)`: `writeln` tells `k.done()` when the printing is finished, allowing us to consume it before pursuing with other printings.

```

exists k,
io.std.writeln("One", k)

```

```

(
  k.done()
->
  exists k',
  io.std.writeln("Two", k')
  (
    k'.done()
  ->
    io.std.writeln("Three", _)
  )
)

```

The notation `do { ... }` allows sequences to be written more conveniently. The previous example is equivalent to the code below.

```

do {
  io.std.writeln("One")
  io.std.writeln("Two")
  io.std.writeln("Three")
}

```

Elements in sequences may use tokens of the form `value(\emph{x})` instead of `done()` to return values. These values are bound with `<-` notation.

```

exists m,
forall k (
  m.a(k)
=>
  k.value(1)
)
do {
  x &lt;- m.a()
  io.std.writeln(x)
}

```

Remark 4. This example is expanded to the following code.

```

exists m,
forall k (
  m.a(k)
=>
  k.value(1)
)
exists k,

```

```
m.a(k)
forall x (k.value(x) -> io.std.writeln(x, _))
```

4.4.3 Arithmetic and functional notation

We can do some arithmetic by using the standard library module `data.arith`. The following code prints the product of 123 by 45. Note that `io.std.writeln` waits that `v` gets a value before printing it.

```
exists v,
data.arith.mul(123, 45, v)
io.std.writeln(v, _)
```

The functional notations allows such code to be written more elegantly without having to give a name to the intermediary variable: when a token is written in term position, the token is moved and put in parallel with the rest of the agent, a fresh variable is added as its last argument, and the term where the token was is replaced by this variable. The previous example is equivalent to the code below.

```
io.std.writeln(data.arith.mul(123, 45), _)
```

Moreover, we can write `x * y` instead of `data.arith.mul(x, y)`, and, more generally, the operators `+`, `-`, `*`, `/`, `^`, `mod` are available as usual notations for arithmetic.

The function definition notation defines procedures with an implicit last argument conventionally called `result`.

```
exists m,
function m.square(x) {
  result = x * x
}
io.std.writeln(m.square(12), _)
```

When the body of the function is reduced to assign a value to `result`, we may use the alternative notation `function \emph{f}(...) = \emph{value}`.

```
exists m,
function m.square(x) = x * x
io.std.writeln(m.square(12), _)
```

4.4.4 Conditional

`data.arith` defines arithmetic comparisons as well. The result of a comparison is represented with a module for either a token `true()` or a token `false()`.

```

exists m,
m.check(22)
m.check(n) {
  exists k,
  data.arith.le(n, 10, k)
  (
    k.true()
  =>
    io.std.writeln(n.string() " is a number less than or equal to 10.", _)
  )
  (
    k.false()
  =>
    io.std.writeln(n.string() " is a number greater than 10.", _)
  )
}

```

Note that all domain values can be converted to character strings by telling a `string()` token in their module.

The idiom of the pair of asks on `true()` and `false()` on the same module can be written more usually as follows. (Note that $x \leq y$ is a notation for `data.arith.le(x, y)`, usual notations for other comparisons are available as well.)

```

exists m,
m.check(22)

m.check(n) {
  if n &lt;= 10 do {
    io.std.writeln(n.string() " is a number less than or equal to 10.")
  }
  else do {
    io.std.writeln(n.string() " is a number greater than 10.")
  }
}

```

Comparisons are allowed in asks as well.

```

exists m,
m.check(22)

forall n (m.check(n) n &lt;= 10 =>
  io.std.writeln(n.string() " is a number less than or equal to 10.", _))

```

```
forall n (m.check(n) n > 10 =>
  io.std.writeln(n.string() " is a number greater than 10.", _))
```

4.4.5 Angelism

Executing an agent can be non-deterministic. Typically, suppose that two asks wait for the same token: when this token is told, there are two possible computations whether the first or the second ask is fired. SiLCC has been designed such that all observable computations are effectively observed. Hence, if each of these two asks leads to an observable behavior (so as to say a side-effect, like printing something), both behaviors will be observed.

```
exists m,
m.token()

m.token() {
  io.std.writeln("First branch.", _)
}

m.token() {
  io.std.writeln("Second branch.", _)
}
```

However, these two branches are two distinct paths in the non-deterministic computation produced by the agent. In particular, they cannot share resources.

```
exists m,
m.token()

m.token() {
  io.std.writeln("First branch.", _)
  m.resource1()
}

m.token() {
  io.std.writeln("Second branch.", _)
  m.resource2()
}

(m.resource1() m.resource2() -> io.std.writeln("Will not be executed!", _))
```

In the previous example, the third ask can only be fired if there exists a computation path where both `resource1` and `resource2` are simultaneously available,

yet each of these resources is produced by a distinct computation path.

The essential consequence of angelism is the decomposability of asks. For example, $\text{forall } x (a(x) \wedge a(x-1) \wedge a(x+2) \Rightarrow \dots)_+$ is equivalent to $\text{forall } x (a(x) \Rightarrow a(x-1) \rightarrow a(x+2) \rightarrow \dots)_+$: since the second agent suspends if one of the three tokens to be consumed is not available, and since there is no side-effects between the three nested asks, the possible partial consumption is just not observed. More generally, every program is rewritten to a kernel language which is indeed limited to asks on single tokens where all arguments are universally quantified. All other forms of asks are decomposed: for example, $\text{forall } (a(x, x-1) \rightarrow \dots)_+$ is decomposed in $\text{forall } x y (a(x, y) \rightarrow \text{if } y = x-1 \dots)_+$.

4.4.6 Modules and the package system

There is always a “current” module, which tokens not prefixed with a dot belong to. This module can always be referred by the keyword `this`. The notation with `with` locally changes the current module.

```
with io.std do {
  write("10 + 2 = ")
  write(10 + 2)
  writeln("")
}
```

The notation `module { ... }` is an expression returning a fresh module. This module is the current module of the agent written between the curly brackets.

```
exists m,
m.operate(a, op, b, f) {
  with io.std do {
    write(a.string() " " op " " b.string() " = ")
    write(f.get(a, b))
    writeln("")
  }
}
m.operate(12, "*", 5, module { function get(a, b) = a * b })
```

The agent defined in the file `\emph{x}.lcc` is read with the free variable `\emph{x}` as current module. When the free variable `\emph{x}` is used, the file `\emph{x}.lcc` is loaded.

If the file `example.lcc` contains:

```
test() {
  io.std.writeln("This is an example.", _)
}
```

the procedure can be called by telling `example.test()`.

Top-level phrases are read with `silcc.top` as current module. This module reacts to some directives like `trace()` for debugging purpose.

Files are sought in the current directory and then in the directories enumerated in the `SILCC_PATH` environment variable, where directories are separated by colons. This variable should typically contains the path to the standard library. If this variable is not defined, the standard library is heuristically sought in the path `../lib` relatively to the location of the “`silcc`” executable.

The free variable `\emph{x}.\emph{y}.\emph{z}` designates the file `\emph{x}/\emph{y}/\emph{z}`. To be called this way, the file should begin by `package \emph{x}.\emph{y}` (in order to ensure that the name of the variable is unique).

4.4.7 Recursion

Here is the classical recursive algorithm for solving the Towers of Hanoi.

```
exists m,

m.Hanoi(n, A, B, C, k) {
  k = do {
    if n = 0 do {
      io.std.writeln("I have nothing to do.")
    }
    else if n = 1 do {
      io.std.writeln("I move a disk from " A " to " C)
    }
    else do {
      m.Hanoi(n - 1, A, C, B)
      io.std.writeln("I move a disk from " A " to " C)
      m.Hanoi(n - 1, B, A, C)
    }
  }
}

m.Hanoi(4, "A", "B", "C", _)
```

When the `do { ... }` is used in term position, it returns the last continuation if the sequence is not empty. If the sequence is empty, it tells the token `done()` in a fresh module variable and returns it.

4.4.8 Lists and iterations

Lists of values can be constructed with the notation `[\emph{v1}, ..., \emph{vn}]`. Lists are constructed recursively from the value `[]` (which is a notation for `data.list.nil()`) and `head :: tail` (which is a notation for `data.list.cons(head, tail)`) where `head` is the first element and `tail` the list of the other elements. The length of a list is measurable with `data.list.length`. List concatenation can be computed with `data.list.concat(l0, l1)` or just by putting lists one after the other like in `(l0 l1)` (which is just a notation for `data.term.concat(l0, l1)`, defined for lists and strings).

Iterations can be written with the `for` notation. In the following example, we use it on lists.

```
do {
  for i in ["one", "two", "three"] do {
    io.std.writeln(i)
  }
}
```

Remark 5. The previous code is just a notation for the following.

```
do {
  ["one", "two", "three"].iter(
    module {
      forall i v (do(i, v) => io.std.writeln(i, v))
    }
  )
}
```

Hence, all modules which implement asks on the `iter` token can be iterated with `for`. There is also a version of the `for` notation without continuation. For iterating on integers, one can use the notation `1 .. 10` (which is rewritten as `data.int.range(1, 10)` which returns a module implementing `iter`).

```
do {
  for _ in 1 .. 10 do {
    io.std.writeln("This will be written ten times.")
  }
}
```

4.4.9 Pattern-matching and terms

Lists can be destructed by pattern-matching: pattern-matching generalizes the `if/else` by allowing the `if` branch to introduce existential variables. (Note that

let $x = v$ is just a notation for $\text{exists } x', x' = v \text{ exists } x, x = x'$ with the introduction of the fresh variable x' to avoid capturing x in v).

```
let l = ["a", "b", "c"]
if exists hd tl, l = hd :: tl do {
  io.std.writeln("Head: " hd)
  io.std.writeln("Tail: " tl.string ())
}
```

Compound terms are constructed by quoting the functor, possibly followed by arguments between parentheses. Functors are mapped to integers, in particular single-character functors are mapped to Unicode code-point. (Note that pattern-matching can be decomposed in several conjunctions to keep intermediate sub-terms.)

```
let t = 'a compound term'(1, ['f'('a ')])
if exists t' c, t = 'a compound term'(_, [t']) and t' = 'f'(c) do {
  io.std.writeln(t.functor ())
  io.std.writeln(t'.functor ())
  io.std.write_char(c)
  io.std.write_char('\n')
}
```

Telling functor to a compound term returns its functor as a character string, telling arity returns its arity. $t[n]$ returns the n th argument (in general, $e[n]$ is just a notation for $e.get(n)$).

Indeed, let is just a particular case of pattern-matching where all free variables are implicitly quantified.

```
let t = 'f'("Hello")
let 'f'(x) = t
do { io.std.writeln(x) }
```

for does pattern-matching as well to filter between enumerated items (note that iterating over a compound term means iterating over its arguments).

```
do {
  for [x] in 'f'(["a"], ["b", "c"], ["d"]) do {
    io.std.writeln(x)
  }
}
```

Remark 6. The previous code is just a notation for the following.

```
do {
  'f'(["a"], ["b", "c"], ["d"]).iter (
```

```

module {
  forall i v (do(i, v) =>
    v = (
      if exists x, i = [x] do { io.std.writeln(x) }
      else do {}
    )
  )
}

```

The `<-` notation for do sequences allows pattern-matching as well. Note the notation for tuples: in the following, the token value has only one argument (a pair). Tuples are just a notation for terms with the functor `''`.

```

exists m,
forall k (
  m.a(k)
=>
  k.value((2, 5))
)
do {
  (x, y) &lt;- m.a()
  io.std.writeln(x + y)
}

```

Remark 7. This example is expanded to the following code.

```

exists m,
forall k (
  m.a(k)
=>
  k.value((2, 5))
)
exists k,
m.a(k)
forall x y (k.value((x, y)) -> io.std.writeln(x + y, _))

```

4.4.10 References

New mutable references may be defined with `data.ref.new(\emph{initial value})`. They act as a value but can be modified with `do { \emph{x} := \emph{new value} }` (`x := e` is just a notation for `x.set(e)`).

```

let x = data.ref.new(1)

do {
  io.std.writeln(x * 2)
  x := x + 1
  io.std.writeln(x * 2)
}

```

4.4.11 Records

Records are denoted $\{ \text{\emph{field}}:\text{\emph{value}}; \dots; \text{\emph{field}}:\text{\emph{value}} \}$. They are modules where each field is a function returning its value, plus support for enumerating fields and values. We can access to the field f in x with the expression $x:f$ which is a notation for $x.f()$ (this notation is not specific to records). In pattern-matching, all fields must match and there should not be other fields except if \dots is used.

```

let x = { a: "hello "; b: 42 }

if exists a, x = { a: a ... } do {
  io.std.writeln(a)
}

```

Remark 8. If the value of a field is omitted, *e.g.* “ $\{ f \}$ ”, then the field takes the value of the variable with the same name: “ $\{ f: f \}$ ”. The previous example can then be rewritten as follows.

```

let x = { a: "hello "; b: 42 }

if exists a, x = { a ... } do {
  io.std.writeln(a)
}

```

4.4.12 Hash-tables and argument indexing

The expression `data.hashtbl.new(\emph{initial size})` (defined in the standard library of the on-going 0.0.1 version) returns a fresh hash-table which associates a fresh variable to each domain value. The expression `\emph{tbl}[\emph{key}]` returns the value associated to *key*.

```

let t = data.hashtbl.new(17)
t['f'({ a: 1})] = "example"

```

```

t[{u: 2; v: 3}] = "test"
with io.std do {
  writeln(t['f'({ a: 1}]))
  writeln(t[{u: 2; v: 3}]))
}

```

Hash-tables are useful to implement indexing on token arguments. Since there is only native indexation on token modules, filtering on arguments in asks can be inefficient. This is a consequence of kernel design: the ask

```
forall i (m.a(i) m.b(i) -> ...)
```

is rewritten in

```
forall i j (m.a(i) m.b(j) -> if i = j { ... })
```

and is therefore fired for every pair of tokens $a(i)$ and $b(j)$. However, a hash-table can be introduced to index the tokens $m.b(i)$ on their argument i : each token $m.b(i)$ is consumed to produce a token $m'.b()$ where m' is indexed on i .

```

exists m,
let b-tbl = data.hashtbl.new(17)
forall i (m.b(i) => b-tbl[i].b())

```

The ask on $m.a(i)$ and $m.b(i)$ can now be written with indexing.

```

forall i (m.a(i) b-tbl[i].b() -> io.std.writeln(i, _))
m.a("Test") m.b("Test")

```

It is worth noticing that since $b-tbl[i].b()$ depends on $b(i)$, consuming one or the other is equivalent. Therefore, angelism ensures that user-indexing on $b-tbl[i].b()$ correctly cohabitates with other asks on $b(i)$.

4.4.13 Asks on variable number of tokens

Let say we want to write an agent with an ask equivalent to, by abuse of notation, $forall\ i\ (m.a(i)\ m.a(i\ 1)\ \dots\ m.a(2 * i)\ =>\ \dots)+$. The basic idea is to iteratively construct chains $m'.chain(i, j)$ in a local module m' by consuming the tokens $m.a(i), \dots, m.a(j)$. Since these chains only exist in computation paths where the corresponding $a()$ tokens have been consumed, consuming a chain is equivalent to consume the individual tokens.

```

exists m,
exists m' (
  forall i (m.a(i) => m'.chain(i, i))
  forall i j (m'.chain(i, j) m.a(j + 1) => m'.chain(i, j + 1))
)

```

```

    forall i (m'.chain(i, 2 * i) => io.std.writeln(i, _))
  )

m.a(7) m.a(8) m.a(9) m.a(10) m.a(11) m.a(12) m.a(13) m.a(14)

```

However, the execution is noticeably slow. The `m.a(i)` tokens should be indexed.

```

exists m,
exists m' (
  let a-tbl = data.hashtbl.new(17)
  forall i (m.a(i) => a-tbl[i].a())

  forall i (m.a(i) => m'.chain(i, i))
  forall i j (m'.chain(i, j) a-tbl[j + 1].a() => m'.chain(i, j + 1))
  forall i (m'.chain(i, 2 * i) => io.std.writeln(i, _))
)

m.a(7) m.a(8) m.a(9) m.a(10) m.a(11) m.a(12) m.a(13) m.a(14)

```

That's better but it is still slow. The number of chains is quadratic compared to the number of `m.a()`. Since angelism gives us the possibility to program our own checking procedure, we are now free to optimise it and change the complexity. All chains were constructed whereas the only interesting chains are starting from `m.a(i)` such that `m.a(2 * i)` is present. Therefore, the computation of chains could be driven by the presence of the two enclosing tokens `m.a(i)` and `m.a(2 * i)`.

```

exists m,

let a-tbl = data.hashtbl.new(17)
forall i (m.a(i) => a-tbl[i].a())

forall i (
  m.a(i) a-tbl[2 * i].a()
=>
  exists m',
  m'.chain(i)
  forall j (
    m'.chain(j) a-tbl[j + 1].a()
=>
    if j + 1 = 2 * i - 1 {
      io.std.writeln(i, _)
    }
  )
)

```



```

    }
    else {
      m'.chain(j + 1)
    }
  )
)

```

m.a(7) m.a(8) m.a(9) m.a(10) m.a(11) m.a(12) m.a(13) m.a(14)

4.5 Related work

The observable of the set of all consequences from a given agent was introduced since the beginning of concurrent constraint programming [73]. However, angelic semantics were never been considered for implementation, preferring committed-choice semantics despite their lack of logical interpretation.

The sequencing construction is a generalization of the sequencing operator introduced in Sylvain Soliman's PhD thesis. It can be seen as a form of continuation passing style.

Chapter 5

Conclusion

The programming paradigms that have been considered throughout this thesis share the “programs as models” approach: LLCC and CHR have logical semantics, constraint models have relational interpretations. The observables of interest involve the resolution of a search problem: the execution involves the search for a proof, the search for logical consequences, or the search for solutions for a constraint satisfaction problem.

Execution models are the algorithmic counter-part of these observables of interest: an execution model describes a search strategy for exploring the space of proofs, the space of consequences or the space of assignments. It is worth noticing that each of the three transformations have been guided by lifting programming into modeling: the committed-choice semantics for LLCC modeled in CHR, search strategies modeled as constraint satisfaction problems, and extra-logical control operators modeled as blocked paths in the angelic execution of concurrent programs, faithful to the logical interpretation of agents.

These new high-level compilation schemes open new insights for compiler implementation: constraint solvers could be implemented focusing on the default fixed enumeration strategy and efficient propagators for reified constraints, and there are new challenges for efficient angelic execution for concurrent languages, with the simplest form for asks. From the logician point of view, there is still a connection to be made between angelic execution and the compilation scheme that transforms Horn clauses to search trees by exploring all successful execution paths. The memory shared between execution paths, which allows programs to go “beyond search trees” in ClpZinc and that extends control in SiLCC, opens the hope for deeper logic formalisms that would be able to capture these behaviors logically.

5.1 Committed-Choice Semantics for LLCC through CHR

We have defined compositional translations from CHR to LLCC and from labeled LLCC to CHR and proved that semantics are preserved with strong bisimilarity. Both CHR and LLCC languages are based on the same model of concurrent computation, where agents communicate through a shared constraint store, with a synchronization mechanism based on constraint entailment. This is a generalization of the previous links between CHR and linear logic. As the work for modules in LLCC suggests[41], variables and CHR constraints are expressive enough to embed a form of closures, and thus lead to a simple encoding for the λ -calculus.

Whereas the state during a CHR derivation is entirely determined by the contents of constraint stores, an LLCC configuration contains suspended agents as well. The *ask-lifting* transformation reveals that suspensions can be reified as linear tokens, which in turns become CHR constraints: tokens acting as transient asks are consumed whereas tokens acting as persistent asks are propagated.

Behaviors of programs or agents obtained by translation are precisely related to their antecedents by (strong) bisimulation. To our knowledge, only weak bisimulation results [20] were formulated in the literature for CHR before. To achieve strong bisimulation in our case, we have managed to circumscribe collaterally the non-determinism in the naive operational semantics of CHR and in the operational semantics of LLCC. Transition systems considered here are non-labeled: this was sufficient for semantics preservation and there are good intuitions about the pair of involved firing rules at each step. Formalizing these intuitions by labeling with rule names seems feasible but with low interest. However, labels usually serve to follow messages that an agent either sends or receives. A challenge would be to label \Rightarrow -transitions by constraints whereas each single transition consumes some while adding others.

The closure encoding may suggest a new programming style, complementary to the imperative RAM-based style recently described [82]. Optimization of the CHR constraints which reify closures could be explored.

These transformations leads to a straightforward implementation of committed-choice LLCC. However, the moot point is to understand the relevance of CHR refined semantics for the translated LLCC agents: the question of control in LLCC has been tackled in the third part of this thesis, with angelic semantics.

5.2 Constraint Model and Search

We have shown that tree search procedures, such as for instance heuristic labeling, dichotomy, interval-splitting, limited discrepancy search, and dynamic symmetry

breaking during search, can be internalized in a constraint model through reified constraints. On the complex dynamic strategy used for solving Korf’s benchmark for square packing, we have shown that the implementation overhead is limited to a factor 2, and can be measured on different CSP solvers without particular support for search. We have also shown with an example that the propagation of search constraints can in fact exhibit an exponential speed up, compared to a classical procedural implementation of the search strategy.

It is worth noting that the conversion of search into constraints opens up a whole field of challenges for CSP solvers with limited built-in search strategies. For instance, Korf’s packing problem with the complex strategy of [80] has been proposed to the MiniZinc contest with the best known strategy. Therefore, the two sentences of [80] stating that this packing problem “nicely tests the generality of a search method” and is a “more attractive benchmark for placement problems than the perfect square” can now apply to compare a broad range of CSP solvers.

Furthermore, work on complex problems with dedicated heuristics can now become more independent of a particular solver through the modeling of the search strategy in our approach. We have added the needed indexicals to the FlatZinc parser of some solvers (Choco [22], JaCoP [50], SICStus [1], Gecode [16], or-tools [37]) and encourage all solver developers to do so in order to tackle these new challenging problems for the MiniZinc community.

Finally, as a perspective for future work, the reification of choice point constraints in our scheme is in principle compatible with lazy clause generation techniques [65] and the learning of nogood by using a SAT solver. Such a combination of modeling search by constraints and learning constraints during search is however quite intriguing and will be the matter of future work.

5.3 Angelic Programming

We have described a new execution model for LLCC and CHR, the angelic semantics, which computes all the reachable configuration from an initial LLCC agent or CHR goal.

We introduced a notion of derivation nets for graphical representation of execution paths. These derivation nets allow the description of sharing strategies which make the angelic semantics tractable in practice. In these settings, we illustrate how angelic semantics can result to a fully declarative language, where control is captured by the logical interpretation. Proposed applications give natural solutions in the angelic execution model to questions which are still open with committed-choice: the existence of CHR meta-interpreters, the redefinition of specific user constraint representations, for indexation in particular, and more generally the interleaving of arbitrary computation between head consumption,

allowing deep guards.

We have shown that by letting the programmer define in CHRat not only satisfiability checks but also entailment checks for constraints, this version of CHR becomes fully modular, i.e. constraints defined in one component can be reused in rules and guards in other components without any restriction. We have shown that some classical examples of constraint solvers defined in CHR could easily be modularized in CHRat and reused for building complex constraint solvers.

Interpreting operational semantics (indifferently CHR or LLCC) as a proof search method in linear logic reveals a parallel between the elimination of solving non-determinism and focalization theory [6] which remains to explore.

This work is a move forward to more declarativity, for reducing the gap between the logical interpretation and the effective implementation of the semantics. We hope for more theoretical development of algorithms taking benefits of the angelic semantics, as well as progress for implementation efficiency.

5.4 Language Implementations

Each of the three sections of this thesis directly leads to an implementation design.

- The ask-lifting transformation from LCC to CHR gives a compiler for LCC under the committed-choice semantics.
- The reification of search strategy gives a compiler for terminating constraint logic programs into constraint solving problems. The implementation, ClpZinc, targetting the MiniZinc middle-level modelling language for constraint satisfaction problems, allows CLP programs to be solved using any solver able to solve MiniZinc models. <http://lifeware.inria.fr/~tmartine/clp2zinc/>
- The collapse of the hierarchy of guard semantics in the angelic settings leads to a minimal and natural choice for a language kernel for LCC under the angelic semantics, keeping only atomic guards. The implementation, SiLCC, combines many layers of syntactic sugar for offering to the programmer high-level constructions that are compiled into the kernel language of agents with atomic guards. <http://lifeware.inria.fr/~tmartine/silcc/>

Bibliography

- [1] SICS AB. Sicstus 4.2.3, 2012.
- [2] Slim Abdennadher. Operational semantics and confluence of constraint propagation rules. In *Proceedings of CP'1997, 3rd International Conference on Principles and Practice of Constraint Programming*, volume 1330 of *Lecture Notes in Computer Science*, pages 252–266, Linz, 1997. Springer-Verlag.
- [3] Slim Abdennadher, Thom W. Frühwirth, and H. Meuss. Confluence and semantics of constraint simplification rules. *Constraints*, 4(2):133–165, 1999.
- [4] Slim Abdennadher and Michael Marte. University course timetabling using constraint handling rules. *Journal of Applied Artificial Intelligence*, 14, 2000.
- [5] Carsten Schürmann Anders Schack-Nielsen. Invited talk: The CHR-Celf connection. In Thom Frühwirth and Tom Schrijvers, editors, *Proceedings of the fifth Constraint Handling Rules Workshop CHR'08*, 2008.
- [6] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3), 1992.
- [7] Rolf Backofen and Sebastian Will. Excluding symmetries in constraint-based search. In Joxan Jaffar, editor, *CP*, volume 1713 of *Lecture Notes in Computer Science*, pages 73–87. Springer, 1999.
- [8] Gérard Berry and Gérard Boudol. The chemical abstract machine. *Theoretical Computer Science*, 96, 1992.
- [9] Eike Best, Frank S. de Boer, and Catuscia Palamidessi. Concurrent constraint programming with information removal. In *Proceedings of Coordination*, *Lecture Notes in Computer Science*. Springer-Verlag, 1997.
- [10] Hariolf Betz and Thom W. Frühwirth. A linear-logic semantics for constraint handling rules. In *Proceeding of CP 2005, 11th*, pages 137–151, 2005.

- [11] Hariolf Betz, Frank Raiser, and Thom Frühwirth. A complete and terminating execution model for Constraint Handling Rules. *Theory and Practice of Logic Programming*, 10(4-6):597–610, 2010.
- [12] K. H. Borgwardt. The average number of pivot steps required by the simplex-method is polynomial. *Mathematical Methods of Operations Research*, 1982.
- [13] M. Bruynooghe, M. Codish, J. P. Gallagher, S. Genaim, and W. Vanhoof. Termination analysis of logic programs through combination of type-based norms. *ACM Transactions on Programming Languages and Systems*, 29(2), 2007.
- [14] Philippe Codognet, Radhia Cousot, François Fages, Roberto Giacobazzi, Jean-Louis Imbert, Jean-François Pique, and Antoine Rauzy. Langages logiques concurrents avec contraintes. Technical report, Ecole Normale Supérieure, Rapport final de l'action inter-PRC Informatique MESR 93S0051, "Opérations de recherches coordonnées en calcul parallèle et distribuée", October 1995.
- [15] A. Colmerauer, Henry Kanoui, Robert Pasero, and Philippe Roussel. Un système de communication en français, rapport préliminaire de fin de contrat iria, groupe intelligence artificielle. Technical report, Technical Report, Faculté des Sciences de Luminy, Université Aix-Marseille II, 1972.
- [16] The Gecode Community. Gecode 4.2.1, 2013.
- [17] Emmanuel Coquery and François Fages. From typing constraints to typed constraint systems in CHR. In *Proceedings of Third workshop on Rule-based Constraint Reasoning and Programming, associated to CP'01*, November 2001.
- [18] B. Courcelle. Fundamental properties of infinite trees. *Theoretical Computer Science*, 25:95–169, 1983.
- [19] Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL'77: Proceedings of the 6th ACM Symposium on Principles of Programming Languages*, pages 238–252, New York, 1977. ACM Press. Los Angeles.
- [20] Leslie De Koninck. Logical algorithms meets CHR: A meta-complexity theorem for Constraint Handling Rules with rule priorities. *Theory and Practice of Logic Programming*, 2009.

- [21] Leslie De Koninck, Tom Schrijvers, and Bart Demoen. User-definable rule priorities for CHR. In *Proceedings of PPDP'07, International Conference on Principles and Practice of Declarative Programming, Wroclaw, Poland*, pages 25–36. ACM Press, 2007.
- [22] École des Mines de Nantes. Choco 3, 2014.
- [23] E.W. Dijkstra. Hierarchical ordering of sequential processes. *Acta Informatica*, 1:115–138, 1971.
- [24] Gregory Duck, Maria Garcia de la Banda, and Peter Stuckey. Compiling ask constraints. In *Proceedings of International Conference on Logic Programming ICLP 2004*, Lecture Notes in Computer Science. Springer-Verlag, 2004.
- [25] Gregory J. Duck, Peter J. Stuckey, Maria Garcia De La Banda, and Christian Holzbaaur. The refined operational semantics of constraint handling rules. In *In 20th International Conference on Logic Programming (ICLP'04)*, pages 90–104. Springer-Verlag, 2004.
- [26] Gregory J. Duck, Peter J. Stuckey, Maria Garcia de la Banda, and Christian Holzbaaur. Extending arbitrary solvers with constraint handling rules. In *Proceedings of PPDP'03, International Conference on Principles and Practice of Declarative Programming, Uppsala, Sweden*, pages 79–90. ACM Press, 2003.
- [27] François Fages, Cleyton Mario de Oliveira Rodrigues, and Thierry Martinez. Modular CHR with ask and tell. In Thom Frühwirth and Tom Schrijvers, editors, *Proceedings of the fifth Constraint Handling Rules Workshop CHR'08*, 2008.
- [28] François Fages, Paul Ruet, and Sylvain Soliman. Linear concurrent constraint programming: operational and phase semantics. *Information and Computation*, 165(1):14–41, February 2001.
- [29] Alan M. Frisch, Warwick Harvey, Chris Jefferson, Bernadette Martinez-Hernandez, and Ian Miguel. Essence: A constraint language for specifying combinatorial problems. *Constraints*, 13:268–306, 2008.
- [30] T. Frühwirth. Theory and practice of Constraint Handling Rules. *Journal of Logic Programming, Special Issue on Constraint Logic Programming*, 37(1-3):95–138, October 1998.
- [31] Thom W. Frühwirth. *Constraint Handling Rules*. Cambridge University Press, 2009.

- [32] Harald Ganzinger. A new metacomplexity theorem for bottom-up logic programs. In *Proceedings of International Joint Conference on Automated Reasoning*, volume 2083 of *Lecture Notes in Computer Science*, pages 514–528. Springer-Verlag, 2001.
- [33] Steven Gay, François Fages, Thierry Martinez, Sylvain Soliman, and Christine Solnon. On the subgraph epimorphism problem. *Discrete Applied Mathematics*, 162:214–228, January 2014.
- [34] Ian P. Gent and Barbara Smith. Symmetry breaking during search in constraint programming. In *Proceedings ECAI'2000*, pages 599–603, 1999.
- [35] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1), 1987.
- [36] Cinzia Giusto, Maurizio Gabbrielli, and Maria Chiara Meo. Expressiveness of multiple heads in CHR. In *SOFSEM '09: Proceedings of the 35th Conference on Current Trends in Theory and Practice of Computer Science*, pages 205–216, Berlin, Heidelberg, 2009. Springer-Verlag.
- [37] Google. or-tools, 2014.
- [38] Rémy Haemmerlé. *Fermetures et Modules dans les Langages Concurrents avec Contraintes fondés sur la Logique Linéaire*. PhD thesis, Univ. Paris 7. Soutenance le 17 janvier 2008, December 2007.
- [39] Rémy Haemmerlé and Hariolf Betz. Verification of constraint handling rules using linear logic phase semantics. In *Proceeding of CHR 2008, the fifth Constraint Handling Rules Workshop*, Report Series 08-10. RICS-Linz, July 2008.
- [40] Rémy Haemmerlé and François Fages. Modules for Prolog revisited. In *Proceedings of International Conference on Logic Programming ICLP 2006*, number 4079 in *Lecture Notes in Computer Science*, pages 41–55. Springer-Verlag, 2006.
- [41] Rémy Haemmerlé, François Fages, and Sylvain Soliman. Closures and modules within linear logic concurrent constraint programming. In V. Arvind and Sanjiva Prasad, editors, *Proceedings of FSTTCS 2007, IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 4855 of *Lecture Notes in Computer Science*, pages 544–556. Springer-Verlag, 2007.

- [42] William D. Harvey and Matthew L. Ginsberg. Limited discrepancy search. In *IJCAI'95: Proceedings of the 14th international joint conference on Artificial intelligence*, pages 607–613, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc.
- [43] Pascal Van Hentenryck and Jean-Philippe Carillon. Generality versus specificity: An experience with ai and or techniques. In Howard E. Shrobe, Tom M. Mitchell, and Reid G. Smith, editors, *AAAI*, pages 660–664. AAAI Press / The MIT Press, 1988.
- [44] J. Jaffar, M. Maher, K. Marriott, and P. Stuckey. The semantics of constraint logic programs. *Journal of Logic Programming*, 37(1-3):1–46, October 1998.
- [45] Joxan Jaffar and Jean-Louis Lassez. Constraint logic programming. In *Proceedings of the 14th ACM Symposium on Principles of Programming Languages, Munich, Germany*, pages 111–119. ACM, January 1987.
- [46] Radha Jagadeesan, Vasant Shanbhogue, and Vijay A. Saraswat. Angelic non-determinism in concurrent constraint programming. Technical report, Xerox Parc, 1991.
- [47] Thomas Johnsson. Lambda lifting: transforming programs to recursive equations. In *Proc. of a conference on Functional programming languages and computer architecture*, pages 190–203, New York, NY, USA, 1985. Springer-Verlag New York, Inc.
- [48] Richard E. Korf. Optimal rectangle packing: Initial results. In Enrico Giunchiglia, Nicola Muscettola, and Dana S. Nau, editors, *ICAPS*, pages 287–295. AAAI, 2003.
- [49] R. Kowalski. Predicate logic as programming language. In *IFIP Congress*, pages 569–574, 1974.
- [50] Krzysztof Kuchcinski and Radosław Szymanek. Jacop 4.0.0, 2013.
- [51] Jean-Louis Laurière. A language and a program for stating and solving combinatorial problems. *Artificial Intelligence*, 10(1):29 – 127, 1978.
- [52] R. J. Lipton. The reachability problem requires exponential space. Technical Report 62, New Haven, Connecticut: Yale University, Department of Computer Science, Research, January 1976.
- [53] Michael J. Maher. Logic semantics for a class of committed-choice programs. In *In 4th International Conference on Logic Programming (ICLP'87)*. MIT Press, 1987.

- [54] Julien Martin, Thierry Martinez, and François Fages. Static expansion vs procedural code generation in rule-based modeling languages. In Barry O’Sullivan Javier Larrosa, editor, *Recent Advances in Constraints*, number 6384 in LNAI, pages 38–58. Springer-Verlag, 2011.
- [55] Thierry Martinez. Semantics-preserving translations between linear concurrent constraint programming and constraint handling rules. In *Proceedings of PPDP’10, International Conference on Principles and Practice of Declarative Programming, Edinburgh, UK*, pages 57–66. ACM, 2010.
- [56] Thierry Martinez. Angelic CHR. In Jon Sneyers, editor, *Proceedings of the Eighth International Workshop on Constraint Handling Rules CHR’11*, September 2011.
- [57] Thierry Martinez and François Fages. On translating minizinc constraint models into fitness function for evolutionary algorithms: Application to continuous placement problems. In *Proceedings of the sixth Workshop on Bin Packing and Placement Constraints BPPC’15, associated to CP’15*, September 2015.
- [58] Thierry Martinez, François Fages, and Sylvain Soliman. Search by constraint propagation. In *Proceedings of the 17th International Conference on Principles and Practice of Declarative Programming, PPDP’15*, pages 173–183. ACM, 2015.
- [59] Thierry Martinez, Lumadaiara Vitorino, François Fages, and Abderrahmane Aggoun. On solving mixed shapes packing problems by continuous optimization with the cma evolution strategy. In *Proceedings of the first Computational Intelligence BRICS Congress BRICS-CCI’13*, pages 515–521. IEEE Press, September 2013.
- [60] Ernst W. Mayr. An algorithm for the general petri net reachability problem. In *Proceedings of the thirteenth annual ACM symposium on Theory of computing, STOC ’81*, pages 238–246, New York, NY, USA, 1981. ACM.
- [61] Pedro Meseguer. Constraint satisfaction problems: an overview. *A.I. Communications*, 2(1):3–17, March 1989.
- [62] Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51:125–157, 1991.

- [63] Faten Nabli, Thierry Martinez, François Fages, and Sylvain Soliman. On enumerating minimal siphons in petri nets using CLP and SAT solvers: Theoretical and practical complexity. *Constraints*, pages 1–26, 2015.
- [64] Nicholas Nethercote, Peter J. Stuckey, Ralph Becket, Sebastian Brand, Gregory J. Duck, and Guido Tack. MiniZinc: Towards a standard CP modelling language. In *CP*, pages 529–543, 2007.
- [65] O. Ohrimenko, P.J. Stuckey, and M. Codish. Propagation via lazy clause generation. *Constraints*, 16(9):357–391, 2009.
- [66] Heiko Oetzberger, , and Er Pretschner. Testing concurrent reactive systems with constraint logic programming. In *In Proc. 2nd workshop on Rule-Based Constraint Reasoning and Programming*, 2000.
- [67] Gilles Pesant. A regular language membership constraint for finite sequences of variables. In Mark Wallace, editor, *Principles and Practice of Constraint Programming – CP 2004*, volume 3258 of *Lecture Notes in Computer Science*, pages 482–495. Springer Berlin Heidelberg, 2004.
- [68] Gordon Plotkin. Call-by-name, call-by-value and the lambda calculus. *Theoretical Computer Science*, 1:125–159, 1975.
- [69] Reza Rafah, Kim Marriott, Maria Garcia de la Banda, Nicholas Nethercote, and Mark Wallace. Adding search to zinc. In *CP*, pages 624–629, 2008.
- [70] Andrea Rendl, Tias Guns, Peter J. Stuckey, and Guido Tack. Minisearch: A solver-independent meta-search language for minizinc. In *Principles and Practice of Constraint Programming - 21st International Conference, CP 2015, Cork, Ireland, August 31 - September 4, 2015, Proceedings*, pages 376–392, 2015.
- [71] Davide Sangiorgi. On the bisimulation proof method. *Mathematical Structures in Computer Science*, 8(5):447–479, 1998.
- [72] Vijay Saraswat. A brief introduction to linear concurrent constraint programming. Xerox PARC, 1993.
- [73] Vijay A. Saraswat. *Concurrent constraint programming*. ACM Doctoral Dissertation Awards. MIT Press, 1993.
- [74] Anders Schack-Nielsen and Carsten Schürmann. Celf — a logical framework for deductive and concurrent systems (system description). In *IJCAR '08: Proceedings of the 4th international joint conference on Automated Reasoning*, pages 320–326, Berlin, Heidelberg, 2008. Springer-Verlag.

- [75] T. Schrijvers, B. Demoen, G. Duck, P. Stuckey, and T. Frühwirth. Automatic implication checking for CHR constraint solvers. *Electronic Notes in Theoretical Computer Science*, 147:93–111, January 2006.
- [76] Tom Schrijvers and Thom W. Frühwirth. Analysing the CHR implementation of unionfind. In *Proceedings of the 19th Workshop on (Constraint) Logic Programming, WCLP'05*, Ulm, Germany, 2005.
- [77] Tom Schrijvers, Peter Stuckey, and Philip Wadler. Monadic constraint programming. *Journal of Functional Programming*, 19(6):663, 2009.
- [78] Tom Schrijvers, Peter J. Stuckey, and Gregory J. Duck. Abstract interpretation for constraint handling rules. In *PPDP '05: Proceedings of the 7th ACM SIGPLAN international conference on Principles and practice of declarative programming*, pages 218–229, New York, NY, USA, 2005. ACM.
- [79] Tom Schrijvers, Guido Tack, Pieter Wuille, Horst Samulowitz, and Peter J. Stuckey. Search combinators. *Constraints*, 18(2):269–305, 2013.
- [80] Helmut Simonis and Barry O’Sullivan. Search strategies for rectangle packing. In Peter J. Stuckey, editor, *Proceedings of CP'08*, volume 5202 of *LNCS*, pages 52–66. Springer-Verlag, 2008.
- [81] Jon Sneyers, Wannes Meert, Joost Vennekens, Yoshitaka Kameya, and Taisuke Sato. CHR(PRISM)-based probabilistic logic learning. *Theory and Practice of Logic Programming*, 10(4-6):433–447, 2010.
- [82] Jon Sneyers, Tom Schrijvers, and Bart Demoen. The computational power and complexity of Constraint Handling Rules. In *Proceedings of the second Constraint Handling Rules Workshop, at ICLP'05*, pages 3–17, 2005.
- [83] Jon Sneyers, Tom Schrijvers, and Bart Demoen. The computational power and complexity of constraint handling rules. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 31(2):1–42, 2009.
- [84] Jon Sneyers, Peter Var Weert, Tom Schrijvers, and Leslie De Koninck. As time goes by: Constraint Handling Rules – a survey of CHR research between 1998 and 2007. *Theory and Practice of Logic Programming*, 2, January 2010.
- [85] Martin Sulzmann, Jeremy Wazny, and Peter J. Stuckey. A framework for extended algebraic data types. In *In Proc. of FLOPS'06, volume 3945 of LNCS*, pages 47–64. Springer-Verlag, 2006.
- [86] Robert Endre Tarjan and Jan van Leeuwen. Worst-case analysis of set union algorithms. *Journal of ACM*, 31(2):245–281, April 1984.

- [87] Pascal Van Hentenryck. *The OPL Optimization programming Language*. MIT Press, 1999.
- [88] David L. Waltz. Generating semantic descriptions from drawings of scenes with shadows. Technical report, Massachusetts Institute of Technology, Cambridge, MA, USA, 1972.
- [89] D. Warren. An abstract Prolog instruction set. Technical Note 309, SRI International, Menlo Park, CA, 1983.
- [90] David S. Warren. The Andorra principle. In *Presented at the Gigalips Workshop, Swedish Institute of Computer Science (SICS), Stockholm, Sweden.*, 1988.