université
PARIS
DIDEROT
PARIS 7

Université de Paris

Logical foundations
of a modelling assistant
for molecular biology

Fondements logiques d'un assistant à la modélisation en biologie moléculaire

par Adrien HUSSON

Thèse de doctorat d'Informatique

Dirigée par Jean KRIVINE

Thèse soutenue publiquement le 16 décembre 2019 devant le jury composé de :

Jean Krivine	Directeur de thèse	Chargé de recherche	Université de Paris
Nicolas Peltier	Rapporteur	Chargé de recherche	Université de Grenoble
François Fages	Rapporteur	Directeur de recherche	INRIA Saclay
Ralf Treinen	Président du jury	Professeur	Université de Paris
Anne Siegel	Examinatrice	Directrice de recherche	Université de Rennes
Walter Fontana	Examinateur	Professeur	Harvard University
Hubert Comon	Examinateur	Professeur	ENS Paris-Saclay

Résumé

Résumé   Cette thèse concerne la "représentation exécutable du savoir"dans le domaine de la biologie moléculaire. Elle introduit les fondements d'un cadre logique appelé iota, dont le but est de décrire et rassembler des faits au sujet d'interactions entre protéines tout en offrant au modeleur la possibilité de "compiler" un fragment raisonnable de la logique vers un ensemble fini de règles de réécriture.

On définit une logique FO[↓] qui décrit des transitions d'états cellulaires. Un état représente le contenu d'une cellule : les éléments du domaine sont des parties de protéines et les relations sont des liaisons entre protéines. L'opérateur logique unaire ↓ sélectionne les transitions où un ensemble minimal de changements a lieu. Les formules qui parlent de transitions dénotent aussi des exécutions, c'est-à-dire des séquences finies ou infinies de transitions. Chaque formule de transition est de plus associée à un ensemble de règles de réécritures équipé d'une sémantique opérationnelle. On introduit deux système déductifs qui permettent de "typer" les formules. On montre que si une formule est typable dans le 1er système, alors l'exécution des règles de réécriture qui lui sont associées produit exactement les exécutions dénotées par la formule ; et que si elle est typable dans le 2nd système, alors son système de règles associé est fini. On introduit une grammaire qui produit des formules typables dans les deux systèmes (à équivalence logique près). Enfin, on étudie la décidabilité et l'expressivité de fragments de FO[↓]. On montre en particulier que les formules typables dans le second système sont définissables dans un petit fragment de FO, ce qui implique que l'opérateur ↓ peut alors être éliminé.

Mots-clefs   Raisonnement non-monotone, representation des connaissances, modélisation basée sur les règles de réécriture, biologie moléculaire, logique du 2nd ordre, circonscription

Abstract

This thesis addresses the issue of "Executable Knowledge Representation" in the context of molecular biology. We introduce the foundation of a logical framework, termed iota, whose aim is to facilitate knowledge collation of molecular interactions at the level of proteins and at the same time allows the modeler to "compile" a reasonable fragment of the logic into a finite set of executable graph rewriting rules.

We define a logic FO[↓] over cell state transitions. States represent cell contents; domain elements are protein parts and relations are protein-protein bindings. The unary logical operator ↓ selects transitions where as little as possible happens. Formulas over transitions also denote runs, which are finite or infinite sequences of transitions. Every transition formula is moreover associated to a set of rewriting rules equipped with an operational semantics. We introduce two deductive systems that act as "typing" for formulas. We show that if a formula is typable in the first system then the execution of its associated rule set produces exactly the runs denoted by the formula, and that if it is typable in the second system then its associated rule set is finite. We introduce a grammar that produces formulas typable in both systems, up to logical equivalence. Finally we study decidability and definability properties of fragments of FO[↓]. In particular, we show that formulas typable in the second system are in a tight fragment of FO, which implies that the operator ↓ can then be eliminated.

Keywords   non-monotonic reasoning, knowledge representation, rule-based modelling, molecular biology, second-order logic, circumscription

Contents

Contents

1. Introduction

1.1. Biological modeling

Biological systems involve many components that interact non-linearly. These components can be molecules, cells, genes, etc. The level of complexity of many of these systems precludes a full mental grasp over all their possible behaviors.

For the past several decades, the field of systems biology has developed mathematical models that have become an important tool for biological research. These models explicitly describe the behaviors and interactions of biological systems; they can then be interpreted by a computer for testing and prediction purposes.

We start with a very simple example designed to motivate biological modelling in general.

A boolean network example   We introduce biological modeling with a simple type of model: a boolean network. A boolean network can represent how proteins influence one another through their activity. The toy model in figure 1.1 shows three protein types: A, B and C, as well as their influence on one another's activity. The activity of a protein is summarised as "active" and "inactive". If a protein or a pair of proteins is active, it can turn on or turn off the activity of another protein. The green arrows in figure 1.1 represent activation (positive influence), and the red, flat arrows represent inhibition (negative influence). A double arrow requires both sources to (de)activate. For instance, A alone has a negative influence on B, but A and B together have a positive influence on C. Double arrows have priority over single arrows.

At any moment in time, each of A, B and C can be active or inactive (but we assume that it is not possible for all 3 to be inactive at the same time). That can be summarised by 3 letters ABC. If a letter is green, the corresponding protein is active. Otherwise it is inative. For instance, ABC means that only A and C are active. Each possible triple of green or black letters is called a *state*.

Using the individual and pair influences on figure 1.1, we can generate a *trace* from any starting state. That means starting with any activation values for A, B and C, and then iteratively
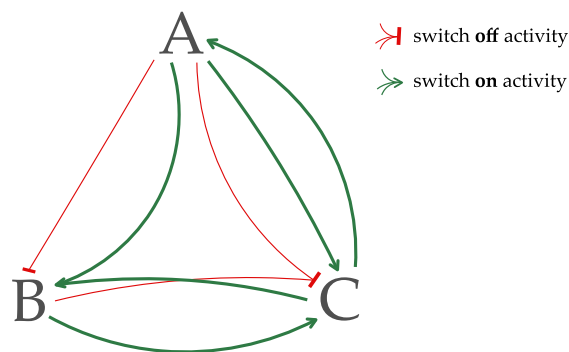

Figure 1.1.: Activity influences between A,B and C. A joint influence has priority.

$$ABC \rightarrow ABC \rightarrow ABC \rightarrow ABC \rightarrow \cdots$$

Figure 1.2.: Generating traces

Figure 1.3.: The resulting network. A green letter is active.

updating those values according to the influence map of figure 1.1.

For instance, figure 1.2 shows a sample trace. The trace starts in state ABC, which means that A is inactive and both B and C are active. Because C has a positive influence on A (cf. the green arrow from C to A in the influence map of figure 1.1), that state can evolve to ABC. Then, because A and B together have a negative influence on C (cf. the double, flat red arrow from (A, B) to C in the influence map), the state can evolve to ABC, and so on.

If we generate many such traces, we will find out that the states ABC, ABC and ABC are not seen very often. In fact, unless we start in one of these states, we never see them at all!

To better understand the situation, we draw a map of all the states and of how traces can move from each one to the others. The resulting diagram is shown in figure 1.3. Each state has A,B,C active or inactive, and arrows are possible moves by the trace-generation process. For instance, there is a (leftward) arrow from ABC to ABC because A and C together have a positive influence on B.

If we look at the diagram carefully, we notice an attractor region (yellow rectangle). On any trace, the system will go to and stay among the attractor's states: since there is no arrow coming out of the yellow region, it is impossible to leave it.

That last observation illustrate the essence of biological modeling: going from a set of individual facts (the joint influences of fig. 1.1) to an insight on the global dynamics (the attractor region of fig. 1.3) through the manipulation and analysis of a mathematical object.

In this thesis, we will be concerned with the earliest part of the modelling process: how to describe the individual facts and collate them together. We also won't be working on boolean networks but on different formal objects, which follow the paradigm of *rule-based modelling* [Boutillier et al., 2018, Chylek et al., 2014, Danos et al., 2007].

1.2. Rule-based models

The purpose of the previous section was to show the biological modelling process, from observations to a formal model to analysis of the model and what that may imply about the real world. We now introduce another biological modelling formalism: *rule-based models*. They will be the starting point of the objects we construct in this thesis.

Boolean networks are typically about modelling gene expression. They are not mechanistic, in the sense that they model statistical interactions at a rather large scale. There is an important, more mechanistic level of abstraction: protein-protein interactions, as described by molecular

biology.

To illustrate the type of knowledge we are aiming at, here is a typical sentence from a molecular biology paper:

> *"The activation of Raf-1 by activated Src requires phosphorylation of Raf-1 on Y340 and/or Y341 [...]. Tyrosine phosphorylation and activation of Raf-1 have been shown to be co-incident. However, others have been unable to detect phosphotyrosine in active Raf-1."*.
> [Mason et al., 1999]

At this level of abstraction, proteins are considered as chains of amino acid residues such as Y340 and Y341, which are identified by their type (Y for tYrosine) and their position in the chain (resp. 340 and 341). Proteins have names, here Raf-1 and Src, and are usually divided into domains or regions that are covering sub-sequences of amino acids. Domains may also be given a name. For instance, Raf-1 has a "Zinc finger" domain in the 137-183 region.[1]

Importantly, *static* names of proteins, domains and residues can be completed with *dynamic* attributes. Here, "phosphorylation" denotes the attachment of a phosphate group to a protein residue, which tends to modify the protein structure. One then talks about a phosphorylated protein, a phosphorylated domain or, as in the example above, a phosphorylated residue. Other dynamic attributes such as "active" are commonplace in molecular biology.

Underlying the snippet of biological literature given above is the notion of protein interactions: "the activation of Raf-1" by "activated Src" indicates that Raf-1 and active Src can bind to each other so that phosphorylation of Raf-1 by Src may occur. Stable binding of proteins requires complementary domains that stick together with various affinities. The binding state of a protein (or a region) is therefore also a dynamic, relational property.

Over time, the field of molecular biology accumulates data and mechanisms suspected to play key roles in the cellular ecosystem. The activity of discovery currently outpaces human abilities to follow and collate new mechanisms [Lazebnik, 2002]. For instance, p53 is a protein family relevant to cell apoptosis and cancer formation. In 2018 alone, it was mentioned in the title or abstract of about 4700 papers (PubMed). The need for principled representations of biological systems has led to the development of formal methods for biology [Fontana and Buss, 1996].

The field of *rule-based modelling* is well-suited to formally modelling mechanistic biological interactions at that lower level of abstraction. Protein-protein interactions are directly described, and gene expression is a global observation on the model's behavior. Rule-based models attempt to collate a large number of known interactions in a single structure: a list of rewrite rules.

Mixtures A rule-based model is a list of *rules* that describe allowed interactions between proteins – think of chemical reactions, but with proteins and protein parts instead of individual atoms. The interactions take place in a *mixture*, a graph which represents the contents of a cell or of a test tube.

Figure 1.4 presents a possible mixture: it contains several *agents* of two types: K for Kinase, E for general enzyme (a kinase is a particular kind of enzyme which helps bind phosphate groups to other proteins). Each agent represents a protein. Some agents have been given names so they can be individually identified: there is an agent a (of type K) and an agent b (of type E). Some agents have *sites* (s, t): each site represents the ability of its protein to interact with another protein. Sites can change in two ways :

---

[1] uniprot.org/uniprot/P09560

Figure 1.4.: A mixture

Figure 1.5.: A sequence of mixtures, changes highlighted in green

- The *label* of a site may change. In figure 1.4, some sites are ON, others not. Labels represent any dynamic aspect of a protein: its state of phosphorylation (that is, whether a phosphate group is attached to part of the protein), its folding conformation, etc.

- The *linkage* of a site may change. If two proteins are linked, they are interacting. As with labels, linkage is dynamic: it may change over time. In the mixture presented in figure 1.4, a and b are linked.

In figure 1.5, we see the sequence of mixtures M_1, M_2 and M_3. They can be seen as three snapshots of a mixture as it evolves in time:

- in M_1, agent a's site s (written a.s) is OFF. a.s and b.s are not linked;

- in M_2, a.s is ON. a.s and b.s are stil not linked.

- in M_3, b.s and a.s are linked.

Rules   Driving the evolution of a mixture are the *rules* mentioned earlier. A rule allows a possible change. In figure 1.6, we see the rule[2] R_1. It allows the site s of any Kinase to spontaneously turn ON when it is not *linked* to anyone.

If an *instance* of a rule is found in a mixture, that mixture can evolve.

Figure 1.8 shows how to find instances of rules: A *match* for R_1 is found by identifying a part of M_1 "similar enough"[3] to the left-hand side of R_1 (cf. red arrow, figure 1.8). Then, the changes

---

[2]In figures, rules are represented with a dashed central triangle in the center, to differentiate from the mere succession of mixtures separated by a grey triangle.

[3]A match is a particular kind of morphism.

Figure 1.6.: The rule R_1

Figure 1.7.: The rule R_2
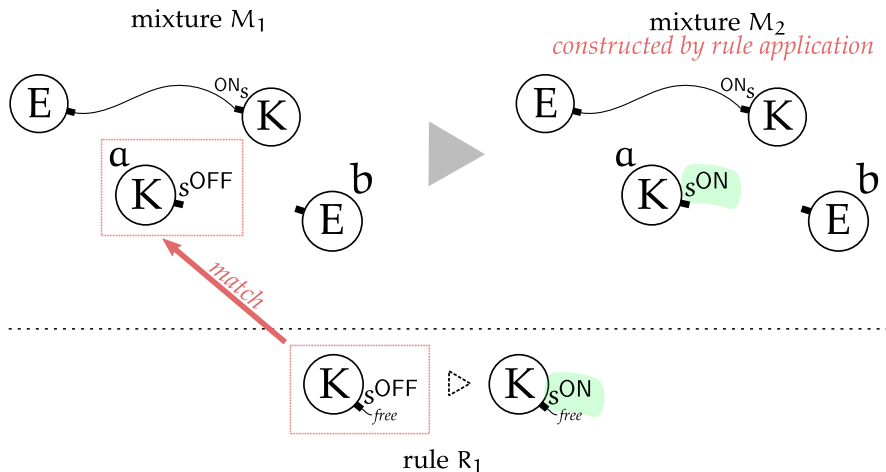
Figure 1.8.: A match from R_1 into M_1

in R_1 (i.e. going from OFF to ON) are applied to M_1 along the match arrow. This constructs the mixture M_2 (top right).

In R_1, notice the word *free* attached to s: it means that for the change to occur, the match image of s in M_1 should not be linked to anyone. The word *free* only appears in rules, not in mixtures, because it is not a property of the graph itself. Instead it is a *application condition* restricting where the rule can be applied.

Figure 1.7 shows the rule R_2, where any two agents can become linked as long as one has an ON site s and that both s sites are free. Following the example of figure 1.8, consider how the rule R_2 could be applied to M_2 in order to construct M_3.

Given a set of rules and an initial mixture one can generate a *trace* by repeating the above operation and at every step nondeterministically choosing one match among the possible ones. Figure 1.5 is actually an example such a sequence, starting from mixture M_1. It uses the rules R_1 and R_2 shown in figures 1.6 and 1.7.

Problems with rule-based models   Rule-based models suffer from two problems that arise as modelling scales up. The first is that models become themselves too large understand. In particular if multiple groups collaborate, it becomes harder to write rules consistent with every other observation, encoding choice and biological invariant of the model. Since model building rests on current biological data, it should be possible to query the existing system; to maintain a coherent model, it is useful to answer queries such as "Is this rule subsumed by a preexisting

Figure 1.9.: R_1 becomes R_1'

one?" or "Is that biological invariant ensured by the current model?". In the next paragraph we introduce a logic-based description of biological knowledge. In such a framework, natural language queries of that type can be encoded as boolean or n-ary logical queries.

A second, related issue is that knowledge should be added *incrementally* as much as possible. Under the rule-based paradigm, it is hard to add new knowledge to an existing set of rules. Sometimes that knowledge is about a new & previously unknown reaction. Then it suffices to create a new rule. But sometimes incorporating the new knowledge means *modifying* existing rules. Take the rule R_1 in figure 1.6, which lets a Kinase go from OFF to ON. Suppose that the following information comes in: *A Kinase turns ON only when the label of its* t *site is OPEN*. We would like that information to exist as a discrete piece of knowledge. But the only way to incorporate it is to edit every single rule where a kinase goes from OFF to ON, as in figure 1.9, where R_1 is edited into R_1'.

Knowledge representation   When one wishes to formally represent information, turning to logic is usually a sensible move. Here is a logic formula; we call if F_0. The x in parenthesis means that it depends on an argument named x. The := symbol means "is defined to be", and $\wedge$ means "and":

$$F_0(x) := OFF(x.s) \wedge ON^\star(x.s)$$

$F_0(x)$ means : *The s site of agent x turns ON.*

- $OFF(x.s)$ means that the site s of x is currently OFF.

- $ON^\star(x.s)$ means that in the *next* mixture, the site s of x will be ON.

Note the presence of a star after the word ON. The star is for things that are susceptible to change (labels and links). No star means we are talking about the *current* mixture, while a star means we are talking about the *next* mixture.

Compare the formula F_0 to the rule R_1 in figure 1.6. Informally, note how they are partially similar: in both cases, some agent's site s goes form OFF to ON. For the rule R_1 the switch is represented by a pair of graph. For the formula F_0, the switch is encoded by two statements taken in conjunction.

We can refine F_0 into another formula, F_1, so that it corresponds more closely to our intuitive understanding of the rule R_1:

$$F_1 := F_0 \wedge K(x) \wedge free(x.s)$$

- $K(x)$ means "x is a Kinase". Names such as Kinase and Enzyme cannot change, so there is no need to also write $K^\star(x)$.

- *free*(x.s) is another formula. It specifies that the site s of x should not be linked to anyone in the *current* mixture (there exists a formula *free** for the *next* mixture).

Figure 1.10.: The transition $\langle M_1, M_2 \rangle$

The meaning of a formula depends on how we *interpret* its variables. $F_0(a)$ means that a is turning ON. $F_0(b)$ means that b is turning ON.

With all its variables interpreted, a formula may or may not be *satisfied* by a pair of mixtures. A pair of mixtures can be represented by a *transition*. One mixture is the *current* one, while the other is the *next*. For instance, $\langle M_1, M_2 \rangle$ is the transition where M_1 is the current mixture, and M_2 is the next one (see figure 1.10)[4]. In $\langle M_1, M_2 \rangle$, $F_1(a)$ is satisfied because a is a Kinase, $a.s$ does turn ON between M_1 and M_2, and s is not linked to anyone. $F_1(b)$ is not satisfied because b does not even have a site s.

One advantage of using logic formulas is that it is easier to incorporate new knowledge. Recall the scenario where a new piece of information comes in: *A Kinase's s site may only turn ON when its t site is OPEN*. Let us write a formula G that says exactly this:

$$G := \forall x. \; (K(x) \wedge OFF(x.s) \wedge ON^\star(x.s)) \to OPEN(x.t)$$

- The universal quantifier $\forall x.$ means that what follows it should be true *for all* interpretations of x, and the arrow $\to$ means "implies".

- As earlier, $K(x)$ means that x has the static name K, OFF is about the current state, and $ON^\star$ is about the next mixture.

- The formula requires all Kinases in the transition with an s site that turns ON to also have a t site OPEN in the current mixture.

In the previous section, the only way to incorporate that new knowledge was to edit R_1 into R_1', as in figure 1.9. Now G is an independent unit of knowledge, and it suffices to write $F_1 \wedge G$.

Traces As we saw earlier, rules can be *executed*. So far there is no notion of execution associated with formulas. Formulas have semantics and we may reason with them but we cannot generate traces with them.

Consider what it would mean for a trace to satisfy (or not) a formula. We wrote the formula F_1 to correspond to the rule R_1. We will also write $F_2(x, y)$ to correspond to the rule R_2:

$$F_2(x, y) := free(x.s) \wedge free(y.t) \wedge ON(y.s) \wedge Link^\star(x.s, y.t)$$

- The formula F_2 depends on two variables, x and y.

- *free*(x.s) stands for "nothing is connected to x.s". Same for y.t.

---

[4]In figures we use a dark grey triangle for transitions. It helps distinguishing them from rules (dashed triangle) and sequences of mixtures (light grey triangle).

Figure 1.11.: The transition from M_1 to M_2 satisfies $F_1(a)$, and therefore F

Figure 1.12.: The transition from M_2 to M_3 satisfies $F_2(b, a)$, and therefore F

- $\mathsf{Link}^\star(x.s, y.t)$ stands for "$x.s$ and $y.t$ are connected in the *next* mixture".

The formula F allows both F_1 and F_2 ($\vee$ means "or"):

$$F(x, y) := F_1(x) \vee F_2(x, y)$$

Consider again the sequence M_1, M_2, M_3 shown in figure 1.5 on page 4. The sequence M_1, M_2, M_3 is correct for the rules R_1 and R_2 because each step corresponds to the application of a rule among $\{R_1, R_2\}$. We will say that M_1, M_2, M_3 is correct for F if every successive pair of mixtures in it satisfies F.

This is illustrated in figures 1.11 and 1.12. Picture a "sliding window" wide enough for 2 mixtures, sliding over the sequence. The pair of mixtures visible under the window is a transition, and either that transition can satisfy F or it cannot. If a visible transition can satisfy F no matter where the window is, then the entire sequence is correct for F.

Change minimisation   With the rule-based approach, a trace is generated by successively choosing which rules to apply where in each successive state. For a formula, the definition we have so far is that every successive *pair* of states should satisfy the formula. This is not a strict enough requirement.

Recall the formula $F_0 = \mathsf{OFF}(x.s) \wedge \mathsf{ON}^\star(x.s)$. Consider the transition $\langle M_1, M_{bad} \rangle$, shown figure 1.13. In this transition, not only does $a.s$ turn ON, but the Enzyme above a disappears (the faint red Enzyme on the right is here to signal that it is missing) and the Kinase above a and b becomes connected to b (and disconnects its previous link).

The transition $\langle M_1, M_{bad} \rangle$ satisfies $F_0(a)$. Logically that is fine, but we want to stay as close the way rule-based trace generation works as possible. Since $\langle M_1, M_{bad} \rangle$ contains superfluous changes, we want to forbid it from appearing in a sequence.

Figure 1.13.: The transition $\langle M_1, M_{bad}\rangle$ satisfies $F_0(a)$

Compare $\langle M_1, M_2\rangle$ on figure 1.10 (p. 7) with $\langle M_1, M_{bad}\rangle$ on figure 1.13. They start from the same mixture M_1. The difference is that $\langle M_1, M_{bad}\rangle$ does strictly more than $\langle M_1, M_2\rangle$ even though $\langle M_1, M_2\rangle$ already satisfies $F_0(a)$.

To fix the issue, we will define a *change order* on transitions. From a same starting mixture, transitions that "do more" will be above transitions that "do less". Then we will introduce a new logical operator, $\downarrow$. With that operator, the formula $\downarrow F_0$ will only be satisfied by transitions that satisfy F_0 with *as few changes as possible*. So $\langle M_1, M_2\rangle$ will satisfy $(\downarrow F_0)(a)$, but $\langle M_1, M_{bad}\rangle$ will not.

To faithfully reproduce the behavior of rules R_1 and R_2, the following formula can replace F:

$$F' := (\downarrow F_1)(x) \vee (\downarrow F_2)(x, y)$$

Goals   Given a formula ϕ that may use $\downarrow$ and a starting mixture, we are interested in generating all the correct traces for ϕ, and only those. Recall the rule-based trace generation process: we start from an initial mixture, find a way to match a rule to the mixture, and create a new mixture from the match. By doing it again and again, we generate a trace. Is there a process where, given a formula ϕ and an initial mixture, we can generate a sequence of new mixtures that together form a correct sequence? Can this process be complete with regards to the set of all correct sequences?

This thesis describes conditions where the answer is positive. More precisely, it starts by rephrasing the question as: *Is it possible to transform a formula into a finite set of rules such that the sequences generated by the rules are exactly those that are correct for the formula?* The rephrasing is particularly relevant because there are existing simulation engines that can take as a input a finite set of rules and efficiently output traces [Boutillier et al., 2018, Chylek et al., 2014] (they run actual simulations in the sense that their trace generation process is parameterised by stochastic rates instead of being nondeterministic). We would like to make use of them. In the next section, we provide an overview of the content of this thesis.

1.3. Plan & Contributions

The thesis is organised as follows: we start by defining transitions, rewrite rules, and how to execute those rules. Then we define a logic for specifying transitions: first-order logic with an operator ($\downarrow$) to minimise changes. We introduce a theory of transitions, a denotational notion of execution (runs), and a canonical mapping from formulas to rules.

A first set of "typing rules", $\vdash_\rho$, produces formulas that can be mapped to rules such that the execution of those rules generates exactly the traces denoted by the formulas.

A second set of typing rules, $\vdash_\lhd$, produces formulas that can be mapped to a finite set of rules. As a result, formulas produced by both $\vdash_\rho$ and $\vdash_\lhd$ can in principle be executed by existing efficient graph rewriting engines [Boutillier et al., 2017] modulo a reasonable theory. We provide a program syntax which always corresponds to $\vdash_\rho$, $\vdash_\lhd$-provable formulas.

The last part deals with decidability and definability issues: we provide a decidable fragment for the theory of transitions, and show when the change-minimisation operator $\downarrow$ can be eliminated.

The contents of chapters 4, 6 and 7, minus full proofs, are in [Husson and Krivine, 2019]. The main contributions of this thesis are *a)* the framework for expressing graph rewriting rules as minimised formulas *b)* the $\vdash_\rho$ deduction rules which produce formulas with infinitarily operational trace semantics, *c)* the $\vdash_\lhd$ deduction rules which, together with $\vdash_\lhd$ produce formulas with finitarily operational run semantics, and *d)* the proof that $\vdash_\rho$-provable formulas are in two very tight fragments of first-order logic, $\exists^*\forall^*$ and $\forall^*\exists^*$.

Plan in more detail:

Structures & Runs   In chapter 2 we introduce *states*, a class of structures, as well as *transitions* between states. Transitions can be summarised by *patterns* and by *rules*. Given a starting state, we describe *a)* an operational notion of rule-based nondeterministic execution that associates a set of *executions* to every rule set and *b)* a denotational notion that associates a set of *runs* to every transition set. The rule-based execution method essentially transports the description of rule-based rewriting of Kappa from [Danos et al., 2012] to signatures and structures. Both runs and executions are sequences of transitions, but we use a different name for each method of generation. We then show conditions when the operational and the denotational notion coincide.

Formulas & Minimisation   In chapter 3 we describe a logic on transitions. We define the semantics of formulas, both as a set of transitions and as a set of runs from initial states. We also introduce the semantics of the minimisation operator $\downarrow$, based on a notion of *change order* between transitions. Except where otherwise noted, the rest of the thesis is in FO[$\downarrow$]. Now that formulas have a sensible notion of execution, we turn to the correspondence between rule sets and formulas.

From formulas to rules   Chapter 3 describes conditions when the execution semantics and the run semantics of formulas are the same. We introduce a first "type system" for formulas, $\vdash_\rho$. We show that for $\vdash_\rho$-provable formulas, there is a set of rules such that, for a class of initial states, their execution produces exactly the runs denoted by the formula.

Bounded changes   The rule sets that correspond to $\vdash_\rho$-provable formulas can be infinite. Morally, the reason is that even after change minimisation the semantics of a formula does not bound the amount of change that can occur during a transition. In chapter 4 we study changes in transitions under a tighter theory, that of n-bounded linked forests, which characterises transitions between forests of bounded-height trees with functional links between tree vertices. We introduce a second "type system" for formulas, $\vdash_\lhd$, such that change-minimal models of $\vdash_\lhd$-provable formulas are bounded in their amount of change.

Compilation: from logic program to rules   In chapter 5 we ask when a formula can effectively be compiled to a finite set of rules. Combining the previous chapters, we find that one answer is:

when the formula is both $\vdash_{\!\!\rho}$ and $\vdash_{\!\!\lhd}$-provable. We introduce a syntax for "programs" that compile to formulas such that any syntatically valid program compiles to a $\vdash_{\!\!\rho}, \vdash_{\!\!\lhd}$-provable formula. More than a toy syntax, programs show admissible rules for $\vdash_{\!\!\lhd}$ and $\vdash_{\!\!\rho}$ which demonstrates the rich syntax available for building compilable formulas.

Decidability   In chapter 6 we consider decidability questions. We show that the FO theory of n-bounded-height linked forest does not have decidable satisfiability, but that its $\exists^*\forall^*$ fragment does.

Change minimality in FO   In chapter 7, we introduce *unified circumscription*. Circumscription dates back to the early 1980s and is a way to characterise the minimal models of formulas along some order, using 2nd-order quantification. We show that the version we dub unified circumscription can express minimality along the change order defined in chapter 3. We then show when the minimality operator $\downarrow$ can be eliminated (thus going from FO[$\downarrow$] to FO). We further show that, modulo the theory of n-bounded-height linked forests, $\vdash_{\!\!\lhd}$-provable formulas are in the in the $\exists^*\forall^*$ and $\forall^*\exists^*$ fragments of FO.

2. Runs and executions for transition structures

The first chapter introduces the main notions used throughout the thesis. In section 2.1 we introduce general notations and notions. Then in section 2.2 we define states and transitions, the basic structures we use in the rest of this thesis and important notions related to them: distance, 0-embeddings, state joining. In section 2.3 we define rules and pushouts, an operational notion of trace generation (executions). Then section 2.4 gives a more denotational version of trace generation with patterns and matches (runs). Section 2.5 shows when the previous two approaches coincide for one step of trace generation. Finally, section 2.6 shows when they coincide on entire traces.

In more detail, this chapter defines *states*, which represent a snapshot of a biological mixture, and *transitions* between those states. A state is a set of elements connected in two ways. First, through *static functions*. They represent the unchanging, rigid structure of a molecule. Second, through *dynamic links*. Those typically represent temporary interactions between molecules; they are subject to change. In addition, unary properties also exist, both in static and dynamic versions. A transition therefore contains only one version of the static information, and two versions of the dynamic information: one for the state "before" the transition, and the other for "after".

Both states and transitions can be generalised to *patterns*: as in rule-based rewriting, a "small" transition can be found in a bigger one. This is called a *match*. A match from a transition to another starts with a structure embedding, but *application conditions* determine when the embedding is a proper match. This supplementary information is found in a pattern.

One can give states and transitions a notion of *execution*. First, given a rule $\imath$ (a variation on the notion of pattern) and a state $\mathfrak{M}$, one may sometimes construct a *pushout*, which is morally a "next state" representing the result of applying $\imath$ to a specific part of $\mathfrak{M}$. In the case of rule-based graph rewriting, this operation was shown figure 1.8 (p.5), where a rule was a pair of graphs equipped with application conditions. From now on, rules will be objects derived from states. This one-step operation can be iterated by repeated application of rules in a set to the most recent "next state", generating an execution.

A more semantic notion is that of pattern *instance*. Instead of being constructed from some "current state" $\mathfrak{M}$, an instance of $\imath$ is any transition in which $\imath$ *could* match. A sequence of such instances can form a *run* – the denotational counterpart to the operational notion of *execution*.

Given a set of patterns, one can derive a corresponding set of rules. The main result of this chapter gives sufficient conditions for the *runs* of a set of patterns to be equal to the *executions* of their derived rules.

A note on figures: in the sequel, figures representing structures use arrows with thick heads for the interpretation of functions and curves for binary relations; in all examples, those relations are symmetric. If a function arrow or a relation curve has no name, then the name is either obvious or unimportant. If a function arrow is not shown on some point, it is the identity there. Names such as a, b, c are for domain elements. If for instance a has a subscript A, it means that a is in the interpretation of A.

2.1. Generalities

In this section we describe various notations and we define the basic model theory concepts used throughout this thesis.

The marker **Notation** (as opposed to **Definition**) introduces notations for objects whose definition is outside the scope of the document, or syntactic sugar for already defined objects. The marker **Implicit parameter** introduces objects in the context of the document. They are then implicitly universally quantified.

Notation (Tuples)**.** If some symbols $(a, x, t, \dots)$ are used to denote a certain kind of objects (variables, nonlogical constants, etc), bold symbols denote tuples of objects of that kind $(\mathbf{a}, \mathbf{x}, \mathbf{t}, \dots)$. Set-like notations apply to tuples; for instance $(a, b) \subseteq (a, x, b, c)$ and $x \in (y, x)$. ∎

Implicit parameter   The universe $\mathcal{U}$ is some recursive, infinite set; the domain of a structure is always a subset of $\mathcal{U}$. An *element* always refers to an element of $\mathcal{U}$.

Basic definitions   A signature Σ is a tuple of symbols. Each symbol $S \in \Sigma$ is equipped with a nature (relational or functional) and an arity $\mathrm{arity}(S) \geqslant 0$.

A Σ-structure $\mathfrak{A}$ (written $\mathfrak{A} : \Sigma$) is a domain $\mathrm{dom}(\mathfrak{A}) \subseteq \mathcal{U}$ together with an interpretation for every symbol of Σ, that is: a k-ary relation on $\mathrm{dom}(\mathfrak{A})$ (resp. k-arguments function on $\mathrm{dom}(\mathfrak{A})$) $[\![S]\!]_{\mathfrak{A}}$ for every k-ary relation symbol (resp. function symbol) $S \in \Sigma$. If $\Sigma' \subseteq \Sigma$, $[\![\Sigma']\!]_{\mathfrak{A}}$ is the tuple of interpretations of Σ'-symbols in $\mathfrak{A}$. The Σ'-reduct of $\mathfrak{A}$ is the Σ'-structure with domain $\mathrm{dom}(\mathfrak{A})$ and interpretations $[\![\Sigma']\!]_{\mathfrak{A}}$.

If $\mathfrak{B} : \Sigma$, $m : \mathrm{dom}(\mathfrak{A}) \to \mathrm{dom}(\mathfrak{B})$ is injective, and $m([\![S]\!]_{\mathfrak{A}}(\mathbf{a})) = [\![S]\!]_{\mathfrak{B}}(m(\mathbf{a}))$ (resp. $\mathbf{a} \in [\![S]\!]_{\mathfrak{A}}$ iff $m(\mathbf{a}) \in [\![S]\!]_{\mathfrak{B}}$) for every function symbol (resp. relation symbol) $S \in \Sigma$ and tuple $\mathbf{a}$ from $\mathrm{dom}(\mathfrak{A})$ of size $\mathrm{arity}(S)$, then m is an embedding from $\mathfrak{A}$ to $\mathfrak{B}$ (written $m : \mathfrak{A} \to \mathfrak{B}$).

If $\mathrm{id}_{\mathfrak{A}} : \mathfrak{A} \to \mathfrak{B}$, then $\mathfrak{A}$ is a substructure of $\mathfrak{B}$. If $K \subseteq \mathcal{U}$, the substructure of $\mathfrak{A}$ induced by K (written $\mathfrak{A} \upharpoonright K$) is the uniquely defined substructure of $\mathfrak{A}$ with $\subseteq$-minimal domain still containing $K \cap \mathrm{dom}(\mathfrak{A})$.

If $\Sigma' \subseteq \Sigma$ and $K \subseteq \mathcal{U}$, $[\![\Sigma']\!]_{\mathfrak{A}} \upharpoonright K := [\![\Sigma']\!]_{\mathfrak{A}' \upharpoonright K}$ where $\mathfrak{A}'$ is the reduct of $\mathfrak{A}$ to Σ'.

2.2. States and transitions

In this section we describe *states* and *transitions*, which are classes of structures, as well as *changes*. We then introduce 3 notions: a pseudodistance δ on structure elements that identifies functionally accessible elements, a join operation for states (when they are similar enough) that produces transitions, and a stricter notion of structure embeddings that respects the identities produced by the pseudodistance δ.

States are the *mixtures* from the introduction. Transitions represent going from one state to another in a single step. First, we describe the class of state and transitions *signatures*:

Definition (State and transition signature)**.** If

- **Dyn** is a tuple of relational symbols such that $|\mathbf{Dyn}| \geqslant 2$ and

  - One distinguished symbol, Link, is binary.

  - Another distinguished symbol, Has, is unary.

  - All other symbols are unary.

- **Dyn*** is **Dyn** with a star* appended to each symbol,

- **Stat** is a tuple of unary symbols, relational or functional,

then $\Theta := (\mathbf{Dyn}, \mathbf{Dyn}^\star, \mathbf{Stat})$ is a *transition signature* and $\Theta^\circ := (\mathbf{Dyn}, \mathbf{Stat})$ is its corresponding *state signature*. ∎

Connection with rule-based models   We explain the connection with the rule-based models from the introduction. Recall mixtures and transitions from the introduction: some names were static, such as K (for Kinase). A static name would be a unary predicate symbol in **Stat**. Other names were dynamic, such as ON. That name would be in **Dyn**. Its counterpart for the next mixture, ON*, would be in **Dyn***. An agent in a mixture would be an element, say a. Its site would be another element, and the name of the site would be the name of a function $s \in$ **Stat**. So $a.s$ refers to the image of a by the interpretation of s.

So an interpretation for **Stat** represents the unchanging part of a structure, while there are two copies of **Dyn**. The interpretation of **Dyn** holds the *precondition*; the interpretation of **Dyn*** holds the *postcondition*. In a biological setting, the static information in the interpretation of **Stat** represents proteins names, internal protein structure, etc. The intepretations of **Dyn** and **Dyn*** hold the dynamical properties: phosphorylation state, folding, link to another protein, etc.

Has and Link   We motivate the distinguished predicate and relation symbols Has, Has*, Link and Link*. The distinguished binary Link symbol is for protein-protein interactions. There is no technical issue with adding more relational symbols, but we have just one to keep things simple. Moreover, we single out the unary Has to encode *presence*: in a logical specification (see later), we would like to say e.g. $\mathsf{Has}(x) \wedge \neg\mathsf{Has}^\star(x)$ to ensure the destruction of x during a transition.

Notation. If Θ is a transition signature and $\mathfrak{A} : \Theta$, then $\mathfrak{A}$ is a *transition*. If $\mathfrak{M} : \Theta^\circ$, then $\mathfrak{M}$ is a *state*. The names $\mathfrak{A}, \mathfrak{B}, \mathfrak{C}$ and variations are reserved for transitions. The names $\mathfrak{M}, \mathfrak{S}$ and variations are reserved for states ($\mathfrak{M}$ for "*m*ixture" and $\mathfrak{S}$ for "*s*tate"). ∎

Definition (Function symbols of **Stat**). The tuple $\mathbf{Stat}_{\mathrm{fun}} \subseteq \mathbf{Stat}$ contains exactly the functional symbols of **Stat**. ∎

Notation. Symbols in **Dyn** and **Dyn*** are said to be *dynamic*; symbols in **Stat** are said to be *static*. In the sequel, the possible fonts are:

- Sans-serif for dynamic symbols.

- Regular for static symbols state/transition signature.

- *Italic* for formula names (see chapter 3).

- Math for other symbols

In any state or transition signature, any symbol $A \in \mathbf{Dyn}$ has its counterpart $A^\star \in \mathbf{Dyn}^\star$ of same nature and same arity. ∎

Changes   We introduce *changes* to capture the difference between the pre- and postconditions. The notation prefigures the use of semantic brackets for formulas.

Definition (Changes). If $\mathfrak{A} : \Theta$, for every $A \in \mathbf{Dyn}$ let $[\![\Delta A]\!]_{\mathfrak{A}} := [\![A]\!]_{\mathfrak{A}} \Delta [\![A^\star]\!]_{\mathfrak{A}}$, where Δ is the symmetric difference. ∎

Definition (Transition changes). If $\mathfrak{A} : \Theta$, let $[\![\Delta]\!]_{\mathfrak{A}}$ be the union of the carrier sets of $[\![\Delta A]\!]_{\mathfrak{A}}$ for all $A \in \mathbf{Dyn}$, and $[\![\neg\Delta]\!]_{\mathfrak{A}}$ its complement in $\mathfrak{A}$. ∎

Implicit parameter $\Theta := (\mathbf{Dyn}, \mathbf{Dyn}^\star, \mathbf{Stat})$ is a transition signature.

Example. The transition $\mathfrak{A}$ is represented graphically figure 2.1. The static structure (interpretation of $\mathbf{Stat}$) is repeated left and right of the arrow ►. The interpretation of $\mathbf{Dyn}$ is left of the arrow: there is a symmetric Link between a and c. The interpretation of $\mathbf{Dyn}^\star$ is right of the arrow: the Link between a and c is gone, and b has gained the $A^\star$ property.

If the image of some function on an element is not shown, it implicitly means that the function is the identity at that point. For instance, $\llbracket f \rrbracket_{\mathfrak{A}}(b) = b$. ∎


$$\mathbf{Dyn} = (A, Link)$$
$$\mathbf{Dyn}^\star = (A^\star, Link^\star)$$
$$\mathbf{Stat}^\star = (f)$$
$$\llbracket A \rrbracket_{\mathfrak{A}} = \emptyset$$
$$\llbracket A^\star \rrbracket_{\mathfrak{A}} = \{b\}$$
$$\llbracket \Delta \rrbracket_{\mathfrak{A}} = \{a, b, c\}$$
$$\llbracket \Delta Link \rrbracket_{\mathfrak{A}} = \{(a, c), (c, a)\}$$

Figure 2.1.: Transition $\mathfrak{A}$.

The next three subsections introduce concepts important to states and transitions: distance, 0-embeddings, and how to create transitions from states. The distance treats paths through $\mathbf{Stat}_{fun}$ differently from paths through $Link, Link^\star$. 0-embeddings relate two structures by defining when one can be "found" in the other. Joins trivially construct a transition from two states, but they are not always well-defined.

2.2.1. Pseudodistance

In this subsection we define a pseudodistance between structure elements.

In the sequel, many properties will depend on identifying a small part of a state or transition which should satisfy some constraint, and beyond which constraints are looser. As we will see later, the difficulty comes from the dynamic part of the structures, because static structures are always conserved through morphisms (see next subsection) and when necessary become bounded in size (see chapter 4). So we will define a notion of pseudodistance which identifies all elements *statically* connected and only counts as a "far away" elements which may become disconnected from one another. In a state or a transition, function symbols are only in $\mathbf{Stat}_{fun}$; they represent immutable structure of a mixture. On the other hand, links ($Link$ and $Link^\star$) are dynamic: they can appear and disappear. So we define a notion of pseudodistance which ignores functions:

Definition (Pseudodistance δ). If $\mathfrak{A} : \Theta$ or $\mathfrak{A} : \Theta^\circ$, let $G_{\mathfrak{A}}$ be the undirected, weighted graph with vertices $\mathrm{dom}(\mathfrak{A})$ and, for every $a, b \in \mathrm{dom}(\mathfrak{A})$:

- an edge $\{a, b\}$ of weight 0 whenever $\llbracket h \rrbracket_{\mathfrak{A}}(a) = b$ for some $h \in \mathbf{Stat}_{fun}$,

- otherwise, an edge $\{a, b\}$ of weight 1 whenever $(a, b) \in \llbracket R \rrbracket_{\mathfrak{A}}$ for some binary, relational $R \in \Theta$ or Θ°.

The following function is a pseudo[1]distance: for every $a, b \in \mathrm{dom}(\mathfrak{A})$, $\delta_{\mathfrak{A}}(a, b)$ is the length of

---

[1]$\delta(a, b) = 0$ does not imply $a = b$.

the shortest a—b path in G. ■

Balls Taking balls of radius d around elements or subdomains of structures will be very useful in the sequel. However, note that the ball of radius 0 is particularly meaningful here since δ is a pseudodistance: for an element a in a structure, the set of elements at distance 0 (its 0-neighbors) are the elements functionally and inverse-functionally accessible from a.

Definition (Balls). If $\mathfrak{A} : \Theta$ of $\mathfrak{A} : \Theta^\circ$, $d \geqslant 0$ and $a \in \mathrm{dom}(\mathfrak{A})$, the *ball of radius d around a* is

$$\mathcal{B}_{\mathfrak{A}}(a, d) := \{\, b \in \mathfrak{A} \mid \delta_{\mathfrak{A}}(a, b) \leqslant d \,\}$$

and naturally extends to subsets of $\mathcal{U}$: if $K \subseteq \mathcal{U}$, $\mathcal{B}_{\mathfrak{A}}(K, d)$ is the union of the balls of radius d around each $a \in K \cap \mathrm{dom}(\mathfrak{A})$. ■

An important feature of δ is that elements that are *statically* connected are considered "the same" from a distance perspective.

Example. In figure 2.2, the state $\mathfrak{S}$ has 4 elements. b is symetrically Linked to d. The names of the functions in $\mathbf{Stat}_{\mathrm{fun}}$ are not shown, we only see arrows, indicating that for some $f_1, f_2 \in \mathbf{Stat}_{\mathrm{fun}}$, $[\![f_1]\!]_{\mathfrak{S}}(a) = c$ and $[\![f_2]\!]_{\mathfrak{S}}(a) = b$. As earlier, function images not shown implicitly mean a loop. For instance, $[\![f_2]\!]_{\mathfrak{S}}(b) = b$. The corresponding graph $G_{\mathfrak{S}}$ is shown on the right, and the ball of radius 0 around b is shown in a semitransparent blue outline. ■

Figure 2.2.: State $\mathfrak{S}$ and distance between its elements.

Using the above notion of distance, it is possible to say when two states are similar enough to be joined into a transition structure.

2.2.2. Joining states into transitions

In this subsection we give the canonical way to assemble states into transitions. The operation is very simple, and the side-condition ensures there is no clash between the static parts of each state. We then give a simple lemma which relates distance in joined states and in each individual state.

For any two states $\mathfrak{S}$, $\mathfrak{S}'$, both of their domains are in a common universe $\mathcal{U}$. We forgo looking for partial isomorphisms between states and simply consider the common elements of $\mathrm{dom}(\mathfrak{S})$ and $\mathrm{dom}(\mathfrak{S}')$. If the static structure (that is, the 0-ball) around them is the same, then $\mathfrak{S}$ and $\mathfrak{S}'$ are "similar enough" to be meaningfully joined into a new transition:

Figure 2.3.: $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ is well-defined.

Figure 2.4.: $\langle \cdot, \cdot \rangle$ not defined on $\mathfrak{S}, \mathfrak{S}'$

Definition ($\langle \cdot, \cdot \rangle$). If $\mathfrak{S}, \mathfrak{S}' : \Theta^\circ$ and $[\![\mathbf{Stat}]\!]_{\mathfrak{S}} \upharpoonright \mathcal{B}_{\mathfrak{S}}(\mathrm{dom}(\mathfrak{S}'), 0) = [\![\mathbf{Stat}]\!]_{\mathfrak{S}'} \upharpoonright \mathcal{B}_{\mathfrak{S}'}(\mathrm{dom}(\mathfrak{S}), 0)$, then $\langle \mathfrak{S}, \mathfrak{S}' \rangle : \Theta$ is:

$$\mathrm{dom}(\langle \mathfrak{S}, \mathfrak{S}' \rangle) := \mathrm{dom}(\mathfrak{S}) \cup \mathrm{dom}(\mathfrak{S}')$$

$$[\![A]\!]_{\langle \mathfrak{S}, \mathfrak{S}' \rangle} := [\![A]\!]_{\mathfrak{S}} \qquad\qquad \text{for all } A \in \mathbf{Dyn}$$

$$[\![A^\star]\!]_{\langle \mathfrak{S}, \mathfrak{S}' \rangle} := [\![A]\!]_{\mathfrak{S}'} \qquad\qquad \text{for all } A \in \mathbf{Dyn}$$

$$[\![T]\!]_{\langle \mathfrak{S}, \mathfrak{S}' \rangle} := [\![T]\!]_{\mathfrak{S}} \cup [\![T]\!]_{\mathfrak{S}'} \qquad\qquad \text{for all } T \in \mathbf{Stat}$$

$\langle \mathfrak{S}, \mathfrak{S}' \rangle$ is well-defined thanks to the proviso on $[\![\mathbf{Stat}]\!]_{\mathfrak{S}}$ and $[\![\mathbf{Stat}]\!]_{\mathfrak{S}'}$. ∎

Example (and counterexample). The states $\mathfrak{S}$ and $\mathfrak{S}'$ in figure 2.3 respect the conditions necessary for $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ to be well-defined.

On the other hand, in figure 2.4, $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ is not defined because $[\![f]\!]_{\mathfrak{S}} \upharpoonright \{a\} = \{(a, a)\}$ and $[\![f]\!]_{\mathfrak{S}'} \upharpoonright \{a, b\} = \{(a, a), (b, a)\}$. ∎

Formally, the condition $[\![\mathbf{Stat}]\!]_{\mathfrak{S}} \upharpoonright \mathcal{B}_{\mathfrak{S}}(\mathrm{dom}(\mathfrak{S}'), 0) = [\![\mathbf{Stat}]\!]_{\mathfrak{S}'} \upharpoonright \mathcal{B}_{\mathfrak{S}'}(\mathrm{dom}(\mathfrak{S}), 0)$ is more than is necessary for $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ to be well-defined. It would suffice to require $[\![\mathbf{Stat}]\!]_{\mathfrak{S}} \upharpoonright \mathfrak{S}' = [\![\mathbf{Stat}]\!]_{\mathfrak{S}'} \upharpoonright \mathfrak{S}$. But then joins such as in figure 2.4 would be allowed. The issue with these joins is that static structures are supposed to *not* change from one state to the next. The creation of d and of a b–d link between $\mathfrak{S}$ and $\mathfrak{S}'$ is allowed because presence and linkage are dynamic properties. But the appearance of an a in the function graph of b violates the design rule about static functions and predicates. More formally, the information loss about the structure of $\mathfrak{S}$ (in figure 2.4) would prevent later operations that extract states back from transitions; those operations use the lemma we are just about to state, which becomes false if we use the looser requirement.

We explain that next lemma. After a join, the distance between two elements can become smaller, for instance by the addition of a link. It is possible for two elements that were infinitely far away in two states $\mathfrak{S}$ and $\mathfrak{S}'$ to be at finite distance in $\langle \mathfrak{S}, \mathfrak{S}' \rangle$. However (keeping with the spirit of static functions being really static), the 0 bound is strict: elements at distance > 0 in a state cannot be at distance 0 after a join.

Lemma 2.2.1. If $\langle \mathfrak{M}, \mathfrak{M}' \rangle$ is defined, $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(a, b) = 0$ and $a \in \mathrm{dom}(\mathfrak{M})$, (resp $a \in \mathrm{dom}(\mathfrak{M}')$) then $\delta_{\mathfrak{M}}(a, b) = 0$ (resp. $\delta_{\mathfrak{M}'}(a, b) = 0$).

Proof. We show for $\mathfrak{M}$, the other version is symmetric.

First, take any $a', b' \in \mathrm{dom}(\langle \mathfrak{M}, \mathfrak{M}' \rangle)$ such that $(a', b') \in [\![h]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$ for some function symbol $h \in \mathbf{Stat}$. We show that if a' or b' is in $\mathrm{dom}(\mathfrak{M})$, the other is as well, and $(a', b') \in [\![h]\!]_{\mathfrak{M}}$.

By definition of $[\![h]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$, the interesting case is $(a', b') \in ([\![h]\!]_{\mathfrak{M}'})$; if one of a', b' is in $\mathfrak{M}$, say a' for concreteness, then since $[\![\mathbf{Stat}]\!]_{\mathfrak{M}} \upharpoonright \mathcal{B}_{\mathfrak{M}}(\text{dom}(\mathfrak{M}'), 0) = [\![\mathbf{Stat}]\!]_{\mathfrak{M}'} \upharpoonright \mathcal{B}_{\mathfrak{M}'}(\text{dom}(\mathfrak{M}), 0)$, $(a', b') \in [\![h]\!]_{\mathfrak{M}}$ as well.

By the above and by induction on the number of 0-edges of any a–b path in $G_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$, we get $a \in \text{dom}(\mathfrak{M})$ and $\delta_{\mathfrak{M}}(a, b) = 0$. □

We will make use of the operation $\langle \cdot, \cdot \rangle$ throughout this chapter and the next.

2.2.3. 0-embedding

In this subsection we define a restriction of embeddings called 0-embeddings and show that 0-embeddings compose.

0-embeddings are the basis for defining when a small transition can be "found" in a bigger one. They follow the mantra started earlier: elements at distance 0 are immutably connected, so a morphism should not be allowed to give new 0-neighbors to an element.

0-embeddings are not powerful enough, however: we will later introduce the concept of a match, which refines 0-embeddings further by adding *application conditions*, as seen in the introduction, of the form "the image of this element may not gain new 1-neighbors". Unlike the constraint of 0-embeddings on the absence of new 0-neighbors, which is universal, matches will specify which elements cannot gain new 1-neighbors.

Definition (0-embedding). If $\mathfrak{A}, \mathfrak{B} : \Theta$ or $\mathfrak{A}, \mathfrak{B} : \Theta°$, an embedding $m : \mathfrak{A} \to \mathfrak{B}$ is a 0-*embedding* if

$$\mathcal{B}_{\mathfrak{B}}(m(\mathfrak{A}), 0) \subseteq m(\mathfrak{A})$$

∎

A 0-embedding is simply an embedding that respects the pseudodistance δ's identities. An alternative design would have been to only consider the class of structures where the following is true: for every pair of elements a, b, b accessible from a in the function graph implies a accessible from b in the function graph. For such structures, all embeddings are 0-embeddings.

Example (Counterexample). In figure 2.5, the identity $\text{id}_{\mathfrak{S}}$ is an embedding but it is *not* a 0-embedding since the 0-ball around b in $\mathfrak{M}$ is not part of $\text{Im}(\text{id}_{\mathfrak{A}})$. ∎

Figure 2.5.: $\text{id}_{\mathfrak{A}} : \mathfrak{S} \to \mathfrak{M}$

The following lemma illustrates the action of a 0-embedding (note that it does not hold for embeddings):

Lemma 2.2.2. If $m : \mathfrak{A} \to \mathfrak{B}$ is a 0-embedding, $a \in \text{Im}(m)$ and $\delta_{\mathfrak{B}}(a, b) = 0$, then $b \in \text{Im}(m)$ and $\delta_{\mathfrak{B}}(m^{-1}(a), m^{-1}(b)) = 0$

Proof. Since $\mathcal{B}_{\mathfrak{B}}(\text{Im}(m), 0) \subseteq \text{Im}(m)$, $b \in \text{Im}(m)$ and by induction on the number of 0-edges in any path a–b in $G_{\mathfrak{B}}$, since m is an embedding $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(b)) = 0$. □

Lemma 2.2.2 immediately implies composition of 0-embeddings:

Lemma 2.2.3. If $\mathfrak{m} : \mathfrak{A} \to \mathfrak{B}$ and $\mathfrak{m}' : \mathfrak{B} \to \mathfrak{C}$ are both 0-embeddings, then $\mathfrak{m}' \circ \mathfrak{m}$ is a 0-embedding as well.

Proof. Immediate by lemma 2.2.2. $\qquad\square$

With the above ingredients (state joining, 0-embeddings and distance), it is possible to define rules and how to apply them to states.

2.3. Rules, pushouts, executions

Now that states, transitions, and the morphisms between them are established, it is time to define our structure-based version of rule application in graph rewriting. In this section, we define those and show that a rule applied to a *current* state always creates a *next* state which can be joined to the current one.

Consider two states $\mathfrak{S}, \mathfrak{S}'$ and another state $\mathfrak{M}$ such that it makes sense to talk about $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ and such that $\mathfrak{S}$ can in some sense be found in $\mathfrak{M}$. We want to *transport the differences between* $\mathfrak{S}$ *and* $\mathfrak{S}'$ *to* $\mathfrak{M}$, thus creating a new state $\mathfrak{M}'$. The operation defined below is meant to be an idealised version of how graph rewriting frameworks such as Kappa work.

Definition (Pushout)**.** If $\mathfrak{S}, \mathfrak{S}', \mathfrak{M} : \Theta^\circ$, $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ is defined, $f : \mathfrak{S} \to \mathfrak{M}$ is a 0-embedding, $\mathcal{B}_\mathfrak{M}(f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}')), 1) \subseteq \text{Im}(f)$ and $g : \text{dom}(\mathfrak{S}') \setminus \text{dom}(\mathfrak{S}) \to \mathcal{U} \setminus \text{dom}(\mathfrak{M})$ is injective, then the *pushout* of $(f, g, \mathfrak{S}, \mathfrak{S}', \mathfrak{M})$ is $\mathfrak{M}'$:

$$\text{dom}(\mathfrak{M}') := \text{dom}(\mathfrak{M}) \setminus f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}')) \uplus \text{Im}(g)$$

$$[\![A]\!]_{\mathfrak{M}'} := [\![A]\!]_\mathfrak{M} \setminus f([\![A]\!]_\mathfrak{S} \setminus [\![A]\!]_{\mathfrak{S}'}) \uplus (f \uplus g)([\![A]\!]_{\mathfrak{S}'} \setminus [\![A]\!]_\mathfrak{S}) \qquad \text{for all } A \in \mathbf{Dyn}$$

$$[\![T]\!]_{\mathfrak{M}'} := [\![T]\!]_\mathfrak{M} \restriction \text{dom}(\mathfrak{M}') \cup (f \uplus g)([\![T]\!]_{\mathfrak{S}'}) \qquad \text{for all } T \in \mathbf{Stat}$$

It is well-defined. First, for all $T \in \mathbf{Stat}$ $[\![T]\!]_\mathfrak{M} \restriction \text{dom}(\mathfrak{M}')$ and $(f \uplus g)([\![T]\!]_{\mathfrak{S}'})$ agree on their intersection because $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ is defined and the image of g does not intersect $\text{dom}(\mathfrak{M})$. Next, the domain of $\mathfrak{M}'$ should cover the interpretation of all symbols in Θ°:

- If $\mathfrak{a} \in [\![A]\!]_\mathfrak{M}$ and for some $a \in \mathfrak{a}$, $a \in f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}'))$ then by $\mathcal{B}_\mathfrak{M}(f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}')), 1) \subseteq \text{Im}(f)$, and the fact that f is an embedding, $\mathfrak{a} \in f([\![A]\!]_\mathfrak{S})$. Since $f^{-1}(a) \notin \text{dom}(\mathfrak{S})$, $f^{-1}(\mathfrak{a}) \notin [\![A]\!]_{\mathfrak{S}'}$, and so $\mathfrak{a} \notin [\![A]\!]_{\mathfrak{M}'}$.

- Suppose there is $\mathfrak{a} \in [\![A]\!]_{\mathfrak{S}'} \setminus [\![A]\!]_\mathfrak{S}$. For every $a \in \mathfrak{a}$, a is not in $\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}')$, moreover if $(f \uplus g)(a) \notin \text{dom}(\mathfrak{M})$ then $a \in \text{dom}(\mathfrak{S}') \setminus \text{dom}(\mathfrak{S})$ by assumption on f, g.

- For the relational symbols of $\mathbf{Stat}$, trivial. We show that if $h \in \mathbf{Stat}_{\text{fun}}$, $a \in \text{dom}(\mathfrak{M}')$ and $(a, b) \in [\![h]\!]_{\mathfrak{M}'}$, then $b \in \text{dom}(\mathfrak{M}')$:
  - If $(a, b) \in (f \uplus g)([\![T]\!]_{\mathfrak{S}'})$ then $b \in (f \uplus g)(\text{dom}(\mathfrak{S}'))$. If $b \in \text{Im}(g)$, $b \in \text{dom}(\mathfrak{M}')$ by definition of $\text{dom}(\mathfrak{M}')$, otherwise $b \in \text{Im}(f)$, so $b \in \text{dom}(\mathfrak{M})$ by definition of f, and $b \notin f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}'))$, so $b \in \text{dom}(\mathfrak{M}')$ by definition of $\text{dom}(\mathfrak{M}')$.
  - If $(a, b) \in [\![T]\!]_\mathfrak{M} \restriction \text{dom}(\mathfrak{M}')$ then $b \in \text{dom}(\mathfrak{M})$, so we need to show $b \notin f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}'))$. Suppose $b \in \text{Im}(f)$. Since f is a 0-embedding, $a \in \text{Im}(f)$ and $(f^{-1}(a), f^{-1}(b)) \in [\![h]\!]_\mathfrak{S}$. Since $a \in \text{dom}(\mathfrak{M}')$, by definition of $\text{dom}(\mathfrak{M}')$ we have $f^{-1}(a) \in \text{dom}(\mathfrak{S}')$ (otherwise a would be in $f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}'))$). So by assumption on $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ (that is, $[\![\mathbf{Stat}]\!]_\mathfrak{S} \restriction \text{dom}(\mathfrak{S}') = [\![\mathbf{Stat}]\!]_{\mathfrak{S}'} \restriction \text{dom}(\mathfrak{S})$) $(f^{-1}(a), f^{-1}(b)) \in [\![h]\!]_{\mathfrak{S}'}$. So $f^{-1}(b) \in \text{dom}(\mathfrak{S}')$, so $b \notin f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}'))$.

■

The definition above realizes literally transports the differences between $\mathfrak{S}$ and $\mathfrak{S}'$ along f and g, and applies those to $\mathfrak{M}$ in order to construct $\mathfrak{M}'$. Morally, $\mathfrak{M}'$ is to $\mathfrak{M}$ what $\mathfrak{S}'$ is to $\mathfrak{S}$.

Design note. To illustrate why the proviso $\mathcal{B}_{\mathfrak{M}}(f(\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}'), 1) \subseteq \text{Im}(f)$ is necessary, consider some $a \in \text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}')$ such that there is $b \in \text{dom}(\mathfrak{M}) \setminus \text{Im}(f)$ and $(f(a), b) \in [\![\text{Link}]\!]_{\mathfrak{M}}$. In other words: a "disappears" in $\langle \mathfrak{S}, \mathfrak{S}' \rangle$, yet at the same time the f-image of a has a neighbor in $\mathfrak{M}$. Here the definition of $\mathfrak{M}'$ would leave a out of $\text{dom}(\mathfrak{M}')$ but keep $(f(a), b)$ in $[\![\text{Link}]\!]_{\mathfrak{M}'}$, which would result in an ill-defined structure.

We believe that it may be possible to have a more complex pushout definition and get rid of the proviso. The more complex definition would discard atoms that mention disappearing elements (such as the pair $(f(a), b)$). As a point of interest, this seems related to the difference between forms of graph rewriting which can allow side-effects (SPO) or forbid them (DPO). ■

Example. This is an example of a pushout. In figure 2.6, $\text{id}_{\mathfrak{S}} : \mathfrak{S} \to \mathfrak{M}$ is a 0-embedding. $\text{dom}(\mathfrak{S}) \setminus \text{dom}(\mathfrak{S}') = \{c\}$, and c does not get any new link when going to $\mathfrak{M}$. $g = \text{id}_{\{e\}}$ maps the new element e of $\mathfrak{S}'$. Since $b \in [\![A]\!]_{\mathfrak{S}}$ but $b \notin [\![A]\!]_{\mathfrak{S}'}$, $b \in [\![A]\!]_{\mathfrak{M}}$ but $b \notin [\![A]\!]_{\mathfrak{M}'}$. ■

Figure 2.6.: $\mathfrak{M}'$ is the pushout of $(\text{id}_{\mathfrak{S}}, g, \mathfrak{S}, \mathfrak{S}', \mathfrak{M})$.

This next lemma shows that the pushout construction (which correspond to rule application, as we will see later) creates a *next* state similar enough to the *current* one for the join operation to be well-defined:

Lemma 2.3.1. If $\mathfrak{M}'$ is the pushout of $(f, g, \mathfrak{S}, \mathfrak{S}', \mathfrak{M})$ then $\langle \mathfrak{M}, \mathfrak{M}' \rangle$ is defined.

Proof. We want $[\![\mathbf{Stat}]\!]_{\mathfrak{M}} \upharpoonright \mathcal{B}_{\mathfrak{M}}(\text{dom}(\mathfrak{M}'), 0) = [\![\mathbf{Stat}]\!]_{\mathfrak{M}'} \upharpoonright \mathcal{B}_{\mathfrak{M}'}(\text{dom}(\mathfrak{M}), 0)$. Let $m := f \uplus g$. First, note that since g does not touch $\mathfrak{M}$, and $\text{dom}(\mathfrak{S}) \cap \text{dom}(\mathfrak{S}') \subseteq \mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0)$, the following holds:

$$\text{dom}(\mathfrak{M}) \cap m(\text{dom}(\mathfrak{S}')) \subseteq m(\mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0)) \tag{2.1}$$

Moreover g defined on $\mathfrak{S}' \setminus \mathfrak{S}$ and $\mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0) \subseteq \text{dom}(\mathfrak{S})$ provide the $\supseteq$-direction in the following equality (the other direction is trivial):

$$f([\![T]\!]_{\mathfrak{S}'} \upharpoonright \mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0)) = m([\![T]\!]_{\mathfrak{S}'}) \upharpoonright m(\mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0)) \tag{2.2}$$

(note that the left-hand side of (2.2) is well-defined since any ball is closed under function application for every function of a structure).

Now:

$$\llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathcal{B}_{\mathfrak{M}}(\text{dom}(\mathfrak{M}'), 0)$$

$$= f(\llbracket T \rrbracket_{\mathfrak{S}} \upharpoonright \mathcal{B}_{\mathfrak{S}}(\text{dom}(\mathfrak{S}'), 0)) \cup \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathcal{B}_{\mathfrak{M}}(\text{dom}(\mathfrak{M}'), 0) \qquad \text{since } f \text{ is a 0-embedding}$$

$$= f(\llbracket T \rrbracket_{\mathfrak{S}'} \upharpoonright \mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0)) \cup \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathcal{B}_{\mathfrak{M}}(\text{dom}(\mathfrak{M}'), 0) \qquad \text{since } \langle \mathfrak{S}, \mathfrak{S}' \rangle \text{ is defined}$$

$$= m(\llbracket T \rrbracket_{\mathfrak{S}'}) \upharpoonright m(\mathcal{B}_{\mathfrak{S}'}(\text{dom}(\mathfrak{S}), 0)) \cup \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \text{dom}(\mathfrak{M}') \qquad \text{by (2.2) (left) and trivially (right)}$$

$$= m(\llbracket T \rrbracket_{\mathfrak{S}'}) \upharpoonright \text{dom}(\mathfrak{M}) \cup \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \text{dom}(\mathfrak{M}') \qquad \text{by (2.1)}$$

$$= \big(m(\llbracket T \rrbracket_{\mathfrak{S}'}) \cup \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \text{dom}(\mathfrak{M}') \big) \upharpoonright \text{dom}(\mathfrak{M}) \qquad \text{trivial}$$

$$= \llbracket T \rrbracket_{\mathfrak{M}'} \upharpoonright \text{dom}(\mathfrak{M}) \qquad \text{by definition of } \llbracket T \rrbracket_{\mathfrak{M}'}$$

$\square$

A rule contains the data necessary to build a pushout plus some *additional* information:

Definition (Rule). A *rule* $\imath$ is a triple of the form $(\mathfrak{S}, K, \mathfrak{S}')$ such that $\mathfrak{S}, \mathfrak{S}' : \Theta^\circ$ and $K \subseteq \mathfrak{S}$. ∎

The K part, called the application condition of a rule, is a form of *application condition* from graph rewriting. When a rule (in the graph rewriting sense) is applied to a graph, the larger context may add edges between nodes present in the rule and nodes from the context. To forbid that, one can specify nodes that should be *free*, i.e. nodes that should not gain edges by going into a larger context. Given a 0-embedding $m : \mathfrak{A} \to \mathfrak{B}$ between two structures with a defined distance, an element $a \in \text{dom}(\mathfrak{A})$ gains no edges (no 1-neighbor) if the following is true :

$$\mathcal{B}_{\mathfrak{B}}(m(a), 1) \subseteq \text{Im}(m)$$

Note that since m is an embedding, new edges must touch the context, that is: the complement of the image of m.

We define what a rule application is and explain more afterwards:

Definition (Rule(s) application). If $\imath := (\mathfrak{S}, K, \mathfrak{S}')$ is a rule and $\mathfrak{M}, \mathfrak{M}' : \Theta^\circ$ then $\mathfrak{M} \twoheadrightarrow_{\imath} \mathfrak{M}'$ if there are f, g such that $\mathfrak{M}'$ is the pushout of $(f, g, \mathfrak{S}, \mathfrak{S}', \mathfrak{M})$ and $\mathcal{B}_{\mathfrak{M}}(f(K), 1) \subseteq \text{Im}(f)$.

In that case, f, g are *pushout witnesses* for $\mathfrak{M} \twoheadrightarrow_{\imath} \mathfrak{M}'$.

If R is a set of rules, then $\mathfrak{M} \twoheadrightarrow_R \mathfrak{M}'$ if $\mathfrak{M} \twoheadrightarrow_{\imath} \mathfrak{M}'$ for some $\imath \in R$

∎

Example. This example shows how to apply the left-hand side of a rule to a state. As illustrated in figure 2.7, the identity on $\mathfrak{S}$ is a good candidate for a successful application of the rule $(\mathfrak{S}, \{a, c\}, \mathfrak{S}')$ (for some $\mathfrak{S}'$ not given here). $d \notin \{a, c\}$ so the new link between d and e is not a problem. Note that c is free even though it has a link; the issue is newly formed links, not existing one (in that sense our notion of *free* is different from the usual one in graph rewriting).

∎

We have formalised the rules and rule applications given in the introduction. In particular, a rule $(\mathfrak{S}, K, \mathfrak{S}')$ must find a 0-embedding between its left-hand side $\mathfrak{S}$ and a state $\mathfrak{M}$, with the application condition $\mathcal{B}_{\mathfrak{A}}(m(K), 1) \subseteq m(\text{dom}(\mathfrak{S}))$. In the language of rule-based models, the rule is said to have a *match* in the larger mixture $\mathfrak{M}$.

However, formulas (introduced later) specify sets of *transitions*. In the next section, we define how the notion of *match* can be defined between transitions. We will later link that denotational notion to the rule application we have just defined.

Figure 2.7.: $\mathrm{id}_{\mathfrak{S}} : \mathfrak{S} \to \mathfrak{M}$

2.4. Patterns and matches

In this section, we introduce patterns, which are transitions equipped with *application conditions* (just as rules had an application in the previous section). Then, the notion of 0-embedding is refined to that of match: a match from a pattern is just a 0-embedding that respects the application conditions of the pattern. The set of transitions where a pattern p can match is called the set of instances of p.

We also give some simple lemmas and show that matches compose.

Patterns

Definition (Pattern). If Σ is a signature, $\mathfrak{A} : \Sigma$ and $K \subseteq \mathrm{dom}(\mathfrak{A})$, the pair a $(\mathfrak{A}, K)$ is a *pattern*. $(\mathfrak{A}, K) : \Sigma$ means that $(\mathfrak{A}, K)$ is a pattern and $\mathfrak{A} : \Sigma$. ∎

A stricter notion of pattern is sometimes necessary because of the proviso $\mathcal{B}_{\mathfrak{M}}(f(\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}')), 1) \subseteq \mathrm{Im}(f)$ in the definition of a pushout. If an element a in a pattern $(\langle \mathfrak{S}, \mathfrak{S}' \rangle, K)$ disappears and has gained an edge after a morphism m into a transition $\langle \mathfrak{M}, \mathfrak{M}' \rangle$, the proviso will not be respected (here f will be $m \upharpoonright \mathrm{dom}(\mathfrak{S})$) and $\mathfrak{M}'$ will not be a pushout. Pure patterns avoid the issue – and the patterns we will most often manipulate later on will all be pure by construction.

Definition (Pure pattern). If $(\mathfrak{A}, K)$ is a pattern and $[\![\Delta]\!]_{\mathfrak{A}} \subseteq K$, then $(\mathfrak{A}, K)$ is a *pure* pattern. ∎

At this point, it is reasonable to ask why we define a notion so similar to rules. Rules and pushouts are operational, and should represent an idealised version of how actual rewriting programs work. But we will later reason using the denotational notion of pattern and match. In the denotational setting, the definition of pattern is natural and easier to manipulate. We will soon make explicit when the operational and the denotational views coincide: in the next section, we show the relationship between patterns and rules, as well as between matches and pushouts.

As a side note, we hope to unify patterns and rules in a later simplification of the formalism where, at the cost of a more complex definition of $\langle \cdot, \cdot \rangle$ (namely keeping track of which elements come from the left and which elements come from the right), there is no difference between the information contained in a pattern and in a rule.

Matches A pattern is like a rule that has already integrated its two states, but there is in general no unique way to decompose a pattern into states. However, there is a much more direct definition of when a pattern can be "found" in a larger one:

Definition (Match). If $(\mathfrak{A}, K_1), (\mathfrak{B}, K_2) : \Theta$, a 0-embedding $m : \mathfrak{A} \to \mathfrak{B}$ is a *match* from $(\mathfrak{A}, K_1)$ to $(\mathfrak{B}, K_2)$ whenever

- $\mathcal{B}_{\mathfrak{B}}(m(K_1), 1) \cup [\![\Delta]\!]_{\mathfrak{B}} \subseteq \mathrm{Im}(m)$ and

- $m(K_1) \subseteq K_2$

The notation is $m : (\mathfrak{A}, K_1) \to (\mathfrak{B}, K_2)$, or $m : (\mathfrak{A}, K_1) \to \mathfrak{B}$ if $K_2 = \mathrm{dom}(\mathfrak{B})$. The set K in $(\mathfrak{A}, K)$ is called the *application condition* of $(\mathfrak{A}, K)$. ∎

Matches give meaning to the application condition of a pattern. As mentioned earlier, in a graph rewrite rule, a node will sometimes be marked *free*; that is: not connected to anyone. Unlike a regular node which may gain edges through a match in a larger graph, *free* nodes must not gain edge. Similarly, an element in the application condition of a pattern cannot gain an edge; as with rule application, this is expressed with a proviso of the form $\mathcal{B}_{\mathfrak{B}}(m(K_1), 1) \subseteq \mathrm{Im}(m)$.

The condition $m(K_1) \subseteq K_2$ ensures matches compose (but we will most often refine to a pattern of the form $(\mathfrak{B}, \mathrm{dom}(\mathfrak{B}))$, thus trivialising the condition).

Finally, $[\![\Delta]\!]_{\mathfrak{B}} \subseteq \mathrm{Im}(m)$ ensures that a match cannot create additional changes. To contrast with rule application: with rules, the next state is mechanically derived in the form of a pushout, so the changes $[\![\Delta]\!]$ are induced by the definition. With patterns, the set of acceptable changes must be explicitly restricted.

Matches are also defined between states; the definition just considers matches between trivial transitions created by joining a single state with itself:

Definition (State match). If $(\mathfrak{S}, K), (\mathfrak{M}, K') : \Theta^\circ$ and $m : (\langle \mathfrak{S}, \mathfrak{S} \rangle, K) \to (\langle \mathfrak{M}, \mathfrak{M} \rangle, K')$ then m *refines* $(\mathfrak{S}, K)$ to $(\mathfrak{M}, K')$ (written $m : (\mathfrak{S}, K) \to (\mathfrak{M}, K')$). ∎

One maybe helpful intuition for rules and patterns is: rules & pushouts define "building a next state from a state and a transition motif" while patterns & matches define "adding context to a transition motif".

Example. The identity $\mathrm{id}_{\mathfrak{A}}$ figure 2.8 is a match from the pattern $(\mathfrak{A}, \{c\})$ to $\mathfrak{B}$. The element c is marked *free* in the precondition of $\mathfrak{A}$ to indicate that c is part of the application condition. One can check that $[\![\Delta]\!]_{\mathfrak{A}} = \{a\} \subseteq \mathrm{Im}(m)$ and that c does not gain any new link. However, b is not in the free set and the match adds a link between d and b. ∎

Figure 2.8.: $\mathrm{id}_{\mathfrak{A}} : (\mathfrak{A}, \{c\}) \to \mathfrak{B}$

Instances Finally it is useful to directly talk about all the transitions into which a pattern can match.

Definition (Instances). If there exists an m such that $m : (\mathfrak{A}, K) \to \mathfrak{B}$, then $\mathfrak{B}$ is an *instance* of $(\mathfrak{A}, K)$. $\mathrm{Insts}((\mathfrak{A}, K))$ is the set of instances of $(\mathfrak{A}, K)$. If P is a set of patterns, $\mathrm{Insts}(P)$ is the union of $\mathrm{Insts}((\mathfrak{A}, K))$ for $(\mathfrak{A}, K) \in P$. ∎

Note that instances only consider transitions, not other patterns. In the match of figure 2.8, $\mathfrak{B}$ is an instance of $(\mathfrak{A}, \{c\})$.

2.4.1. Properties of matches

In this subsection we give useful properties of matches. The first is that an isomorphism is always a match:

Lemma 2.4.1. If $m : \mathfrak{A} \to \mathfrak{B}$ is an isomorphism, then $m : (\mathfrak{A}, K_A) \to (\mathfrak{B}, K_B)$ whenever $m(K_A) \subseteq K_B$.

Proof. $m(\text{dom}(\mathfrak{A})) = \text{dom}(\mathfrak{B})$. $\qquad\square$

The next lemma shows the effect of application conditions. Recall that with a 0-embedding, an element can not gain new 0-neighbors. With a match from a pattern $(\mathfrak{A}, K)$, elements outside of K are allowed to gain new 1-neighbors, but elements in K are *not*:

Lemma 2.4.2. If $m : (\mathfrak{A}, K) \to \mathfrak{B}$, $a \in m(K), b \in \text{dom}(\mathfrak{B})$ and $\delta_{\mathfrak{B}}(a, b) = 1$ then $b \in \text{Im}(m)$ and $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(b)) = 1$.

Proof. By definition of a match $b \in \text{Im}(m)$. There are c, d such that $\delta_{\mathfrak{B}}(a, c) = 0$, (c, d) or $(d, c) \in \llbracket \text{Link} \rrbracket_{\mathfrak{B}} \cup \llbracket \text{Link}^\star \rrbracket_{\mathfrak{B}}$ and $\delta_{\mathfrak{B}}(d, b) = 0$. Applying lemma 2.2.2 yields $c, d \in \text{Im}(m)$ and $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(c)) = 0$ and $\delta_{\mathfrak{A}}(m^{-1}(b), m^{-1}(d)) = 0$. Since m is an embedding, $\delta_{\mathfrak{A}}(m^{-1}(c), m^{-1}(d)) \leqslant 1$, so $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(b)) \leqslant 1$. $\delta_{\mathfrak{B}}(a, b) \leqslant \delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(b))$ so $\delta_{\mathfrak{A}}(a, b) = 1$. $\qquad\square$

Just as the "no new 0-neighbor" property helps show that 0-embeddings compose, the previous lemma helps show that matches compose:

Lemma 2.4.3. If $m_1 : (\mathfrak{A}_1, K_1) \to (\mathfrak{A}_2, K_2)$ and $m_2 : (\mathfrak{A}_2, K_2) \to (\mathfrak{A}_3, K_3)$, then $m_2 \circ m_1 : (\mathfrak{A}_1, K_1) \to (\mathfrak{A}_3, K_3)$.

Proof. The nontrivial part is showing $\mathcal{B}_{\mathfrak{A}_3}(m_2 \circ m_1(K_1), 1) \subseteq m_2 \circ m_1(\mathfrak{A}_1)$. Consider $a \in m_2(m_1(K_1))$ and b such that $\delta_{\mathfrak{A}_3}(a, b) \leqslant 1$. Since $m_1(K_1) \subseteq K_2$ and $\mathcal{B}_{\mathfrak{A}_3}(m_2(K_2), 1) \subseteq m_2(\mathfrak{A}_2)$, $b \in m_2(\mathfrak{A}_2)$. By lemmas 2.2.2 or 2.4.2 (depending on whether $\delta_{\mathfrak{A}_3}(a, b) = 0$ or 1), $\delta_{\mathfrak{A}_2}(m_2^{-1}(a), m_2^{-1}(b)) \leqslant 1$. Since $m_2^{-1}(a) \in m_1(K_1)$ and $\mathcal{B}_{\mathfrak{A}_2}(m_1(K_1), 1) \subseteq m_1(\mathfrak{A}_1)$, $m_2^{-1}(b) \in m_1(\mathfrak{A}_1)$.
The rest is by lemma 2.2.3. $\qquad\square$

The following lemma is a form of weakening: if a match guarantees an application condition K_A and relies on an application condition K_B, then it guarantees any condition K_A' weaker than K_A and it relies on any condition K_B' stronger than K_B:

Lemma 2.4.4. If $m : (\mathfrak{A}, K_A) \to (\mathfrak{B}, K_B)$, $K_A' \subseteq K_A$, and $K_B \subseteq K_B'$, then $m : (\mathfrak{A}, K_A') \to (\mathfrak{B}, K_B')$.

Proof. Apply lemma 2.4.1 twice on the identity, and lemma 2.4.3 once.
$\qquad\square$

The next lemma is important. It generalises the "no new 1-neighbor" consequence of application conditions: if an element is at distance d of the application condition boundary (on the inside), then it has no new $(d + 1)$-neighbor:

Lemma 2.4.5. If $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}(K, d)) \to \mathfrak{B}$, $a \in m(K), b \in \text{dom}(\mathfrak{B})$ and $\delta_{\mathfrak{B}}(a, b) \leqslant d + 1$ then $b \in m(\mathfrak{A})$ and $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(b)) = \delta_{\mathfrak{B}}(a, b)$.

Proof. By induction on $k := \delta_{\mathfrak{B}}(a, b)$. If 0, lemma 2.2.2. Otherwise true for $k' < k$ and there is $c \in \text{dom}(\mathfrak{B})$ such that $\delta_{\mathfrak{B}}(a, c) = k - 1$ and $\delta_{\mathfrak{B}}(b, c) = 1$. By induction $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(c)) = k - 1 \leqslant d$, so $c \in m(\mathcal{B}_{\mathfrak{A}}(K, d))$, so by lemma 2.4.2: $b \in m(\mathfrak{A})$ and $\delta_{\mathfrak{A}}(m^{-1}(c), m^{-1}(b)) = 1$; so $\delta_{\mathfrak{A}}(m^{-1}(a), m^{-1}(b)) = k$ (distances do not get longer in an embedding). $\qquad\square$

Now that the proper notions are introduced, we can connect patterns and matches on the one hand, with rules and pushouts on the other.

2.5. Patterns and rules

In this section, we show that we can go back and forth between rule application and pattern match if the pattern is pure and the transition respects a technical condition.

We start with simple technical definitions and introduce a partial converse of joining states together: a way to split a transition into a precondition state and a postcondition state. We also introduce the notion of support, which is central in this thesis. Recall the distinguished predicates Has and Has* from section 2.2: a transition is supported if for all its elements, at least of one Has and Has* is true.

We start with some general notions necessary for the sequel. The symmetry of transition signatures can be used to flip a transition around:

Definition (Mirror of a transition)**.** If $\mathfrak{A} : \Theta$, its *mirror* $\mathfrak{A}^{-1}$ is $\mathfrak{A}$ with the interpretations of **Dyn** and **Dyn*** flipped; that is, $[\![A]\!]_{\mathfrak{A}^{-1}} = [\![A^\star]\!]_{\mathfrak{A}}$ and $[\![A^\star]\!]_{\mathfrak{A}^{-1}} = [\![A]\!]_{\mathfrak{A}}$ for every $A \in$ **Dyn**. ∎

The next definition singles out the dynamic unary predicate Has. There is no unique way to decompose a transition into states, so given a transition, we pick the functional closure of the interpretation of Has as the substructure representing the "before" state, and the functional closure of the interpretation of Has* as the substructure representing the "after" state.

Definition (Pre and post states)**.** If $\mathfrak{A} : \Theta$, $\mathrm{Pre}(\mathfrak{A}) : \Theta^\circ$ is the reduct of $\mathfrak{A} \upharpoonright [\![\mathrm{Has}]\!]_{\mathfrak{A}}$ to Θ° and $\mathrm{Post}(\mathfrak{A})$ is $\mathrm{Pre}(\mathfrak{A}^{-1})$. ∎

Note that in general, it is *not* the case that $\mathrm{Pre}(\langle \mathfrak{S}, \mathfrak{S}' \rangle) = \mathfrak{S}$ and $\mathrm{Post}(\langle \mathfrak{S}, \mathfrak{S}' \rangle) = \mathfrak{S}'$.

The next definition is used throughout this thesis:

Definition (Supported transition)**.** If $\mathfrak{A} : \Theta$, $\mathfrak{A}$ is *supported* whenever $\mathrm{dom}(\mathfrak{A}) = [\![\mathrm{Has}]\!]_{\mathfrak{A}} \cup [\![\mathrm{Has}^\star]\!]_{\mathfrak{A}}$.
∎

Definition (Strongly supported transition)**.** If $\mathfrak{A} : \Theta$, $\mathfrak{A}$ is *strongly supported* whenever for all $a \in \mathrm{dom}(\mathfrak{A})$, $\mathcal{B}_{\mathfrak{A}}(a, 0) \subseteq [\![\mathrm{Has}]\!]_{\mathfrak{A}}$ or $\mathcal{B}_{\mathfrak{A}}(a, 0) \subseteq [\![\mathrm{Has}^\star]\!]_{\mathfrak{A}}$. ∎

Any strongly supported transition is of course supported. Support means that a transition can be split into two states without information loss. Strong support means that those two states can be merged back together using $\langle \cdot, \cdot \rangle$.

Lemma 2.5.1. If $\mathfrak{A}$ is strongly supported then $\langle \mathrm{Pre}(\mathfrak{A}), \mathrm{Post}(\mathfrak{A}) \rangle = \mathfrak{A}$

Proof. Trivial. ☐

The next lemma goes *from pushout to match*: given the application of some rule $(\mathfrak{S}, K, \mathfrak{S}')$ to a state $\mathfrak{M}$, producing a new state $\mathfrak{M}'$, when can one build a match from $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ to $\langle \mathfrak{M}, \mathfrak{M}' \rangle$?

Lemma 2.5.2. If $\mathfrak{M} \dashrightarrow_{(\mathfrak{S}, K, \mathfrak{S}')} \mathfrak{M}'$ with pushout witnesses f, g, then for any $K' \subseteq \mathrm{dom}(\mathfrak{S}') \setminus \mathrm{dom}(\mathfrak{S})$, $(f \uplus g) : (\langle \mathfrak{S}, \mathfrak{S}' \rangle, K \uplus K') \to \langle \mathfrak{M}, \mathfrak{M}' \rangle$.

Proof.

$f \uplus g$ is a 0-embedding $f \uplus g$ is injective. Now we show the structure-preserving properties. First, for all $A \in \mathbf{Dyn}$:

$$\llbracket A \rrbracket_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \upharpoonright \mathrm{Im}(f \uplus g)$$
$$= \llbracket A \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{Im}(f \uplus g) \qquad\qquad \text{by def. of } \langle \mathfrak{M}, \mathfrak{M}' \rangle$$
$$= \llbracket A \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{Im}(f) \qquad\qquad \text{by } \mathrm{dom}(\mathfrak{M}) \cap \mathrm{Im}(g) = \emptyset$$
$$= f(\llbracket A \rrbracket_{\mathfrak{S}}) \qquad\qquad \text{since } f \text{ is an embedding}$$
$$= (f \uplus g)(\llbracket A \rrbracket_{\mathfrak{S}}) \qquad\qquad g \text{ defined nowhere on } \mathrm{dom}(\mathfrak{S})$$
$$= (f \uplus g)(\llbracket A \rrbracket_{\langle \mathfrak{S}, \mathfrak{S}' \rangle}) \qquad\qquad \text{by def. of } \langle \mathfrak{S}, \mathfrak{S}' \rangle$$

Now for all $A^\star \in \mathbf{Dyn}^\star$:

$$\llbracket A^\star \rrbracket_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \upharpoonright \mathrm{Im}(f \uplus g)$$
$$= \llbracket A \rrbracket_{\mathfrak{M}'} \upharpoonright \mathrm{Im}(f \uplus g) \qquad\qquad \text{by def. of } \langle \mathfrak{M}, \mathfrak{M}' \rangle$$
$$= (\llbracket A \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{Im}(f \uplus g)) \setminus f(\llbracket A \rrbracket_{\mathfrak{S}} \setminus \llbracket A \rrbracket_{\mathfrak{S}'}) \cup (f \uplus g)(\llbracket A \rrbracket_{\mathfrak{S}'} \setminus \llbracket A \rrbracket_{\mathfrak{S}}) \qquad \text{by def. of } \llbracket A \rrbracket_{\mathfrak{M}'}$$
$$= f(\llbracket A \rrbracket_{\mathfrak{S}}) \setminus f(\llbracket A \rrbracket_{\mathfrak{S}} \setminus \llbracket A \rrbracket_{\mathfrak{S}'}) \cup (f \uplus g)(\llbracket A \rrbracket_{\mathfrak{S}'} \setminus \llbracket A \rrbracket_{\mathfrak{S}}) \qquad \text{by assumption on } f, g$$
$$= (f \uplus g)(\llbracket A \rrbracket_{\mathfrak{S}} \setminus (\llbracket A \rrbracket_{\mathfrak{S}} \setminus \llbracket A \rrbracket_{\mathfrak{S}'}) \cup (\llbracket A \rrbracket_{\mathfrak{S}'} \setminus \llbracket A \rrbracket_{\mathfrak{S}})) \qquad \text{by injectivity of } f \uplus g$$
$$= (f \uplus g)(\llbracket A \rrbracket_{\mathfrak{S}'}) \qquad\qquad \text{trivial set operations}$$
$$= (f \uplus g)(\llbracket A^\star \rrbracket_{\langle \mathfrak{S}, \mathfrak{S}' \rangle}) \qquad\qquad \text{by def. of } \langle \mathfrak{S}, \mathfrak{S}' \rangle$$

For all $T \in \mathbf{Stat}$, by definition of $f \uplus g$ and the fact that f is an embedding :

$$\llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{Im}(f \uplus g) = \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{Im}(f) = f(\llbracket T \rrbracket_{\mathfrak{S}}) = (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}}) \tag{2.3}$$

Moreover, by definition of $\mathrm{dom}(\mathfrak{M}')$:

$$f(\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}')) \cap \mathrm{dom}(\mathfrak{M}') = \emptyset \tag{2.4}$$

So:

$$\llbracket T \rrbracket_{\mathfrak{M}'} \upharpoonright \mathrm{Im}(f \uplus g)$$
$$= \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{dom}(\mathfrak{M}') \upharpoonright \mathrm{Im}(f \uplus g) \cup (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad \text{by def. of } \llbracket T \rrbracket_{\mathfrak{M}'}$$
$$= \llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{dom}(\mathfrak{M}') \upharpoonright \mathrm{Im}(f) \cup (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad \text{since } \mathrm{Im}(g) \cap \mathrm{dom}(\mathfrak{M}') = \emptyset$$
$$= f(\llbracket T \rrbracket_{\mathfrak{S}}) \upharpoonright \mathrm{dom}(\mathfrak{M}') \cup (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad \text{since } f \text{ is an embedding}$$
$$= f(\llbracket T \rrbracket_{\mathfrak{S}} \upharpoonright \mathrm{dom}(\mathfrak{S}')) \cup (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad \text{by (2.4)}$$
$$= f(\llbracket T \rrbracket_{\mathfrak{S}'} \upharpoonright \mathrm{dom}(\mathfrak{S})) \cup (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad \text{since } \langle \mathfrak{S}, \mathfrak{S}' \rangle \text{ is defined}$$
$$= (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad\qquad \text{trivial} \tag{2.5}$$

Therefore :

$$\llbracket T \rrbracket_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \upharpoonright \mathrm{Im}(f \uplus g)$$
$$= (\llbracket T \rrbracket_{\mathfrak{M}} \upharpoonright \mathrm{Im}(f \uplus g)) \cup (\llbracket T \rrbracket_{\mathfrak{M}'} \upharpoonright \mathrm{Im}(f \uplus g)) \qquad \text{by def. of } \langle \mathfrak{M}, \mathfrak{M}' \rangle$$
$$= (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}}) \cup (f \uplus g)(\llbracket T \rrbracket_{\mathfrak{S}'}) \qquad \text{by (2.3) (left), (2.5) (right)}$$
$$= (f \uplus g)(\llbracket T \rrbracket_{\langle \mathfrak{S}, \mathfrak{S}' \rangle}) \qquad\qquad \text{by def. of } \langle \mathfrak{S}, \mathfrak{S}' \rangle$$

Now we know that $(f \uplus g)$ is an embedding. To show that it is a 0-embedding, we show that if $(b, a) \in [\![h]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$ and $a \in \text{Im}(f \uplus g)$, then $b \in \text{Im}(f \uplus g)$ (the other direction is obvious). The rest is by induction on the minimal number of 0-edges between a and b in $G_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$.

If $a \in \text{Im}(f)$, suppose $b \in \text{dom}(\mathfrak{M})$. Then $b \in \text{Im}(f)$ since f is a 0-embedding. If $b \notin \text{dom}(\mathfrak{M})$, then by definition of $\text{dom}(\mathfrak{M}')$, $b \in \text{Im}(g)$.

If $a \in \text{Im}(g)$, then $(f \uplus g)^{-1}(a) \notin \text{dom}(\mathfrak{M})$, so $(b, a) \in (f \uplus g)([\![h]\!]_{\mathfrak{S}'})$, so $b \in \text{Im}(f \uplus g)$.

$[\![\Delta]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \subseteq \textbf{Im} \, (\boldsymbol{f} \uplus \boldsymbol{g})$ For all $A \in \textbf{Dyn}$, by definition of $\mathfrak{M}, \mathfrak{M}'$, $a \in [\![A]\!]_{\mathfrak{M}} \setminus [\![A]\!]_{\mathfrak{M}'}$ implies $a \in \text{Im}(f)$ and $a \in [\![A]\!]_{\mathfrak{M}'} \setminus [\![A]\!]_{\mathfrak{M}}$ implies $(f \uplus g)([\![A]\!]_{\mathfrak{S}'} \setminus [\![A]\!]_{\mathfrak{S}})$.

$\mathcal{B}_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \, ((\boldsymbol{f} \uplus \boldsymbol{g}) \, (\boldsymbol{K} \uplus \boldsymbol{K}') \, , \boldsymbol{1}) \subseteq \textbf{Im} \, (\boldsymbol{f} \uplus \boldsymbol{g})$ Let $e \in K \uplus K'$, $d \in \text{dom}(\langle \mathfrak{M}, \mathfrak{M}' \rangle)$, $a := (f \uplus g)(e)$ and $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(a', b) \leqslant 1$.

By definition there are b, c such that $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(a, b) = 0$, (b, c) or $(c, b) \in [\![\text{Link}]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \cup [\![\text{Link}^\star]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$, and $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(c, d) = 0$. Since $f \uplus g$ is a 0-embedding, we have: $c \in \text{Im}(f \uplus g)$ implies $d \in \text{Im}(f \uplus g)$

So it now suffices to show $c \in \text{Im}(f \uplus g)$.

- If the link is in $[\![\text{Link}^\star]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} \setminus [\![\text{Link}]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}$ then by definition of $\langle \mathfrak{M}, \mathfrak{M}' \rangle$, $c \in \text{Im}(f \uplus g)$.

- Otherwise the link is in $[\![\text{Link}]\!]_{\langle \mathfrak{M}, \mathfrak{M}' \rangle} = [\![\text{Link}]\!]_{\mathfrak{M}}$, so $b, c \in \text{dom}(\mathfrak{M})$. Now suppose $a \in \text{dom}(\mathfrak{M})$ and $\delta_{\mathfrak{M}}(a, b) = 0$. Therefore $e \in K$ (since $(f \uplus g)(K')$ does not touch $\text{dom}(\mathfrak{M})$) and $\delta_{\mathfrak{M}}(a, c) \leqslant 1$, so by definition of a match $c \in \text{Im}(f) \subseteq \text{Im}(f \uplus g)$.

  So we will show that $a \in \text{dom}(\mathfrak{M})$ and $\delta_{\mathfrak{M}}(a, b) = 0$. We know $b \in \text{dom}(\mathfrak{M})$ and $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(a, b) = 0$ so by lemma 2.2.1, $a \in \text{dom}(\mathfrak{M})$ and $\delta_{\mathfrak{M}}(a, b) = 0$.

$\square$

The next lemma complements the previous one by going *from match to pushout*: given a match from a pattern to a transition, when can one derive a rule from the pattern that can be applied to construct exactly that transition?

Lemma 2.5.3. If $(\mathfrak{A}, K)$ is a pure pattern, $m : (\mathfrak{A}, K) \to \langle \mathfrak{M}, \mathfrak{M}' \rangle$, $[\![\text{Has}]\!]_{\mathfrak{M}} = \text{dom}(\mathfrak{M})$, $[\![\text{Has}]\!]_{\mathfrak{M}'} = \text{dom}(\mathfrak{M}')$ and $\varkappa := (\text{Pre}(\mathfrak{A}), K \restriction \text{dom}(\text{Pre}(\mathfrak{A})), \text{Post}(\mathfrak{A}))$ then $\mathfrak{M} \twoheadrightarrow_\varkappa \mathfrak{M}'$.

Proof. Let $\mathfrak{S}$ be the reduct of $\mathfrak{A} \restriction m^{-1}(\text{dom}(\mathfrak{M}))$ to Θ°. Let $\mathfrak{S}'$ be the reduct of $\mathfrak{A} \restriction m^{-1}(\text{dom}(\mathfrak{M}'))$ to Θ°.

First, we show $\mathfrak{S} = \text{Pre}(\mathfrak{A})$. Since $\mathfrak{S}$ and $\text{Pre}(\mathfrak{A})$ are both Θ°-reducts of induced substructures of $\mathfrak{A}$, $\text{dom}(\mathfrak{S}) = \text{dom}(\text{Pre}(\mathfrak{A}))$ implies $\mathfrak{S} = \text{Pre}(\mathfrak{A})$.

- If $a \in \mathfrak{S}$ then $m(a) \in \text{dom}(\mathfrak{M})$ so $m(a) \in [\![\text{Has}]\!]_{\mathfrak{M}}$ by assumption on $[\![\text{Has}]\!]_{\mathfrak{M}}$. Since m is an embedding, $a \in \text{Pre}(\mathfrak{A})$.

- If $a \in \text{Pre}(\mathfrak{A})$ then $a \in \mathfrak{A} \restriction \{b\}$ for some $b \in [\![\text{Has}]\!]_{\mathfrak{A}}$; by assumption on $\mathfrak{M}, \mathfrak{M}'$, $m(b) \in \text{dom}(\mathfrak{M})$. Since m is an embedding, $m(a) \in \langle \mathfrak{M}, \mathfrak{M}' \rangle \restriction \{m(b)\}$, so $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(a, b) = 0$. So by lemma 2.2.1, $m(a) \in \text{dom}(\mathfrak{M})$ as well.

By the same reasoning $\mathfrak{S}' = \text{Post}(\mathfrak{A})$. So $\varkappa = (\mathfrak{S}, K \restriction \text{dom}(\mathfrak{S}), \mathfrak{S}')$.

We now show $\mathfrak{M} \twoheadrightarrow_{\varkappa'} \mathfrak{M}'$. Note that by lemma 2.5.1, $\langle \mathfrak{S}, \mathfrak{S}' \rangle$ is defined and equal to $\mathfrak{A}$.

Since m is an embedding and $\mathfrak{M}$ is a substructure of $\langle \mathfrak{M}, \mathfrak{M}' \rangle$, $\text{dom}(\mathfrak{S}) = m^{-1}(\text{dom}(\mathfrak{M}))$.

Let $f := m \restriction \text{dom}(\mathfrak{S})$. $f : \mathfrak{S} \to \mathfrak{M}$ is an embedding, by restriction of m. To show that $\mathcal{B}_{\mathfrak{M}}(f(\text{dom}(\mathfrak{S})), 0) \subseteq \text{Im}(f)$, it suffices to show that for every function symbol $h \in \textbf{Stat}$, if

$b \in \mathfrak{A} \upharpoonright m^{-1}(\mathfrak{M})$, $a \in \mathrm{dom}(\mathfrak{M})$ and $(a, b) \in [\![h]\!]_{\mathfrak{M}}$ then $a \in \mathrm{Im}(m)$. Since $\Theta^{\circ} \subseteq \Theta$ and by definition of $\langle \mathfrak{M}, \mathfrak{M}' \rangle$ we have $\delta_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(a, b) = 0$ and m a 0-embedding, so by definition of a 0-embedding, $a \in \mathrm{Im}(m)$.

To get $\mathcal{B}_{\mathfrak{M}}(\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}'), 1) \subseteq \mathrm{Im}(f)$, by the above it suffices to show that if $(a, b) \in [\![\mathsf{Link}]\!]_{\mathfrak{M}}$ and $f^{-1}(a)$ is defined and in $\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}')$, then $b \in \mathrm{Im}(f)$. Take such an a. Since m is an embedding, $a \in [\![\mathsf{Has}]\!]_{\mathfrak{S}}$ and $a \notin [\![\mathsf{Has}]\!]_{\mathfrak{S}'}$, so $a \in [\![\Delta]\!]_{\mathfrak{A}}$. Since $(\mathfrak{A}, K)$ is pure, $[\![\Delta]\!]_{\mathfrak{A}} \subseteq K$, so $b \in \mathrm{Im}(f)$.

We obtain the rule application condition $\mathcal{B}_{\mathfrak{M}}(f(K \upharpoonright \mathrm{dom}(\mathfrak{S})), 1) \subseteq \mathrm{Im}(f)$ from $\mathcal{B}_{\langle \mathfrak{M}, \mathfrak{M}' \rangle}(m(K), 1) \subseteq \mathrm{Im}(m)$ since $\mathfrak{M}$ is a substructure of $\langle \mathfrak{M}, \mathfrak{M}' \rangle$ and $\mathrm{Im}(m) \setminus \mathrm{Im}(f) \subseteq \mathrm{dom}(\mathfrak{M}') \setminus \mathrm{dom}(\mathfrak{M})$. So $f : (\mathfrak{S}, K \upharpoonright \mathrm{dom}(\mathfrak{S})) \to \mathfrak{M}$.

Again, since m is an embedding and $\mathfrak{M}'$ is a substructure of $\langle \mathfrak{M}, \mathfrak{M}' \rangle$, $\mathrm{dom}(\mathfrak{S}') = m^{-1}(\mathrm{dom}(\mathfrak{M}'))$. By definition of $\mathfrak{S}, \mathfrak{S}'$, $m(\mathrm{dom}(\mathfrak{S}') \setminus \mathrm{dom}(\mathfrak{S})) \cap \mathrm{dom}(\mathfrak{M}) = \emptyset$; so let $g : \mathrm{dom}(\mathfrak{S}') \setminus \mathrm{dom}(\mathfrak{S}) \to \mathcal{U} \setminus \mathrm{dom}(\mathfrak{M}) := x \mapsto m(x)$. As a restriction of m, g is injective. Trivially, $m = f \uplus g$.

Similarly, $m(\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}')) \cap \mathrm{dom}(\mathfrak{M}') = \emptyset$. Moreover, since m is a match and $\langle \mathfrak{M}, \mathfrak{M}' \rangle$ is supported (by assumption on $\mathrm{dom}(\mathfrak{M})$ and $\mathrm{dom}(\mathfrak{M}')$), $\mathrm{dom}(\mathfrak{M}') \setminus \mathrm{dom}(\mathfrak{M}) = \mathrm{Im}(g)$ and $\mathrm{dom}(\mathfrak{M}) \setminus \mathrm{dom}(\mathfrak{M}') = f(\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}'))$. Put together :

$$\mathrm{dom}(\mathfrak{M}') = \mathrm{dom}(\mathfrak{M}) \setminus f(\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}')) \uplus \mathrm{Im}(g) \tag{2.6}$$

(note that f is defined and equal to m on $\mathrm{dom}(\mathfrak{S}) \setminus \mathrm{dom}(\mathfrak{S}')$, and that g does not map to $\mathrm{dom}(\mathfrak{M})$).

For all $A \in \mathbf{Dyn}$:

$$
\begin{aligned}
[\![A]\!]_{\mathfrak{M}'} &= [\![A]\!]_{\mathfrak{M}} \setminus ([\![A]\!]_{\mathfrak{M}} \setminus [\![A]\!]_{\mathfrak{M}'}) \uplus ([\![A]\!]_{\mathfrak{M}'} \setminus [\![A]\!]_{\mathfrak{M}}) && \\
&= [\![A]\!]_{\mathfrak{M}} \setminus m([\![A]\!]_{\mathfrak{S}} \setminus [\![A]\!]_{\mathfrak{S}'}) \uplus m([\![A]\!]_{\mathfrak{S}'} \setminus [\![A]\!]_{\mathfrak{S}}) && \text{since } m : (\mathfrak{A}, K) \to \langle \mathfrak{M}, \mathfrak{M}' \rangle \\
&= [\![A]\!]_{\mathfrak{M}} \setminus f([\![A]\!]_{\mathfrak{S}} \setminus [\![A]\!]_{\mathfrak{S}'}) \uplus m([\![A]\!]_{\mathfrak{S}'} \setminus [\![A]\!]_{\mathfrak{S}}) && \text{by definition of } f
\end{aligned}
$$

The decomposition of $\mathrm{dom}(\mathfrak{M}')$ given by (2.6) yields:

$$[\![T]\!]_{\mathfrak{M}'} = [\![T]\!]_{\mathfrak{M}} \upharpoonright \mathrm{dom}(\mathfrak{M}') \cup m([\![T]\!]_{\mathfrak{S}'})$$

By $[\![\mathbf{Stat}]\!]_{\mathfrak{M}} \upharpoonright \mathcal{B}_{\mathfrak{M}}(\mathrm{dom}(\mathfrak{M}'), 0) = [\![\mathbf{Stat}]\!]_{\mathfrak{M}'} \upharpoonright \mathcal{B}_{\mathfrak{M}'}(\mathrm{dom}(\mathfrak{M}), 0)$ (by assumption on $\langle \mathfrak{M}, \mathfrak{M}' \rangle$) for the $[\![T]\!]_{\mathfrak{M}} \upharpoonright \mathrm{dom}(\mathfrak{M}')$ part, and by $g = m \upharpoonright \mathrm{dom}(\mathfrak{S}')$ and the fact that m is an embedding for the $m([\![T]\!]_{\mathfrak{S}'})$ part.

So $\mathfrak{M}'$ is a pushout of $(f, g, \mathfrak{S}, \mathfrak{S}', \mathfrak{M})$.

$\square$

To summarise the last two lemmas:

- If a rule applied to a state produces a pushout with witnesses f, g, then $f \uplus g$ is a match from (morally) the rule to the $\langle \text{state}, \text{pushout} \rangle$ transition.

- If there is a match m between a pure pattern and $\langle \mathfrak{M}, \mathfrak{M}' \rangle$, then the rule induced by the pattern, when applied to $\mathfrak{M}$, produces $\mathfrak{M}'$ as pushout. Moreover, the proof defines the pushout witnesses f and g such that $m = f \uplus g$.

It is easy to check that under the right conditions (pure pattern, strongly supported transitions), the composition of those lemmas in any order yields the identity function. Now that the 1-step operational and denotational views are connected, we move to traces.

2.6. Runs and executions

In this section we extend the operational and denotational views to traces (as seen for rule-based models in the introduction) and we give conditions that are sufficient to make both views coincide.

Operational view   When given a starting state, a *set of rules* induces a pushout-based trace we call an execution. Since rule application (by pushout) generates states that can be joined into transitions, a set of rules can now generate sequences of transitions from a given starting state:

Definition (Executions). If R is a set of rules, and $\mathfrak{M}_0, \mathfrak{M}_1, \ldots$ is a possibly infinite sequence such that $\mathfrak{M}_0 \twoheadrightarrow_R \mathfrak{M}_1 \twoheadrightarrow_R \mathfrak{M}_2 \twoheadrightarrow_R \ldots$, then $\langle \mathfrak{M}_0, \mathfrak{M}_1 \rangle, \langle \mathfrak{M}_1, \mathfrak{M}_2 \rangle, \ldots$ is an *execution* starting from $\mathfrak{M}_0$ induced by R (well-defined by lemma 2.3.1). The set of all executions starting from a state $\mathfrak{M} : \Theta^\circ$ induced by R is $\text{exec}_{\mathfrak{M}}(R)$ ∎

Denotational view   Instead of being iteratively generated, like executions, runs are a function of *sets of transitions*. They must make sure that their Pre and Post match pairwise. This denotational view paves the way for runs defined on transitions that are models of a formula.

Definition (Run). If M is a set of transitions, the *runs* $\rho(M)$ of M are the possibly infinite sequences $r := \mathfrak{A}_1, \mathfrak{A}_2, \ldots$ of supported transitions from M such that for all $\mathfrak{A}_i, \mathfrak{A}_{i+1} \in r$ we have $\text{Post}(\mathfrak{A}_i) = \text{Pre}(\mathfrak{A}_{i+1})$. ∎

Definition (Run from initial state). If $\mathfrak{S} : \Theta^\circ$ the $\mathfrak{S}$-runs $\rho_{\mathfrak{S}}(M)$ of M are the runs of the form $\mathfrak{A}_1, \ldots$ of M such that $\mathfrak{S} = \text{Pre}(\mathfrak{A}_1)$.

∎

Definition (From patterns to rules). If P is a set of patterns the set of rules $\text{rules}(P)$ is:

$$\text{rules}(P) := \{(\text{Pre}(\mathfrak{A}), K \upharpoonright \text{dom}(\text{Pre}(\mathfrak{A})), \text{Post}(\mathfrak{A})) \mid (\mathfrak{A}, K) \in P\}$$

∎

Contrast runs and executions: both are sequences of transitions; one is described operationally and the other denotationally. The task ahead is to find out when they are the same.

The next few technical lemmas prepare the theorem at the end of this section; in particular they show that when Has consistently encodes *presence* in a state, the operational and denotational trace generation processes maintain that consistency.

Lemma 2.6.1. If $\mathfrak{M} : \Theta^\circ$ is such that $\text{dom}(\mathfrak{M}) = [\![\text{Has}]\!]_{\mathfrak{M}}$, $\mathfrak{A}$ is strongly supported, and $\mathfrak{M}'$ is the pushout of $(f, g, \text{Pre}(\mathfrak{A}), \text{Post}(\mathfrak{A}), \mathfrak{M})$ then $\text{dom}(\mathfrak{M}') = [\![\text{Has}]\!]_{\mathfrak{M}'}$.

Proof. Since $\mathfrak{A}$ is strongly supported, every element in $\text{dom}(\mathfrak{A}) \upharpoonright [\![\text{Has}]\!]_{\mathfrak{A}}$ is in $[\![\text{Has}]\!]_{\mathfrak{A}}$ and so $[\![\text{Has}]\!]_{\text{Pre}(\mathfrak{A})} = \text{dom}(\text{Pre}(\mathfrak{A}))$. The same goes for $\text{Has}^\star$ and $\text{Post}(\mathfrak{A})$. The rest is trivial by definition of a pushout. □

Lemma 2.6.2. If $\mathfrak{A}$ is strongly supported, $\mathfrak{B}$ is supported and $\mathfrak{B}$ is an instance of $(\mathfrak{A}, K)$ for some K, then $\mathfrak{B}$ is strongly supported.

Proof. Suppose there are b_1, b_2 such that $b_1 \in \text{dom}(\mathfrak{B} \upharpoonright \{b_2\})$, $b_1 \notin [\![\text{Has}^\star]\!]_{\mathfrak{B}}$, $b_2 \in [\![\text{Has}^\star]\!]_{\mathfrak{B}}$.

As $\mathfrak{B}$ is supported, $b_1 \in [\![\Delta]\!]_{\mathfrak{B}}$, so there is $a_1 \in \text{dom}(\mathfrak{A})$ such that $f(a_1) = b_1$ for some embedding f of $\mathfrak{A}$ into $\mathfrak{B}$, and some $a_2 \in \text{dom}(\mathfrak{A})$ such that $a_2 \in \text{dom}(\mathfrak{A} \upharpoonright \{a_1\})$ and $a_2 \in [\![\text{Has}]\!]_{\mathfrak{A}}$, $f(a_2) = b_2$. By definition of an embedding, $a_1 \notin [\![\text{Has}^\star]\!]_{\mathfrak{A}}$, which contradicts the definition of strong support. □

This technical lemma simply ensures that $\text{Pre}(\cdot)$, $\text{Post}(\cdot)$ and $\langle \cdot, \cdot \rangle$ interact well together.

Lemma 2.6.3. *If* $\mathfrak{M}, \mathfrak{M}', \mathfrak{M}'' : \Theta^\circ$ *and* $\langle \mathfrak{M}, \mathfrak{M}' \rangle, \langle \mathfrak{M}', \mathfrak{M}'' \rangle$ *are defined, then* $\text{Post}(\langle \mathfrak{M}, \mathfrak{M}' \rangle) = \text{Pre}(\langle \mathfrak{M}', \mathfrak{M}'' \rangle)$.

Proof. If $a \in \text{dom}(\text{Post}(\langle \mathfrak{M}, \mathfrak{M}' \rangle))$ there is $b \in [\![\text{Has}]\!]_{\mathfrak{M}'}$ such that $a \in \text{dom}(\langle \mathfrak{M}, \mathfrak{M}' \rangle \restriction [\![\text{Has}]\!]_{\mathfrak{M}'})$. Since $[\![\textbf{Stat}]\!]_{\mathfrak{M}} \restriction \mathcal{B}_{\mathfrak{M}}(\text{dom}(\mathfrak{M}'), 0) = [\![\textbf{Stat}]\!]_{\mathfrak{M}'} \restriction \mathcal{B}_{\mathfrak{M}'}(\text{dom}(\mathfrak{M}), 0)$, $a \in \text{dom}(\mathfrak{M}')$. Similarly, if $a \in \text{Pre}(\langle \mathfrak{M}', \mathfrak{M}'' \rangle)$, $a \in \text{dom}(\mathfrak{M}')$. So $\text{dom}(\text{Post}(\langle \mathfrak{M}, \mathfrak{M}' \rangle)) = \text{dom}(\text{Pre}(\langle \mathfrak{M}', \mathfrak{M}'' \rangle))$. The equality of **Dyn** and **Stat** is by further unrolling of definitions. $\qquad \square$

Main theorem   We now prove the main theorem of this chapter: if patterns include all changes in their application conditions and both patterns and the starting state treat Has consistently, then the operational and denotational trace semantics of patterns are equal.

The proof is by double inclusion: from (pushout-based) executions to (match-based) runs, we show using lemma 2.5.2 that if a rule $\imath$ is induced by a pattern p, then a pushout based on $\imath$ induces a match for p. From (match-based) runs to (pushout-based) executions, we show using lemma 2.5.3 that the target of a match from a pattern p can be decomposed into the ingredients of a pushout starting from a rule induced by p.

Theorem 2.6.4. *If* P *is a set of pure and strongly supported patterns,* $\mathfrak{M}_0 : \Theta^\circ$ *and* $[\![\text{Has}]\!]_{\mathfrak{M}_0} = \text{dom}(\mathfrak{M}_0)$ *then* $\texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P)) = \rho_{\mathfrak{M}_0}(\text{Insts}(P))$

Proof.

$\texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P)) \subseteq \rho_{\mathfrak{M}_0}(\textbf{Insts}(P))$ Let $r := \langle \mathfrak{M}_0, \mathfrak{M}_1 \rangle, \ldots \in \texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P))$.

By lemma 2.6.1, $[\![\text{Has}]\!]_{\mathfrak{M}} = \text{dom}(\mathfrak{M})$ for any $\mathfrak{M}$ that is $\twoheadrightarrow_{\texttt{rules}(P)}$-accessible from $\mathfrak{M}_0$; this immediately implies $\langle \mathfrak{M}_i, \mathfrak{M}_{i+1} \rangle$ is supported for all $\langle \mathfrak{M}_i, \mathfrak{M}_{i+1} \rangle \in r$. By lemma 2.6.3 $\text{Post}(\langle \mathfrak{M}_i, \mathfrak{M}_{i+1} \rangle) = \text{Pre}(\langle \mathfrak{M}_{i+1}, \mathfrak{M}_{i+2} \rangle)$ whenever $\langle \mathfrak{M}_{i+1}, \mathfrak{M}_{i+2} \rangle \in r$.

Take any $\langle \mathfrak{M}_i, \mathfrak{M}_{i+1} \rangle \in r$. By definition of $\texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P))$ there are $(\mathfrak{A}, K) \in P$, $f : (\text{Pre}(\mathfrak{A}), K \restriction \text{dom}(\text{Pre}(\mathfrak{A}))) \to \mathfrak{M}_i$, and an injection $g : \text{dom}(\text{Post}(\mathfrak{A})) \setminus \text{dom}(\text{Pre}(\mathfrak{A})) \to \mathcal{U} \setminus \mathfrak{M}_i$ such that $\mathfrak{M}_{i+1}$ is the pushout of $(f, g, \text{Pre}(\mathfrak{A}), \text{Post}(\mathfrak{A}), \mathfrak{M}_i)$.

$\mathfrak{A}$ is supported so by lemma 2.5.1, $\langle \text{Pre}(\mathfrak{A}), \text{Post}(\mathfrak{A}) \rangle = \mathfrak{A}$. Note that $K \setminus (K \restriction \text{dom}(\text{Pre}(\mathfrak{A}))) \subseteq \text{dom}(\text{Post}(\mathfrak{A})) \setminus \text{dom}(\text{Pre}(\mathfrak{A}))$. Let $m := f \uplus g$

By lemma 2.5.2, $m : (\mathfrak{A}, K) \to \langle \mathfrak{M}_i, \mathfrak{M}_{i+1} \rangle$. So $r \in \rho_{\mathfrak{M}_0}(\text{Insts}(P))$.

$\rho_{\mathfrak{M}_0}(\text{Insts}(P)) \subseteq \texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P))$ Let $r := \mathfrak{B}_1, \ldots \in \rho_{\mathfrak{M}_0}(\text{Insts}(P))$.

By definition of $\rho_{\mathfrak{M}_0}(\text{Insts}(P))$, each $\mathfrak{B}_i$ is supported, $\text{Pre}(\mathfrak{B}_1) = \mathfrak{M}_0$ and $\text{Post}(\mathfrak{B}_i) = \text{Pre}(\mathfrak{B}_{i+1})$ whenever $B_{i+1} \in r$; and by lemma 2.5.1 $\langle \text{Pre}(\mathfrak{B}_i), \text{Post}(\mathfrak{B}_i) \rangle = \mathfrak{B}_i$.

By lemma 2.6.2, each $\mathfrak{B}_i$ is strongly supported, so $[\![\text{Has}]\!]_{\text{Pre}(B_i)} = \text{dom}(\text{Pre}(\mathfrak{B}_i))$ for every $\mathfrak{B}_i \in r$, and $[\![\text{Has}]\!]_{\text{Post}(B_i)} = \text{dom}(\text{Post}(\mathfrak{B}_i))$.

By definition of $\rho_{\mathfrak{M}_0}(\text{Insts}(P))$ there is some $(\mathfrak{A}, K) \in P$ and some $m : (\mathfrak{A}, K) \to \langle \text{Pre}(\mathfrak{B}_i), \text{Post}(\mathfrak{B}_i) \rangle$. By lemma 2.5.3, with $\imath := (\text{Pre}(\mathfrak{A}), K \restriction \text{dom}(\text{Pre}(\mathfrak{A})), \text{Post}(\mathfrak{A}))$ we get $\text{Pre}(\mathfrak{B}_i) \twoheadrightarrow_\imath \text{Post}(\mathfrak{B}_i)$; and $\imath \in \texttt{rules}(P)$ by definition so $\text{Pre}(\mathfrak{B}_i) \twoheadrightarrow_{\texttt{rules}(P)} \text{Post}(\mathfrak{B}_i)$.

So $r \in \texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P))$.

$\qquad \square$

In figure 2.9, we see the connection between the execution of the rules for a pattern set P and the runs on the instance of P. Note that $r = (\text{Pre}(\mathfrak{A}), K \restriction \text{dom}(\text{Pre}(\mathfrak{A})), \text{Post}(\mathfrak{A}))$, and similarly for r' and r''.

$$
\begin{array}{ccccccc}
r & r' & r'' & \text{induced by} & (\mathfrak{A}, K) & (\mathfrak{A}', K') & (\mathfrak{A}'', K'') \\
\end{array}
$$

Figure 2.9.: Left: an execution in $\mathtt{exec}_{\mathfrak{M}_0}(\mathtt{rules}(P))$. Right a run in $\rho_{\mathfrak{M}_0}(\mathsf{Insts}(P))$.

3. Execution of change-minimal specifications

In this chapter, we study formulas on transitions and states and further explore the connection between runs and executions.

Formulas and ↓   We start in section 3.1 by defining a change order on transitions which describes "how much" a transition does relative to another. Then in section 3.2 we define logic formulas on transitions, including the minimisation operator, which specifies the change-minimal models of a formula, as hinted at in the introduction. This defines FO[↓], which is first-order logic extended with a unary minimisation operator. After section 3.2, we derive a notion of *run* on formulas, and the rest of the chapter is about the runs associated with formulas and finding rules which execute in a way that corresponds with the runs.

Runs on models, runs on instances   First, we will introduce a parameterised function from formulas to patterns (the canonical patterns). We define those in section 3.3. For now consider that the canonical patterns provide a good "summary" of what the formula is about. Patterns can be instantiated, and the set of instances of the canonical patterns of a formula overapproximates the models of that formula.

Now there are three ways to look at runs based on ϕ:

1. We can consider the runs on the *models* of ϕ.

2. We can instantiate the canonical patterns of ϕ and consider the runs of those *instances*.

3. We can convert those canonical patterns into rules and consider the *execution* of those rules.

Thanks to the previous chapter, we know conditions where 2. and 3. coincide. The rest of the chapter after section 3.2 will be dedicated to building formulas where 1. and 2. coincide.

This will not be the end of the story. The resulting execution may require an infinite number of rules. The following chapter will address the issue.

Note on figures: in the sequel, we move to a more comfortable representation of transitions. Elements in figures belong to a single unary predicate from **Stat***, such as K or E. This predicate is shown within a large circle. Around the circle, identifers such as* $a, b, \ldots$ *are as before, and sans-serif letters represent belonging to a dynamic predicate of* **Dyn** *(left of the arrow) or of* **Dyn*** *(right of the arrow).*

3.1. Change order

Here we define an order $\trianglelefteq$ between transitions such that if $\mathfrak{A} \trianglelefteq \mathfrak{B}$, $\mathfrak{A}$ can be said to *do less than* $\mathfrak{B}$. Recall that change is the difference between the interpretation of **Dyn** and the interpretation of **Dyn***, so we consider the product subset relation to compare changes. With the order $\trianglelefteq$, we can then define minimisation, as we shall see in the next section.

Figure 3.1.: $\mathfrak{A} \trianglelefteq \mathfrak{B}$. $\mathsf{F} = \neg\mathsf{Link}(a, b) \wedge \mathsf{Link}^\star(a, b)$

As mentioned in the introduction, specifications on transitions are too lax when it comes to having any hope of generating traces. To motivate the next section, we informally use formulas. If we write $\neg\mathsf{Link}(a, b) \wedge \mathsf{Link}^\star(a, b)$ (to require that a and b become linked), we are not explicitly forbidding additional changes; such as the creation of a brand new element c.

However, if we wish to generate traces using this formula, we should make the additional assumption that *all observations have been accounted for*. If a c could appear exactly when a and b become linked, we would know it.

Take two transitions $\mathfrak{A}, \mathfrak{B}$ with the exact same precondition but where

- $\mathfrak{A}$ creates an a–b link.

- $\mathfrak{B}$ creates a–b link *and* creates c.

We would like a way to *compare* $\mathfrak{A}$ to $\mathfrak{B}$ and conclude that while $\mathfrak{A}$ is a good model of our formula when *all observations have been accounted for*, $\mathfrak{B}$ is not. $\mathfrak{A}$ and $\mathfrak{B}$ can be seen on figure 3.1.

Definition (Change order). If $\mathfrak{A}, \mathfrak{B} : \Theta$, we say that $\mathfrak{A} \trianglelefteq \mathfrak{B}$ whenever:

(1) $[\![\mathbf{Dyn}]\!]_\mathfrak{A} = [\![\mathbf{Dyn}]\!]_\mathfrak{B}$

(2) $[\![\mathbf{Stat}]\!]_\mathfrak{A} = [\![\mathbf{Stat}]\!]_\mathfrak{B} \upharpoonright \mathrm{dom}(\mathfrak{A})$

(3) $[\![\Delta\mathsf{A}]\!]_\mathfrak{A} \subseteq [\![\Delta\mathsf{A}]\!]_\mathfrak{B}$ for all $\mathsf{A} \in \mathbf{Dyn}$.

(4) $\mathrm{dom}(\mathfrak{A}) \subseteq \mathrm{dom}(\mathfrak{B})$

◾

Example. In the biology-inspired example of figure 3.1, the transition $\mathfrak{B}$ has two changes: the enzyme b and the kinase a become linked, and an enzyme c is created (changes are highlighted in green). $\mathfrak{A}$ is strictly below $\mathfrak{B}$ in the change order. It has the same precondition ($[\![\mathbf{Dyn}]\!]_\mathfrak{A} = [\![\mathbf{Dyn}]\!]_\mathfrak{B}$), but c is not in $\mathrm{dom}(\mathfrak{A})$ ($\mathrm{dom}(\mathfrak{A}) \subseteq \mathrm{dom}(\mathfrak{B})$). On the *remaining* elements, static interpretations are the same ($[\![\mathbf{Stat}]\!]_\mathfrak{A} = [\![\mathbf{Stat}]\!]_\mathfrak{B} \upharpoonright \mathrm{dom}(\mathfrak{A})$). Finally it contains strictly fewer changes ($[\![\Delta\mathsf{A}]\!]_\mathfrak{A} \subseteq [\![\Delta\mathsf{A}]\!]_\mathfrak{B}$ for all $\mathsf{A} \in \mathbf{Dyn}$); indeed the a–b link still appears, but no enzyme is created. Importantly, $\mathfrak{A}$ contains no additional change.

Formally, the "creation" of c displayed at the top of figure 3.1 is a representation of $c \notin [\![\mathsf{A}]\!]_\mathfrak{B}$ for all $\mathsf{A} \in \mathbf{Dyn}$ and $c \in [\![\mathsf{A}^\star]\!]_\mathfrak{B}$ for some $\mathsf{A}^\star \in \mathbf{Dyn}^\star$, typically $\mathsf{Has}^\star$. It is possible for c to not be in $\mathrm{dom}(\mathfrak{A})$ precisely because c does not appear anywhere in the interpretation of $\mathbf{Dyn}$.

In addition, note that support enforces some degree of coherence : since every element is in the interpretation of Has or of $\mathsf{Has}^\star$, no element is a "ghost", formally present, but that could be removed by going down the change order $\trianglelefteq$ even though the strictly smaller model would have no fewer changes than the larger one.

◾

Lemma 3.1.1. $\trianglelefteq$ is a partial order

Proof. Reflexivity and transitivity are trivial. For antisymmetry, assume $\mathfrak{A} \trianglelefteq \mathfrak{B}$ and $\mathfrak{B} \trianglelefteq \mathfrak{A}$. We immediately have $[\![\mathbf{Dyn}]\!]_{\mathfrak{A}} = [\![\mathbf{Dyn}]\!]_{\mathfrak{B}}$ as well as $[\![\mathbf{Stat}]\!]_{\mathfrak{A}} = [\![\mathbf{Stat}]\!]_{\mathfrak{B}}$ and $\mathrm{dom}(\mathfrak{A}) = \mathrm{dom}(\mathfrak{B})$. We just need $[\![\mathbf{Dyn}^{\star}]\!]_{\mathfrak{A}} = [\![\mathbf{Dyn}^{\star}]\!]_{\mathfrak{B}}$. For every $A \in \mathbf{Dyn}$, $[\![A^{\star}]\!]_{\mathfrak{A}} = [\![A^{\star}]\!]_{\mathfrak{B}}$ is given by $[\![\Delta A]\!]_{\mathfrak{A}} = [\![\Delta A]\!]_{\mathfrak{B}}$, $[\![A]\!]_{\mathfrak{A}} = [\![A]\!]_{\mathfrak{B}}$, and simple set-theoretic considerations. So $\mathfrak{A} = \mathfrak{B}$. $\square$

3.2. First-order logic on transitions, $\downarrow$

In this section we introduce FO[$\downarrow$], which is first-order logic augmented with a $\trianglelefteq$-minimising unary operator, $\downarrow$. Remember that $\trianglelefteq$ orders transitions according to how much change they contain.

We start with first-order logic with equality (FO) on the signature Θ. Here are some examples of FO formulas on transitions:

Examples.

$\neg \mathsf{Has}(x) \wedge \mathsf{Has}^{\star}(x)$	x appears
$\exists y.\ \mathsf{Link}(x, y) \wedge (\mathsf{On}(y) \leftrightarrow \mathsf{On}(x))$	some y is linked to x, x is On iff y is
$\forall x.\ \mathsf{Off}(x) \rightarrow \forall y.\ \neg \mathsf{Link}(x, y)$	Any Off element is free

∎

3.2.1. Minimisation ($\downarrow$)

To talk about the change order defined earlier, we introduce a unary operator $\downarrow$. Its intended meaning is precisely : *all observations have been accounted for*. Going back to figure 3.1, note that both $\mathfrak{A}$ and $\mathfrak{B}$ are models of F. But we would like only $\mathfrak{A}$ to be a model of $\downarrow\mathsf{F}$, not $\mathfrak{B}$.

We extend first-order logic with equality to FO[$\downarrow$], where $\downarrow$ is a unary logical constant whose semantics are defined by induction on the structure of ϕ.

Definition ($\downarrow$). If $\phi \in$ FO[$\downarrow$], $\mathfrak{A} : \Theta$ and $\mu : \langle \phi \rangle \rightarrow \mathrm{dom}(\mathfrak{A})$:

$$\mathfrak{A}, \mu \vDash \downarrow \phi$$

$$:=$$

$$\mathfrak{A}, \mu \vDash \phi \text{ and } \mathfrak{B}, \mu \nvDash \phi \text{ for all } \mathfrak{B} \lhd \mathfrak{A} \text{ such that } \mathrm{Im}(\mu) \subseteq \mathrm{dom}(\mathfrak{B})$$

∎

$\downarrow \phi$ is pronounced "min ϕ".

From now on (and until we explicitly move to 2nd order logic in section 7.1), "a formula" means "a Θ-formula over FO[$\downarrow$]".

3.2.2. FO[$\downarrow$] generalities

In this section we informally clarify logic notions. An FO[$\downarrow$] formula is built from terms, atoms, and the usual $\wedge, \vee, \neg, \forall, \exists$ and $\downarrow$. The satisfaction relation is defined on non-empty structures only.

Notation (Literal sets). If Σ is a signature, $\mathcal{L}_{\Sigma}$ is the set of Σ-literals, and $\mathcal{L}_{\Sigma}^{=}$ the union of all equality literals and $\mathcal{L}_{\Sigma}$. ∎

Notation (Free variables). If ϑ is a term, a set of terms, or a formula, $\langle\vartheta\rangle$ is the set of free variables in ϑ.

If we write $\phi(x_1, \ldots, x_n)$, we mean that $\langle\phi\rangle \subseteq \{x_1, \ldots x_n\}$. ∎

Notation (Quantifier closure). If ϕ is a formula, $\exists\phi$ is the existential closure of ϕ and $\forall\phi$ is its universal closure. ∎

Notation (Quantifier rank). If ϕ is a formula, $\|\phi\|$ is the quantifier rank of ϕ. ∎

Convention. We assume that variables are always fresh. We reuse names such as $x, y, \ldots$ etc to make reading easier, but name capture only happens through binders such as $\forall x$ or explicitly when variables are surfaced as arguments of formulas, such as in $\phi_1(x) \wedge \phi_2(x)$. ∎

Definition (Semantic bracket). We use the semantic bracket for formulas and terms, e.g. $[\![\phi]\!]$ is the set of pairs $(\mathfrak{A}, \mu)$ with $\mu : \langle\phi\rangle \to \mathrm{dom}(\mathfrak{A})$ such that $\mathfrak{A}, \mu \vDash \phi$. With $\langle\phi\rangle = x_1, \ldots x_n$ (we assume some ordering on variables), $[\![\phi]\!]_{\mathfrak{A}}$ is the set of n-tuples $(a_1, \ldots, a_n)$ of elements of $\mathrm{dom}(\mathfrak{A})$ such that $\mathfrak{A}, x_1 \mapsto a_1 \ldots x_n \mapsto a_n \vDash \phi$.

If t is a term and $\mu : \langle t \rangle \to \mathrm{dom}(\mathfrak{A})$, $[\![t]\!]_{\mathfrak{A},\mu}$ is the interpretation of t in $\mathfrak{A}, \mu$ defined by induction on the structure of t. ∎

Guards

Definition (Guards). If $x \neq y$, a *guard* $\alpha(x, y)$ is a binary atom (including "=") where both x and y occur. ∎

Note that unlike with the notation $\phi(x, y)$, $\alpha(x, y)$ means that x and y *occur* in the guard. The order is not specified, and the arguments of the underlying atom can be any terms over x and y.

Lemma 3.2.1. If $\mathfrak{A} : \Theta$, $t(x)$ is a term over a variable x, $u(y)$ is a term over a variable y, $A(t, u)$ is a binary atom over terms t and u, and $\mathfrak{A}, \mu \vDash A(t, u)$ then $\delta_{\mathfrak{A}}(\mu(x), \mu(y)) \leqslant 1$.

Proof. Let $a = [\![t]\!]_{\mathfrak{A},\mu}$ and $b = [\![u]\!]_{\mathfrak{A},\mu}$. By assumption $\delta_{\mathfrak{A}}(a, b) \leqslant 1$. By definition of $\delta_{\mathfrak{A}}$, $\delta_{\mathfrak{A}}(a, \mu(x)) = 0$ and $\delta_{\mathfrak{A}}(b, \mu(y)) = 0$. So $\delta_{\mathfrak{A}}(\mu(x), \mu(y)) \leqslant 1$. □

3.2.3. FO[↓] on transitions

A simple syntactic change can shift all the constraints expressed by a formula towards either the precondition or the postcondition of its models. Consider $\phi(x) := \mathsf{Active}(x) \vee \mathsf{Active}^\star(x)$. $\phi(x)$ requires that x is active at least at some point during the transition. However the starred version of ϕ, $\phi^\star(x) := \mathsf{Active}^\star(x) \vee \mathsf{Active}^\star(x)$, requires x to be active in the postcondition;

Definition (Starred and unstarred formulas). If ϕ is a formula, $\phi^\star$ is ϕ with every symbol from **Dyn** replaced with its counterpart from **Dyn***. Conversely, ϕ° is ϕ with every symbol from **Dyn*** replaced with its counterpart from **Dyn***. ∎

Formulas may mention symbols in **Stat**, **Dyn** and **Dyn***. A formula which does not mention symbols in **Dyn** is not talking about the precondition of its models. A formula which avoids **Dyn*** ignores the postcondition of its models.

If $\phi = \phi^\star$, then ϕ is a *post* formula. If $\phi = \phi^\circ$, it is a *pre* formula.

Convention. If ϕ is a pre, FO formula and $\mathfrak{S} : \Theta^\circ$, we freely talk about the semantics in ϕ in $\mathfrak{S}$ by taking ϕ to be a formula in the language of Θ°. ∎

3.2.4. Runs on formulas

Definition (Runs on formulas). Continuing an earlier definition (p.30), if ϕ is a formula, the runs of ϕ are the runs of $[\![\exists\phi]\!]$. We abuse notation and write $\rho(\phi)$ for $\rho([\![\exists\phi]\!])$. Similarly, if $\mathfrak{S} : \Theta^\circ$ we write $\rho_\mathfrak{S}(\phi)$ for $\rho_\mathfrak{S}([\![\exists\phi]\!])$. ∎

3.2.5. FO characterisation of support and strong support

Recall that a supported transition $\mathfrak{A}$ has its domain covered by $[\![\mathsf{Has}]\!]_\mathfrak{A} \cup [\![\mathsf{Has}^\star]\!]_\mathfrak{A}$. Supported transitions can be characterized by a universal sentence:

Definition (*Support*). The formula $Support := \forall x.\ \mathsf{Has}(x) \vee \mathsf{Has}^\star(x)$ characterizes supported transitions (no proof needed). ∎

Moreover, strongly supported transitions are supported and the interpretations of Has and $\mathsf{Has}^\star$ are 0-closed, which can also be characterized by a universal sentence:

Definition (*Support$_0$*). Let

$$\mathsf{Has}_0 := \forall x.\ \mathsf{Has}(x) \leftrightarrow \bigwedge_{h \in \mathbf{Stat}_{\mathrm{fun}}} \mathsf{Has}(h(x))$$

$$Support_0 := Support \wedge \mathsf{Has}_0 \wedge \mathsf{Has}_0^\star$$

The formula *Support$_0$* characterizes strongly supported structures (no proof needed).
If ϕ is any formula:

$$\phi_{S0} := \phi \wedge Support_0$$

■

3.3. Context-freeness & Retractability

A note on runs and executions We start with a note on the difference between ρ and `exec` to motivate the existence of each notion.

Suppose we start with a definition of *what happens* given a set of rewrite rules `rules` and a starting state $\mathfrak{M}$. This *execution semantics* $\mathtt{exec}_\mathfrak{M}(\mathtt{rules})$ is based on an operational pushout construction.

Now suppose there is some algorithm which, given a formula ϕ and a set of patterns P_ϕ induced by ϕ, yields the rules $\mathtt{rules}_{P_\phi}$. To evaluate the correctness of the rules, a reasonable criterion could be

> *The rewrite steps allowed by* $\mathtt{rules}_{P_\phi}$ *are the models of* ϕ*.*

But that would restrict ϕ too much. To widen the set of acceptable formulas, we choose to weaken the criterion to

> *Modulo a "nice" starting state, the rewrite steps allowed by* $\mathtt{rules}_{P_\phi}$ *are the models of* ϕ*.*

This naturally leads to defining, in addition to the transition semantics $[\![\phi]\!]$, the *run semantics* $\rho_\mathfrak{M}(\phi)$ that describe exactly the traces with starting state $\mathfrak{M}$ composed of models of ϕ.

So `exec` is an idealised execution method working on an input `rules`. To define the correctness of `exec` when given some rules induced by a formula ϕ, we defined a notion of execution *purely*

in terms of the formula. Now we have a framework to evaluate transformations from formulas to sets of rewrite rules.

Starting state aside, to connect $\rho(\phi)$ and $\texttt{exec}(\texttt{rules}_{P_\phi})$, we will go through the patterns P_ϕ. In the past chapter, we have shown when $\texttt{exec}(\texttt{rules}_{P_\phi}) = \rho(\text{Insts}(P_\phi))$. Now, we will show when $\rho(\text{Insts}(P_\phi)) = \rho(\phi)$.

Context-freeness and retractability   Equipped with a language for specifying transitions, we move in this section to the relationship between pattern matches and satisfaction of a formula. We define the dual notions of context-freeness and retractability. The first transports satisfiability forward along a match arrow, the other backward. We then introduce a set of deduction rules, $\vdash_{\!*}$, such that $\vdash_{\!*}$-provable formulas are both context-free and retractable. Those rules act as a "typing mechanism" for the formulas.

In chapter 2 we described how to relate executions on rules derived from a set of pattern P to the runs on *instances* of those patterns. In the previous section, we described runs on *models of formulas*. This section is a first step towards connecting runs on formulas and runs derived from pattern instances.

We need a brief definition, then we continue with the explanations:

Figure 3.2.: A canonical pattern. $\mathfrak{A}$ is in full outline. $\mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)$ is in dashed outline. The ball encompasses a radius of d around both the image of μ and elements in $[\![\Delta]\!]_{\mathfrak{A}}$.

Canonical patterns   To construct canonical patterns given a formula ϕ and $d \geqslant 0$, we take all models $\mathfrak{A}$ of ϕ that do not go beyond a radius $d + 1$ around:

- The interpretation of the variables that made $\mathfrak{A}$ satisfy ϕ.

- The changing elements of $\mathfrak{A}$, that is the elements $a \in [\![\Delta]\!]_{\mathfrak{A}}$.

For each of those, the application condition is the ball of radius d around those same elements. So as illustrated in figure 3.2, the application condition leaves a ring of radius 1 around the edges of the model.

Notation. If $\mathfrak{A} : \Theta$, μ is a function into $\text{dom}(\mathfrak{A})$ and $d \geqslant 0$,

$$\mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d) \coloneqq \mathcal{B}_{\mathfrak{A}}([\![\Delta]\!]_{\mathfrak{A}} \cup \text{Im}(\mu), d)$$

∎

Definition (Canonical patterns $\mathcal{P}(\phi, d)$). If ϕ is a formula and $d \geqslant 0$, the canonical patterns $\mathcal{P}(\phi, d)$ are the patterns of the form $(\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d))$ such that $\text{dom}(\mathfrak{A}) = \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d+1)$ and $\mathfrak{A}, \mu \vDash \phi$. ∎

Relationship between instances of canonical patterns and models Our goal is to find out when the runs associated with a formula are equal to the executions associated with some set of rules. Under the right conditions, we will show that for a formula ϕ and a starting state $\mathfrak{M}_0$, the right set of rules is

$$\texttt{rules}(\mathcal{P}(\phi, \|\phi\|))$$

and obtain the following equality:

$$\rho_{\mathfrak{M}_0}(\phi) = \texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(\mathcal{P}(\phi, \|\phi\|)))$$

From the previous chapter, we know conditions where, for P a set of patterns,

$$\rho_{\mathfrak{M}_0}(\mathsf{Insts}(P)) = \texttt{exec}_{\mathfrak{M}_0}(\texttt{rules}(P))$$

Let $P = \mathcal{P}(\phi, \|\phi\|)$. All we need to complete the equality is

$$\rho_{\mathfrak{M}}(\phi) = \rho_{\mathfrak{M}}(\mathsf{Insts}(\mathcal{P}(\phi, \|\phi\|)))$$

On the left-hand side, the runs are directly built from the models of the formula. On the right-hand side, we start by only looking at *some* models of the formula, then turn these models into patterns, then generate *all* possible instances of those patterns and build runs from those instances.

We will need the following two properties:

1. For the inclusion $\rho_{\mathfrak{M}}(\phi) \subseteq \rho_{\mathfrak{M}}(\mathsf{Insts}(\mathcal{P}(\phi, \|\phi\|)))$, we want any model of ϕ to be an instance of some pattern in $\mathcal{P}(\phi, \|\phi\|)$. The solution will be, given a transition $\mathfrak{A}$, to construct a pattern p and *deduce* $p \in \mathcal{P}(\phi, \|\phi\|)$ from the fact that $\mathfrak{A}$ satisfies ϕ.

2. For the inclusion $\rho_{\mathfrak{M}}(\phi) \supseteq \rho_{\mathfrak{M}}(\mathsf{Insts}(\mathcal{P}(\phi, \|\phi\|)))$, we want *enough* instances of $\mathcal{P}(\phi, \|\phi\|)$ to be models of ϕ. The solution will be, given a pattern p with an instance $\mathfrak{A}$ and some side-conditions, to *deduce* that $\mathfrak{A}$ satisfies ϕ from the fact that $p \in \mathcal{P}(\phi, \|\phi\|)$.

Recall the following general fact: purely existential formulas are preserved by embeddings, and purely universal ones are preserved by retraction along embeddings. See figure 3.3 for an illustration in our setting. We will define two parameterised sets: context-free and retractable formulas. If $(\mathfrak{A}, K)$ is a pattern and $m : (\mathfrak{A}, K) \to \mathfrak{B}$, $\mathfrak{A}$ may satisfy ϕ but $\mathfrak{B}$ may not (and vice-versa). Both directions are important: if a transition $\mathfrak{B}$ is part of a run in $\rho_{\mathfrak{M}}(\phi)$, it will be in a run of $\rho_{\mathfrak{M}}(\mathsf{Insts}(\mathcal{P}(\phi, \|\phi\|)))$ only if some pattern from $\mathcal{P}(\phi, \|\phi\|)$ can instantiate to it. If ϕ is *retractable*, it will be satisfied by a pattern of the right shape. In the other direction, a transition $\mathfrak{A}$ part of a run induced by pattern instances will be in $\rho_{\mathfrak{M}}(\phi)$ only if it satisfies ϕ. If ϕ is *context-free*, it will be satisfied by $\mathfrak{A}$.

In the next section, we will study the dual notions of context-freeness and retractability, defined on formulas, which correspond to maintaining satisfiability by going forwards or backwards along a match, modulo a proviso on the *precondition* of the match target. We will introduce a first system of rules, $\vdash_{\!*}$, which acts as a "typing" for formulas: $\vdash_{\!*}$-provable formulas are both context-free and retractable.

After that, there will still be some work left to do: the proviso which guarantees preservation of satisfaction needs to be true for all the elements of a run. We will introduce a restriction of the rules of $\vdash_{\!*}$, which we call $\vdash_{\!\rho}$. For $\vdash_{\!\rho}$-provable formulas, the truth or falsity of the proviso is invariant along runs. So we will show how, modulo some side conditions, a $\vdash_{\!\rho}$-provable formula ϕ induces the same runs as the instances of its canonical patterns.

$$\exists x.\ \mathsf{K}(x) \dashv \quad \textcircled{K} \; \textcircled{E}^{\,b} \quad \nvDash \forall x.\ \mathsf{E}(x)$$

$$\uparrow$$

$$\exists x.\ \mathsf{K}(x) \nvdash \quad \left(\textcircled{E}^{\,b},\{b\}\right) \quad \vDash \forall x.\ \mathsf{E}(x)$$

Figure 3.3.: A match and two formulas

3.3.1. Context-free, retractable formulas

Both notions of context-free and retractable use the following definition. The idea is that the target of a function f may need to verify some proviso *outside* of the image of f, on its *precondition*, but the image itself may be dispensed from verifying the proviso.

Definition (Validity in image complement). For any structure $\mathfrak{B}$, function f that maps to $\mathrm{dom}(\mathfrak{B})$ and any formula I, we write $\mathfrak{B} \vDash \forall (I \vee f)$ to mean that for all $\nu : \langle I \rangle \to \mathrm{dom}(\mathfrak{B})$, $\mathrm{Im}(\nu) \subseteq \mathrm{Im}(f)$ or $\mathfrak{B}, \nu \vDash I$. ∎

Given a radius d, a proviso I, two transitions $\mathfrak{A}, \mathfrak{B}$ and a particular application condition K. Suppose $m : (\mathfrak{A}, K) \to B$, and suppose $\mathfrak{B}$ satisfies I at least outside of the image of m. Then ϕ being (I, d)-context-free means that m transports ϕ-satisfaction forwards, and ϕ being (I, d)-retractable means that m transports it backwards.

Definition (Context-freeness). If ϕ is a formula, I is a pre formula, and $d \geqslant 0$, ϕ is (I, d)-*context-free* if for all $\mathfrak{A}, \mathfrak{B}$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d)) \to \mathfrak{B}$, if $\mathfrak{A}, \mu \vDash \phi$ and $\mathfrak{B} \vDash \forall(I \vee m)$ then $\mathfrak{B}, m \circ \mu \vDash \phi$. ∎

Any suitable d is a *context radius* of (ϕ, I).

Definition (Retractability). If ϕ is a formula, I is a pre formula I, $d \geqslant 0$, ϕ is (I, d)-*retractable* if for all $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \to \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d)) \to \mathfrak{B}$, if $\mathfrak{B}, m \circ \mu \vDash \phi$ and $\mathfrak{B} \vDash \forall(I \vee m)$ then $\mathfrak{A}, \mu \vDash \phi$. ∎

Any suitable d is a *retraction radius* of (ϕ, I).

We informally rephrase context-free and retractable in plain english. For some parameters (I, d) and μ, consider any "safe enough" match m from some $\mathfrak{A}$ to some $\mathfrak{B}$, namely:

a) m transports a d-radius ball containing $\mathrm{Im}(\mu)$ and the changes in $\mathfrak{A}$ to $\mathfrak{B}$

b) any tuple from $\mathfrak{B}$ fully outside of $\mathrm{Im}(m)$ satisfies the precondition I.

Then:

- ϕ is context-free if $\mathfrak{A}, \mu \vDash \phi$ implies $\mathfrak{B}, m \circ \mu \vDash \phi$

- ϕ is retractable if $\mathfrak{B}, m \circ \mu \vDash \phi$ implies $\mathfrak{A}, \mu \vDash \phi$

Example. Consider the formulas from figure 3.3. $\forall x.\ \mathsf{E}(x)$ is not context-free for any (I, d), but it is $(\top, 0)$-retractable. Conversely, $\exists x.\ \mathsf{K}(x)$ is not retractable for any (I, d) but it is $(\top, 0)$-context-free. ∎

Good clauses

We will soon introduce rules that can build context-free and retractable formulas. Before that we need to define a class of clauses: good clauses. A refinement of good clauses, called great clauses, will come into play later.

Definition (Subliterals of a clause). In a clause $\mathcal{C} = \bigvee_{1 \leqslant i \leqslant n} L_i$, let $\mathcal{L}_{\mathbf{Dyn}}(\mathcal{C})$ be the set of literals of $\mathcal{C}$ that use a symbol from $\mathbf{Dyn}$, and $\mathcal{L}^{\star}_{\mathbf{Dyn}}(\mathcal{C})$ be the set of literals of $\mathcal{C}$ that use a symbol from $\mathbf{Dyn}^{\star}$. ∎

Definition (Graph of a clause). Let $G_{\mathcal{C}}$ be the undirected graph with vertices $\langle \mathcal{C} \rangle$, and an edge $\{x, y\}$ iff there is $\neg\mathsf{Link}(t, u)$ or $\neg\mathsf{Link}^{\star}(t, u)$ in $\mathcal{C}$ such that $\{\langle t \rangle, \langle u \rangle\} = \{x, y\}$. ∎

Definition (Good clause). A clause $\mathcal{C}$ is *good* if $\mathcal{L}^{\star}_{\mathbf{Dyn}}(\mathcal{C})$ nonempty implies $G_{\mathcal{C}}$ connected. ∎

Example. The clause $\mathsf{A}(x) \vee \mathsf{Link}(x, y)$ is good. The clause $\mathsf{A}^{\star}(x) \vee \mathsf{Link}(x, y)$ is not good. The clause $\mathsf{A}^{\star}(x) \vee \neg\mathsf{Link}(x, y)$ is good. ∎

Rules $\vDash_{*}$ and first theorem

The rest of this subsection displays the rules $\vDash_{*}$ and states the main theorem of this section. Then, we list each pair of lemmas that proves the correctness of each rules and conclude with the main proof of the theorem.

In a sequent $\Gamma \vDash_{*} \phi$, Γ is a set of clauses $\mathcal{C}_1, \ldots, \mathcal{C}_k$ and ϕ is a formula. The rules $\vDash_{*}$ are an intermediate step towards the rules that actually guarantee correct execution. The asterisk $*$ is a little training wheel that will come off eventually.

$$\frac{}{\vDash_{*} L} \ \text{LITERAL}_{*} \qquad\qquad \mathcal{C} \text{ good} \ \frac{}{\mathcal{C}^{\circ} \vDash_{*} \forall \mathcal{C}} \ \text{CLAUSE}_{*}$$

$$\frac{\Gamma_1 \vDash_{*} \phi_1 \qquad \Gamma_2 \vDash_{*} \phi_2}{\Gamma_1, \Gamma_2 \vDash_{*} \phi_1 \vee \phi_2} \ \text{DISJ}_{*} \qquad\qquad \frac{\Gamma_1 \vDash_{*} \phi_1 \qquad \Gamma_2 \vDash_{*} \phi_2}{\Gamma_1, \Gamma_2 \vDash_{*} \phi_1 \wedge \phi_2} \ \text{CONJ}_{*}$$

$$\frac{\Gamma \vDash_{*} \phi}{\Gamma \vDash_{*} \forall x. \ \alpha(x, y) \rightarrow \phi} \ \forall\text{GUARD}_{*} \qquad\qquad \frac{\Gamma \vDash_{*} \phi}{\Gamma \vDash_{*} \exists x. \ \alpha(x, y) \wedge \phi} \ \exists\text{GUARD}_{*}$$

$$\frac{\Gamma \vDash_{*} \phi}{\Gamma \vDash_{*} \downarrow(\phi \wedge \textit{Support})} \ \text{CIRCUM}_{*}$$

Notation (Clause set to formula). If $\Gamma = \mathcal{C}_1, \ldots, \mathcal{C}_k$ is a set of clauses, $\wedge\Gamma = \bigwedge_{1 \leqslant i \leqslant k} \mathcal{C}_i$. ∎

The theorem to prove is:

Theorem 3.3.22. If $\Gamma \vDash_{*} \phi$ then ϕ is $(\wedge\Gamma \wedge \mathsf{Has}, \|\phi\|)$-context-free and retractable.

The proof will be found in section 3.3.3.

3.3.2. Building context-free and retractable formulas

Weakening

The weakening says that what works with a certain "protection radius" and certain proviso works with a larger radius and an even stronger proviso.

The proofs up to paragraph **Context-free clauses** are routine and can be found in the appendix, section A.1. Skip to page 42 for the more interesting part.

Lemma 3.3.1. If ϕ is (I, d)-context-free then ϕ is (I', d')-context-free for any $d' \geqslant d$, $I' \vDash I$.

Lemma 3.3.2. If ϕ is (I, d)-retractable then ϕ is (I', d')-retractable for any $d' \geqslant d$, $I' \vDash I$.

There is no need for an explicit weakening rule in $\vdash_*$ since DISJ$_*$ and CONJ$_*$ are multiplicative.

Quantifier-free formulas

Lemma 3.3.3. A quantifier-free formula ϕ is $(\top, 0)$-context-free.

Lemma 3.3.4. A quantifier-free formula ϕ is $(\top, 0)$-retractable.

Boolean composition

Lemma 3.3.5. If ϕ_1 is (I, d_1)-context-free (resp. retractable) and ϕ_2 is (I, d_2)-context-free (resp. retractable) then $\phi_1 \wedge \phi_2$ is $(I, \max(d_1, d_2))$-context-free (resp. retractable).

Lemma 3.3.6. If ϕ_1 is (I, d_1)-context-free (resp. retractable) and ϕ_2 is (I, d_2)-context-free (resp. retractable) then $\phi_1 \vee \phi_2$ is $(I, \max(d_1, d_2))$-context-free (resp. retractable).

Universal quantification

Lemma 3.3.7. If ϕ is (I, d)-context-free and $\alpha(x, y)$ is a guard, then $\forall x.\ \alpha(x, y) \to \phi$ is $(I, d + 1)$-context-free.

Lemma 3.3.8. If ϕ is (I, d)-retractable then for any x, $\forall x.\ \phi$ is (I, d)-retractable.

Existential quantification

Lemma 3.3.9. If ϕ is (I, d)-context-free then for any x, $\exists x.\ \phi$ is (I, d)-retractable.

Lemma 3.3.10. If ϕ is (I, d)-retractable and $\alpha(x, y)$ is a guard then $\exists x.\ \alpha(x, y) \wedge \phi$ is $(I, d + 1)$-retractable.

Context-free clauses

The main difficulty is the CLAUSE$_*$ rule for the context-free part. To transport a universal constraint *forward* along a morphism, we do need to some conditions to hold on the destination.

The following lemma simply rephrases the definition of $[\![\Delta]\!]$ as the *changes* of a transition by asserting that things outside of $[\![\Delta]\!]$ do not change indeed.

Lemma 3.3.11. If $L(t, u)$ is a pre literal ($t = u$ if L unary) and $\mathfrak{A}, \mu$ is such that $\delta_{\mathfrak{A}}([\![\langle t \rangle]\!]_{\mathfrak{A}, \mu}, [\![\Delta]\!]_{\mathfrak{A}}) > 0$ then $\mathfrak{A}, \mu \vDash L(t, u) \leftrightarrow L^\star(t, u)$.

Proof. Let S be the symbol used in $L(t, u)$. If $S \notin \mathbf{Dyn}$, $L(t, u) = L^\star(t, u)$ so we are done. Otherwise, since $\delta_{\mathfrak{A}}([\![\langle t \rangle]\!]_{\mathfrak{A}, \mu}, [\![\Delta]\!]_{\mathfrak{A}}) > 0$, we have $\delta_{\mathfrak{A}}([\![t]\!]_{\mathfrak{A}, \mu}, [\![\Delta]\!]_{\mathfrak{A}}) > 0$. Since $\mathfrak{A}, \mu \vDash \Delta L(t, u)$ implies $[\![t]\!]_{\mathfrak{A}, \mu} \in [\![\Delta]\!]_{\mathfrak{A}}$, we are done. $\qquad \square$

The following lemma gives a key property of good clauses. Consider $G_{\mathcal{C}}$. Assuming it is connected, $\mathcal{C}$ can only be false when the interpretation of $\langle \mathcal{C} \rangle$ fits in a small enough ball.

Lemma 3.3.12. If $\mathcal{C}$ is a clause such that $G_{\mathcal{C}}$ is connected and $\mathfrak{A}, \mu$ is such that $\mathfrak{A}, \mu \not\models L$ for every literal $L \in \mathcal{C}$ of the form $\neg\mathsf{Link}(t, u)$ or $\neg\mathsf{Link}^\star(t, u)$ then $d_{\mathfrak{A}}(a, b) < |\mu(\langle \mathcal{C} \rangle)|$ for every $a, b \in \mu(\langle \mathcal{C} \rangle)$.

Proof. Let G be the undirected graph with vertices $\mu(\langle \mathcal{C} \rangle)$ and an edge $\{a, b\}$ whenever $d_{\mathfrak{A}}(a, b) \leqslant 1$. By assumption, $G_{\mathcal{C}}$ is connected. We show that G is connected for any two $a, b \in G$ by induction on $\delta_{G_{\mathcal{C}}}(\mu^{-1}(a), \mu^{-1}(b))$.

If the distance is 0, then $a = b$ and we are done. Otherwise, there is y such that $\delta_{G_{\mathcal{C}}}(\mu^{-1}(a), y) = k$, $\delta_{G_{\mathcal{C}}}(y, \mu^{-1}(b)) = 1$ and by induction hypothesis, there is a path in G from a to $\mu(y)$. Now by definition of $G_{\mathcal{C}}$, $\mathfrak{A}, \mu \models \mathsf{Link}(t, u) \vee \mathsf{Link}^\star(t, u)$ with $\{\langle t \rangle, \langle u \rangle\} = \{y, m^{-1}(b)\}$. So $d_{\mathfrak{A}}(\mu(y), b) \leqslant 1$. So there is a path from a to b in G. So G is connected, so its diameter is $\leqslant |\mu(\langle \mathcal{C} \rangle)|$. $\quad\square$

This lemma gives a simple consequence of the definition of a match: in a match, changes in the target have a preimage in the changes of the source.

Lemma 3.3.13. If $m : (\mathfrak{A}, K) \to \mathfrak{B}$, then $[\![\Delta]\!]_{\mathfrak{B}} \subseteq m([\![\Delta]\!]_{\mathfrak{A}})$.

Proof. For every $A \in \mathbf{Dyn}$, every $\mathbf{b} \in [\![\Delta A]\!]_{\mathfrak{B}}$ and every $b \in \mathbf{b}$, $b \in [\![\Delta]\!]_{\mathfrak{B}}$. So $b \in m(\mathfrak{A})$ by definition of a match. So $m^{-1}(b)$ is well-defined. So by definition of an embedding, $m^{-1}(b) \in [\![\Delta A]\!]_{\mathfrak{A}}$. $\quad\square$

The following lemma explains the importance of good clauses and the presence of a proviso in the definition of context-free. By default, universal formulas are very much not context-free – the new context might betray the universal specification. Here we see that for a good clause, it is enough to check that the match target satisfies the *pre* version of the clause. Note that the context around the image of a match contains no changes, so it should be enough to know that the *pre* version of the clause is satisfied. The $|\langle \mathcal{C} \rangle|$ radius makes sure that the new context doesn't come too close to the changes (which are all defined by the domain of the match).

Lemma 3.3.14. If $\mathcal{C}$ is good and $\forall x. \, \mathcal{C}$ is a sentence, $\forall x. \, \mathcal{C}$ is $(\mathcal{C}^\circ, |\langle \mathcal{C} \rangle| - 1)$-context-free.

Proof. Let $d := |\langle \mathcal{C} \rangle|$. Take any $\mathfrak{A}, \mathfrak{B}, m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \emptyset}(d-1)) \to \mathfrak{B}$ such that $\mathfrak{A} \models \forall x. \, \mathcal{C}$, and for all $v : \langle \mathcal{C}^\circ \rangle \to \mathrm{dom}(\mathfrak{B})$, $\mathrm{Im}(v) \subseteq m(\mathfrak{A})$ or $\mathfrak{B}, v \models \mathcal{C}^\circ$. We want $\mathfrak{B} \models \forall x. \, \mathcal{C}$.

Let $\eta : \langle \mathcal{C} \rangle \to \mathrm{dom}(\mathfrak{B})$. If $\mathrm{Im}(\eta) \subseteq m(\mathfrak{A})$, let $\mu := m^{-1} \circ \eta$. By assumption, $\mathfrak{A}, \mu \models \mathcal{C}$. Since $\mathcal{C}$ is quantifier-free and m is an embedding, $\mathfrak{B}, \eta \models \mathcal{C}$ and we are done.

Otherwise there is $a \in \mathrm{Im}(\eta) \setminus m(\mathfrak{A})$. Moreover $\langle \mathcal{C} \rangle = \langle \mathcal{C}^\circ \rangle$ (trivially), so $\eta : \langle \mathcal{C}^\circ \rangle \to \mathrm{dom}(\mathfrak{B})$. So by assumption on $\mathfrak{A}, \mathfrak{B}, m$ and η, $\mathfrak{B}, \eta \models \mathcal{C}^\circ$.

If $\mathcal{L}_{\mathbf{Dyn}}^\star(\mathcal{C})$ is empty, then $\mathcal{C} = \mathcal{C}^\circ$ and we are done.

Otherwise, by definition of good, $G_{\mathcal{C}}$ is connected. We get the following inclusions by the lemmas indicated below:

$$\mathcal{B}_{\mathfrak{B}}([\![\Delta]\!]_{\mathfrak{B}}, d) \underset{\text{lemma 3.3.13}}{\subseteq} \mathcal{B}_{\mathfrak{B}}(m([\![\Delta]\!]_{\mathfrak{A}}), d) \underset{\text{lemma 2.4.5}}{\subseteq} m(\mathfrak{A})$$

Since $a \notin m(\mathfrak{A})$, $\delta_{\mathfrak{B}}(a, [\![\Delta]\!]_{\mathfrak{B}}) > d$.

If there is a literal L in $\mathcal{C}$ of the form $\neg\mathsf{Link}^\star(t, u)$ or $\neg\mathsf{Link}(t, u)$ such that $\mathfrak{B}, \eta \models L$, then $\mathfrak{B}, \eta \models \mathcal{C}$ and we are done.

If there is no such L then by lemma 3.3.12 $d_{\mathfrak{B}}(b, b') < |\mathrm{Im}(\eta)| \leqslant |\langle \mathcal{C} \rangle| = d$ for every $b, b' \in \mathrm{Im}(\eta)$. Since $\delta_{\mathfrak{B}}(a, [\![\Delta]\!]_{\mathfrak{B}}) > d$, we get $\delta_{\mathfrak{B}}(b, [\![\Delta]\!]_{\mathfrak{B}}) > 0$ for every $b \in \mathrm{Im}(\eta)$. Consider any L in $\mathcal{C}$ such that $\mathfrak{B}, \eta \models L^\circ$ (by assumption, there is at least one).

If L is not a post literal, then $L^\circ = L$ so $\mathfrak{B}, \eta \vDash \mathcal{C}$.

Otherwise, since $G_\mathcal{C}$ is connected, we apply lemma 3.3.11 and get $\mathfrak{B}, \eta \vDash L$, so $\mathfrak{B}, \eta \vDash \mathcal{C}$.

□

Example. In figure 3.4, we illustrate good clauses with an example of a match between two structures $\mathfrak{A}$ and $\mathfrak{B}$. $\mathfrak{A}$ satisfies $F_1 := \forall x.\ \mathsf{On}(x) \vee \mathsf{Active}^\star(x)$. We explain why $\mathfrak{B}$ does as well.

By CLAUSE$_*$, $\mathsf{On}(x) \vee \mathsf{Active}(x) \not\vDash_* F_1$. So by lemma 3.3.14, F_1 is $(\mathsf{On} \vee \mathsf{Active}, 0)$-context-free. Note that $\mathfrak{A}, x \mapsto a \nvDash \mathsf{On}(x) \vee \mathsf{Active}(x)$, yet $\mathfrak{A} \vDash F_1$. The elements of $\mathrm{dom}(\mathfrak{B})$ that are not in $\mathrm{dom}(\mathfrak{A})$ *do* satisfy $\mathsf{On}(x) \vee \mathsf{Active}(x)$. Crucially, b is not involved in a change, and thus is still Active in the postcondition. Thus, we get $\mathfrak{B} \vDash F_1$. Generally speaking, in the rule CLAUSE$_*$, having $\mathcal{C}$ be good ensures that, if a tuple satisfies $\mathcal{C}^\circ$ and is not involved in a change, then it satisfies $\mathcal{C}$. ∎


Figure 3.4.: $\mathfrak{A}$ and $\mathfrak{B}$ satisfy $F_1 := \forall x.\ \mathsf{On}(x) \vee \mathsf{Active}^\star(x)$. Changes highlighted in green.

Example. In figure 3.5, we see a more complex example. The transition $\mathfrak{A}$ vacuously satisfies the formula $F_2 := \forall xy.\ \mathsf{Link}(x, y) \rightarrow \mathsf{E}(x) \rightarrow \mathsf{On}^\star(x)$ (with $\mathsf{E} \in \mathbf{Stat}$). The match $\mathrm{id}_\mathfrak{A} : (\mathfrak{A}, \mathfrak{B}_\mathfrak{A}^{\Delta, \emptyset}(1)) \rightarrow \mathfrak{B}$ ensures $\mathfrak{B} \vDash F_2$ provided $\mathfrak{B}, \mu \vDash \mathsf{Link}(x, y) \rightarrow \mathsf{E}(x) \rightarrow \mathsf{On}(x)$ whenever $\mathrm{Im}(\mu) \not\subseteq \mathrm{Im}(\mathrm{id}_\mathfrak{A})$. a gains an edge (through the match, *not* during the transition) and is allowed to since it is not in $[\![\Delta]\!]_\mathfrak{A}$. The proviso only requires $a \in [\![\mathsf{On}]\!]_\mathfrak{A}$: since a is not involved in a change, that implies $a \in [\![\mathsf{On}^\star]\!]_\mathfrak{A}$. However b is involved in a change, and thus cannot become linked through the match. ∎

Figure 3.5.: $\mathfrak{A}$ and $\mathfrak{B}$ satisfy $F_2 := \forall xy.\ \mathsf{Link}(x, y) \rightarrow \mathsf{E}(x) \rightarrow \mathsf{On}^\star(x)$

Minimisation

We are in FO[↓], not just FO, so we must study how to build context-free and retractable minimised formulas. Here is the main idea, informally: suppose ϕ is both context-free and retractable. Is $\downarrow \phi$ context-free?

Take some $\mathfrak{m} : (\mathfrak{A}, K) \rightarrow \mathfrak{B}$. We know $\mathfrak{A} \vDash \downarrow \phi$, so (modulo a proviso and some facts about K) $\mathfrak{B} \vDash \phi$. The hard question is: is $\mathfrak{B}$ minimal among models ϕ?

Suppose it is not, that is: suppose there is $\mathfrak{B}' \lhd \mathfrak{B}$ such that $\mathfrak{B}' \vDash \phi$. Lemma 3.3.18 will show that under the right circumstances, we can build a transition $\mathfrak{A}' \lhd \mathfrak{A}$ such that there is a "nice" match from $\mathfrak{A}'$ to $\mathfrak{B}'$. Now, since we supposed that ϕ was not only context-free but also retractable, $\mathfrak{B}' \vDash \phi$ yields $\mathfrak{A}' \vDash \phi$, which contradicts the minimality of $\mathfrak{A}$ among models of ϕ. That reasoning made precise can be found in lemma 3.3.19.

There is also a technical requirement that we only consider supported models, due to the definition of $\unlhd$. An alternative design (but with fairly wide-ranging consequences) would be to remove the $\mathrm{dom}(\mathfrak{A}) \subseteq \mathrm{dom}(\mathfrak{B})$ in the definition of $\mathfrak{A} \unlhd \mathfrak{B}$.

The same reasoning works if one asks whether $\downarrow\phi$ is retractable, as shown in lemmas 3.3.20 and 3.3.21.

The following lemma simply rests on the fact that the change order cannot touch preconditions, so pre formulas are not affected by it.

Lemma 3.3.15. If ψ is a pre, quantifier-free formula, $\mathfrak{A} \unlhd \mathfrak{B}$ and $\nu : \langle\psi\rangle \to \mathrm{dom}(\mathfrak{A})$, then $\mathfrak{A}, \nu \vDash \psi$ iff $\mathfrak{B}, \nu \vDash \psi$.

Proof. Trivial induction on the structure of ψ, base cases for atoms use the fact that, by definition of $\unlhd$, $[\![\mathbf{Dyn}]\!]_{\mathfrak{A}} = [\![\mathbf{Dyn}]\!]_{\mathfrak{B}}$ and $[\![\mathbf{Stat}]\!]_{\mathfrak{A}} = [\![\mathbf{Stat}]\!]_{\mathfrak{B}} \restriction \mathrm{dom}(\mathfrak{A})$. □

This lemma gives an immediate consequence of the fact that $\mathsf{Has}(x) \vee \mathsf{Has}^\star(x)$ is good. This is an ad hoc proof of something $\vDash_{\bar{*}}$ is designed to prove, but we need it before we can prove all of $\vDash_{\bar{*}}$ correct.

Lemma 3.3.16. If ϕ is (I, d)-context-free (resp. (I, d)-retractable), then $\phi \wedge \mathit{Support}$ is $(I \wedge \mathsf{Has}, d)$-context-free (resp. $(I \wedge \mathsf{Has}, d)$-retractable).

Proof. For retractable, lemmas 3.3.4, 3.3.8 and 3.3.2 immediately give the conclusion. For context-free, we note that $\mathsf{Has}(x) \vee \mathsf{Has}^\star(x)$ is good and $|\langle\mathsf{Has}(x) \vee \mathsf{Has}^\star(x)\rangle| = 1$. For context-free, by lemma 3.3.14, $\mathit{Support}$ is $(\mathsf{Has}, 0)$-context-free. By lemmas 3.3.5 and 3.3.1, we are done. □

This simple lemma gives another way to state the effect of application conditions on 0-embeddings.

Lemma 3.3.17. If $m : (\mathfrak{A}, K) \to \mathfrak{B}$, $A \in \mathbf{Dyn}$, $b \in [\![A]\!]_{\mathfrak{B}}$, and $b \in m(K)$ for some $b \in \mathbf{b}$ then there is $\mathbf{a} \in [\![A]\!]_{\mathfrak{A}}$ such that $m(\mathbf{a}) = \mathbf{b}$.

Proof. If $b \in m(K)$ then by definition of a match and $d_{\mathfrak{B}}(b, b') \leqslant 1$ for every $b' \in \mathbf{b}$, $m^{-1}(\mathbf{b})$ is well-defined. So by definition of an embedding, $m^{-1}(\mathbf{b}) \in [\![A]\!]_{\mathfrak{A}}$. □

We now move to the main lemmas for proving that minimisation does not impact context-freeness and retractability. Illustrations are included.

Context-freeness of $\downarrow\phi$ This next pair of lemmas shows how to obtain context-freeness of $\downarrow\phi$ from both retractability and context-freeness of ϕ. Figure 3.6 illustrates the first one.

Lemma 3.3.18. If $\mathfrak{B} : \Theta$ is supported, $\mathfrak{B}' \lhd \mathfrak{B}$, $\mathfrak{A} : \Theta$, $\mu : X \to \mathrm{dom}(\mathfrak{A})$ for X some set of variables, and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)) \to \mathfrak{B}$ such that $m(\mu(X)) \subseteq \mathrm{dom}(\mathfrak{B}')$, then there are $\mathfrak{A}'$ and $m' : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)) \to \mathfrak{B}'$ such that $\mathfrak{A}' \lhd \mathfrak{A}$ and $m' = m \restriction m^{-1}(\mathfrak{B}')$.

Figure 3.6.: Illustration of the statement of lemma 3.3.18

Proof. We start by defining $\mathfrak{A}'$.

$$\operatorname{dom}(\mathfrak{A}') := \operatorname{dom}(\mathfrak{A}) \setminus m^{-1}(\operatorname{dom}(\mathfrak{B}) \setminus \operatorname{dom}(\mathfrak{B}'))$$

$$\llbracket A \rrbracket_{\mathfrak{A}'} := \llbracket A \rrbracket_{\mathfrak{A}} \qquad\qquad \text{for all } A \in \textbf{Dyn}$$

$$\llbracket T \rrbracket_{\mathfrak{A}'} := \llbracket T \rrbracket_{\mathfrak{A}} \restriction \operatorname{dom}(\mathfrak{A}') \qquad\qquad \text{for all } T \in \textbf{Stat}$$

$$\llbracket \Delta A \rrbracket_{\mathfrak{A}'} := \llbracket \Delta A \rrbracket_{\mathfrak{A}} \setminus m^{-1}(\llbracket \Delta A \rrbracket_{\mathfrak{B}} \setminus \llbracket \Delta A \rrbracket_{\mathfrak{B}'}) \qquad\qquad \text{for all } A \in \textbf{Dyn}$$

We show that $\mathfrak{A}'$ is well-defined, i.e. that the above definitions induce a single structure and that the domain covers the elements mentioned in the relations.

- The definitions for every $A \in \textbf{Dyn}$ of $\llbracket A \rrbracket_{\mathfrak{A}'}$ and $\llbracket \Delta A \rrbracket_{\mathfrak{A}'}$ induce

$$\llbracket A^\star \rrbracket_{\mathfrak{A}'} = (\llbracket \Delta A \rrbracket_{\mathfrak{A}'} \setminus \llbracket A \rrbracket_{\mathfrak{A}'}) \cup (\llbracket A \rrbracket_{\mathfrak{A}'} \setminus \llbracket \Delta A \rrbracket_{\mathfrak{A}'})$$

- Suppose there is $\mathfrak{a} \in \llbracket A \rrbracket_{\mathfrak{A}'}$. $\mathfrak{a} \in \llbracket A \rrbracket_{\mathfrak{A}}$ by definition of $\trianglelefteq$. $m(\mathfrak{a}) \in \llbracket A \rrbracket_{\mathfrak{B}}$ by definition of an embedding. $m(\mathfrak{a}) \in \llbracket A \rrbracket_{\mathfrak{B}'}$ by definition of $\llbracket A \rrbracket_{\mathfrak{B}'}$. So for every $a \in \mathfrak{a}$, $m(a) \in \operatorname{dom}(\mathfrak{B}')$, and thus by definiton of $\operatorname{dom}(\mathfrak{A}')$, $\mathfrak{a} \in \operatorname{dom}(\mathfrak{A}')$.

- Suppose there is $\mathfrak{a} \in \llbracket A^\star \rrbracket_{\mathfrak{A}'}$. Wlog $\mathfrak{a} \notin \llbracket A \rrbracket_{\mathfrak{A}'}$ (see previous case), so $\mathfrak{a} \in \llbracket \Delta A \rrbracket_{\mathfrak{A}'}$. $\mathfrak{a} \in \llbracket \Delta A \rrbracket_{\mathfrak{A}}$ by definition of $\trianglelefteq$. $m(\mathfrak{a}) \in \llbracket \Delta A \rrbracket_{\mathfrak{B}}$ by definition of an embedding. So, by definition of $\llbracket \Delta A \rrbracket_{\mathfrak{A}'}$, $m(\mathfrak{a}) \in \llbracket \Delta A \rrbracket_{\mathfrak{B}'}$. So $a \in \operatorname{dom}(\mathfrak{B}')$ for every $a \in \mathfrak{a}$, and thus by definition of $\operatorname{dom}(\mathfrak{A}')$, $\mathfrak{a} \in \operatorname{dom}(\mathfrak{A}')$.

- The case $\mathfrak{a} \in \llbracket T \rrbracket_{\mathfrak{B}'}$ is trivial.

Now we show $\mathfrak{A}' \lhd \mathfrak{A}$. By definition $\mathfrak{A}' \trianglelefteq \mathfrak{A}$. Moreover $\mathfrak{B}' \lhd \mathfrak{B}$, so $\llbracket \Delta A \rrbracket_{\mathfrak{B}'} \subset \llbracket \Delta A \rrbracket_{\mathfrak{B}}$ for some $A \in \textbf{Dyn}$ or $\operatorname{dom}(\mathfrak{B}') \subset \operatorname{dom}(\mathfrak{B})$.

- If there is $b \in \operatorname{dom}(\mathfrak{B}) \setminus \operatorname{dom}(\mathfrak{B}')$, by definition of $\lhd$ we have $b \notin \llbracket \text{Has} \rrbracket_{\mathfrak{B}}$. Since $\mathfrak{B}$ is supported, $b \in \llbracket \Delta \text{Has} \rrbracket_{\mathfrak{B}}$, and so by lemma 3.3.13 there is $a \in \llbracket \Delta \text{Has} \rrbracket_{\mathfrak{A}}$ such that $m(a) = b$. So by definition of $\operatorname{dom}(\mathfrak{A}')$, $a \notin \operatorname{dom}(\mathfrak{A}')$ and thus $\mathfrak{A}' \lhd \mathfrak{A}$.

- If there is $b \in \llbracket \Delta A \rrbracket_{\mathfrak{B}} \setminus \llbracket \Delta A \rrbracket_{\mathfrak{B}'}$ for some $A \in \textbf{Dyn}$, by lemma 3.3.13 there is $a \in \llbracket \Delta A \rrbracket_{\mathfrak{A}}$ such that $m(a) = b$. By definition of $\llbracket \Delta A \rrbracket_{\mathfrak{A}'}$, $a \notin \llbracket \Delta A \rrbracket_{\mathfrak{A}'}$ and thus $\mathfrak{A}' \lhd \mathfrak{A}$.

With $m' = m \restriction m^{-1}(\mathfrak{B}')$, we need $m' : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)) \to \mathfrak{B}'$, i.e. that $m^{-1}(\mathfrak{B}') = \operatorname{dom}(\mathfrak{A}')$, $\mu(X) \subseteq \operatorname{dom}(\mathfrak{A}')$, that m is an embedding, that $\mathcal{B}_{\mathfrak{B}'}(m'(\mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)), 1) \subseteq m'(\mathfrak{A}')$ and that $\llbracket \Delta \rrbracket_{\mathfrak{B}'} \subseteq m'(\mathfrak{A}')$.

- We show $m^{-1}(\mathfrak{B}') = \text{dom}(\mathfrak{A}')$. If $a \in m^{-1}(\mathfrak{B}')$ then $m(a) \in m(\mathfrak{A})$, so by definition of $\text{dom}(\mathfrak{A}')$, $a \in \text{dom}(\mathfrak{A}')$. In the other direction, if $a \in \text{dom}(\mathfrak{A}')$ then $a \in \text{dom}(\mathfrak{A})$ as well by definition of $\trianglelefteq$, so $a \in m^{-1}(\mathfrak{B})$. Thus by definition of $\text{dom}(\mathfrak{A}')$, $a \in m^{-1}(\mathfrak{B}')$.

- $\mu(X) \subseteq \text{dom}(\mathfrak{A}')$ follows from $m^{-1}(\mathfrak{B}') \subseteq \text{dom}(\mathfrak{A}')$ and the assumption $m(\mu(X)) \subseteq \text{dom}(\mathfrak{B}')$.

- Since m is an embedding and by definition of $\mathfrak{A}'$, $m' : \mathfrak{A}' \to \mathfrak{B}'$ is an embedding. To show that it is a 0-embedding, it suffices to show that for all $a \in \text{dom}(\mathfrak{B}')$, function symbol $h \in \mathbf{Stat}$, $b \in \text{Im}(m')$ such that $(a,b) \in [\![h]\!]_{\mathfrak{B}'}$, we have $a \in \text{Im}(m)$. By definition of $\mathfrak{B}'$, that premise implies $\delta_{\mathfrak{B}}(a,b) = 0$, so, since m is a 0-embedding and m' is a restriction of m, we get $a \in \text{Im}(m)$.

- Let $b \in \text{dom}(\mathfrak{B}')$ be at distance 1 (in $\mathfrak{B}'$) from $m'(a)$ with $a \in \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)$. Since $[\![\Delta]\!]_{\mathfrak{A}'} \subseteq [\![\Delta]\!]_{\mathfrak{A}}$ and by definition of m, $b \in m(\mathfrak{A})$. Suppose $b \notin \text{dom}(\mathfrak{A}')$: by definition of $\text{dom}(\mathfrak{A}')$, $b \notin \text{dom}(\mathfrak{B}')$. Absurd.

- Finally by $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq [\![\Delta]\!]_{\mathfrak{B}}$ and lemma 3.3.13, $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m([\![\Delta]\!]_{\mathfrak{A}})$. It is easy to see that $[\![\Delta]\!]_{\mathfrak{A}'} = [\![\Delta]\!]_{\mathfrak{A}} \setminus m^{-1}([\![\Delta]\!]_{\mathfrak{B}} \setminus [\![\Delta]\!]_{\mathfrak{B}'})$, so $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m([\![\Delta]\!]_{\mathfrak{A}'})$. Since $\text{dom}(\mathfrak{A}') \subseteq m^{-1}(\mathfrak{B}')$, $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m'([\![\Delta]\!]_{\mathfrak{A}'})$.

$\square$

Recall the reasoning at the beginning of the **minimisation** paragraph above and keep in mind figure 3.6, which summarises the previous lemma. We use the contrapositive of the previous lemma to prove the next one.

Lemma 3.3.19. *If ϕ is (I, d_c)-context-free and (I, d_r)-retractable, then $\downarrow(\phi \wedge Support)$ is $(I \wedge Has, \max(d_c, d_r))$-context-free.*

Proof. Let d_c be any context radius and d_r be any retraction radius for ϕ, I, and let $d := \max(d_c, d_r)$. Consider $\mathfrak{A}, \mathfrak{B}, \mu : \langle\phi\rangle \to \text{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}$ such that $\mathfrak{A}, \mu \vDash_{\downarrow} (\phi \wedge Support)$ and $\mathfrak{B} \vDash \forall((I \wedge Has) \vee m)$. We want $\mathfrak{B}, m \circ \mu \vDash_{\downarrow} (\phi \wedge Support)$. Assume not. By lemma 3.3.16, $\phi \wedge Support$ is $(I \wedge Has, d_c)$-context-free. Since $d \geqslant d_c$, by lemma 3.3.1 $\mathfrak{B}, m \circ \mu \vDash \phi \wedge Support$. So there is $\mathfrak{B}' \lhd \mathfrak{B}$ such that $\mathfrak{B}', m \circ \mu \vDash \phi \wedge Support$.

Since $\mathfrak{B}$ is supported and $m \circ \mu$ maps to $\text{dom}(\mathfrak{B}')$, by lemma 3.3.18, there is $\mathfrak{A}' \lhd \mathfrak{A}$ and $m' : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)) \to \mathfrak{B}'$ such that $m \restriction m^{-1}(\mathfrak{B}') = m'$. We show $\mathfrak{A}', \mu \vDash \phi \wedge Support$, which contradicts $\mathfrak{A}, \mu \vDash_{\downarrow} (\phi \wedge Support)$:

Take any $\nu : \langle I \wedge Has \rangle \to \text{dom}(\mathfrak{B}')$. Since $I \wedge Has$ is a pre, quantifier free formula and $\mathfrak{B}' \trianglelefteq \mathfrak{B}$, by lemma 3.3.15 $\mathfrak{B}, \nu \vDash I \wedge Has$ implies $\mathfrak{B}', \nu \vDash I \wedge Has$. Otherwise, by assumption $\text{Im}(\nu) \subseteq \text{Im}(m)$. Since $\text{Im}(\nu) \subseteq \text{dom}(\mathfrak{B}')$, $\text{Im}(\nu) \subseteq \text{Im}(m \restriction m^{-1}(\mathfrak{B}')) = m'(\mathfrak{A}')$. So $\mathfrak{B}' \vDash \forall((I \wedge Has) \vee m')$. By lemma 3.3.16, $\phi \wedge Support$ is $(I \wedge Has, d_r)$-retractable. Since $d \geqslant d_r$, by lemma 3.3.2 we get $\mathfrak{A}', \mu \vDash \phi \wedge Support$.

$\square$

Retractability of $\downarrow\phi$ This next pair of lemmas shows how to obtain retractability of $\downarrow\phi$ from both retractability and context-freeness of ϕ. Figure 3.7 illustrates the first one.

Lemma 3.3.20. *If $\mathfrak{A} : \Theta$ is supported, $\mathfrak{A}' \lhd \mathfrak{A}$, $\mathfrak{B} : \Theta$, $\mu : X \to \text{dom}(\mathfrak{A}')$ for X some set of variables, and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}$, then there are $\mathfrak{B}'$ and $m' : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)) \to \mathfrak{B}'$ such that $\mathfrak{B}' \lhd \mathfrak{B}$ and $m' = m \restriction m^{-1}(\mathfrak{B}')$.*

Figure 3.7.: Illustration of the statement of lemma 3.3.20

Proof. We start by defining $\mathfrak{B}'$.

$$\mathrm{dom}(\mathfrak{B}') := \mathrm{dom}(\mathfrak{B}) \setminus m(\mathrm{dom}(\mathfrak{A}) \setminus \mathrm{dom}(\mathfrak{A}'))$$

$$[\![A]\!]_{\mathfrak{B}'} := [\![A]\!]_{\mathfrak{B}} \qquad\qquad\qquad \text{for all } A \in \mathbf{Dyn}$$

$$[\![T]\!]_{\mathfrak{B}'} := [\![T]\!]_{\mathfrak{B}} \restriction \mathrm{dom}(\mathfrak{B}') \qquad\qquad \text{for all } T \in \mathbf{Stat}$$

$$[\![\Delta A]\!]_{\mathfrak{B}'} := [\![\Delta A]\!]_{\mathfrak{B}} \setminus m([\![\Delta A]\!]_{\mathfrak{A}} \setminus [\![\Delta A]\!]_{\mathfrak{A}'}) \qquad \text{for all } A \in \mathbf{Dyn}$$

We show that $\mathfrak{B}'$ is well-defined, i.e. that the above definitions induce a single structure and that the domain covers the elements mentioned in the relations.

- The definitions for every $A \in \mathbf{Dyn}$ of $[\![A]\!]_{\mathfrak{B}'}$ and $[\![\Delta A]\!]_{\mathfrak{B}'}$ induce

$$[\![A^\star]\!]_{\mathfrak{B}'} = ([\![\Delta A]\!]_{\mathfrak{B}'} \setminus [\![A]\!]_{\mathfrak{B}'}) \cup ([\![A]\!]_{\mathfrak{B}'} \setminus [\![\Delta A]\!]_{\mathfrak{B}'})$$

- Suppose there is $b \in m(\mathfrak{A})$ and $\mathbf{b}$ such that $b \in \mathbf{b}$ and $\mathbf{b} \in [\![A]\!]_{\mathfrak{B}'}$. If $m(a) = b$ for some $a \in [\![A]\!]_{\mathfrak{A}}$, then $m^{-1}(b) \in \mathrm{dom}(\mathfrak{A}')$ by definition of $\trianglelefteq$. Otherwise by lemma 3.3.17 (taking the contrapositive) $m^{-1}(b) \notin \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)$. So $m^{-1}(b) \notin [\![\Delta]\!]_{\mathfrak{A}}$. Since $\mathfrak{A}$ is supported, $m^{-1}(b) \in [\![Has]\!]_{\mathfrak{A}}$, and so $m^{-1}(b) \in \mathrm{dom}(\mathfrak{A}')$ by definition of $\trianglelefteq$.

- Suppose that for some $b \in m(\mathfrak{A})$, $b \in \mathbf{b} \in [\![A^\star]\!]_{\mathfrak{B}'}$. Wlog $\mathbf{b} \notin [\![A]\!]_{\mathfrak{B}'}$ (see previous case). So $\mathbf{b} \in [\![\Delta A]\!]_{\mathfrak{B}'}$. By lemma 3.3.13, $\mathbf{b} \in m([\![\Delta A]\!]_{\mathfrak{A}})$. So by definition of $[\![\Delta A]\!]_{\mathfrak{B}'}$, $\mathbf{b} \in m([\![\Delta A]\!]_{\mathfrak{A}'})$. So $\mathbf{b} \in m(\mathfrak{A}')$.

- The case $\mathbf{b} \in [\![T]\!]_{\mathfrak{B}'}$ is trivial.

Now we show $\mathfrak{B}' \vartriangleleft \mathfrak{B}$. By definition $\mathfrak{B}' \trianglelefteq \mathfrak{B}$. Moreover $\mathfrak{A}' \vartriangleleft \mathfrak{A}$, so $[\![\Delta A]\!]_{\mathfrak{A}'} \subset [\![\Delta A]\!]_{\mathfrak{A}}$ for some $A \in \mathbf{Dyn}$ or $\mathrm{dom}(\mathfrak{A}') \subset \mathrm{dom}(\mathfrak{A})$. In any case $\mathfrak{B}' \vartriangleleft \mathfrak{B}$.

Finally we show $m \restriction m^{-1}(\mathfrak{B}') : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}'$.

With $m' = m \restriction m^{-1}(\mathfrak{B}')$, we need $m' : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}'$, i.e. that $m^{-1}(\mathfrak{B}') = \mathrm{dom}(\mathfrak{A}')$, that $\mu(X) \subseteq \mathrm{dom}(\mathfrak{A}')$, that m is an embedding, that $\mathcal{B}_{\mathfrak{B}'}(m'(\mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)), 1) \subseteq m'(\mathfrak{A}')$ and that $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m'(\mathfrak{A}')$.

- We show $m^{-1}(\mathfrak{B}') = \mathrm{dom}(\mathfrak{A}')$. If $a \in m^{-1}(\mathfrak{B}')$ then $m(a) \in m(\mathfrak{A})$, and by definition of $\mathrm{dom}(\mathfrak{B}')$, $m(a) \notin m(\mathrm{dom}(\mathfrak{A}) \setminus \mathrm{dom}(\mathfrak{A}'))$ so $a \in \mathrm{dom}(\mathfrak{A}')$. In the other direction, if $a \in \mathrm{dom}(\mathfrak{A}')$ then $a \in \mathrm{dom}(\mathfrak{A})$ as well by definition of $\trianglelefteq$, so $m(a) \in \mathrm{dom}(\mathfrak{B}')$ by definition of $\mathrm{dom}(\mathfrak{B}')$.

- $\mu(X) \subseteq \text{dom}(\mathfrak{A}')$ by assumption on $\mathfrak{A}$

- Since m is an embedding and by definition of $\mathfrak{B}'$, $m' : \mathfrak{A}' \to \mathfrak{B}'$ is an embedding. To show that it is a 0-embedding, it suffices to show that for all $a \in \text{dom}(\mathfrak{B}')$, function symbol $h \in \mathbf{Stat}$, $b \in \text{Im}(m')$ such that $(a, b) \in [\![h]\!]_{\mathfrak{B}'}$, we have $a \in \text{Im}(m)$. By definition of $\mathfrak{B}'$, that premise implies $\delta_{\mathfrak{B}}(a, b) = 0$, so, since m is a 0-embedding and m' is a restriction of m, we get $a \in \text{Im}(m)$.

- Let $b \in \text{dom}(\mathfrak{B}')$ be at distance 1 (in $\mathfrak{B}'$) from $m'(a)$ with $a \in \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)$. Since $[\![\Delta]\!]_{\mathfrak{A}'} \subseteq [\![\Delta]\!]_{\mathfrak{A}}$ and by definition of m, $b \in m(\mathfrak{A})$. Suppose $b \notin \text{dom}(\mathfrak{A}')$: by definition of $\text{dom}(\mathfrak{B}')$, $b \notin \text{dom}(\mathfrak{B}')$. Absurd.

- Finally by $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq [\![\Delta]\!]_{\mathfrak{B}}$ and lemma 3.3.13, $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m([\![\Delta]\!]_{\mathfrak{A}})$. It is easy to see that $[\![\Delta]\!]_{\mathfrak{B}'} = [\![\Delta]\!]_{\mathfrak{B}} \setminus m([\![\Delta]\!]_{\mathfrak{A}} \setminus [\![\Delta]\!]_{\mathfrak{A}'})$, so $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m([\![\Delta]\!]_{\mathfrak{A}'})$. Since $\text{dom}(\mathfrak{A}') \subseteq m^{-1}(\mathfrak{B}')$, $[\![\Delta]\!]_{\mathfrak{B}'} \subseteq m'([\![\Delta]\!]_{\mathfrak{A}'})$.

$\square$

Recall the reasoning at the beginning of the **minimisation** paragraph above and keep in mind figure 3.7, which summarises the previous lemma. We use the contrapositive of the previous lemma to prove the next one.

Lemma 3.3.21. If ϕ is (I, d_c)-context-free and (I, d_r)-retractable, then $\downarrow(\phi \wedge \textit{Support})$ is $(I \wedge \textit{Has}, \max(d_c, d_r))$-retractable.

Proof. Let $d := \max(d_c, d_r)$. Consider $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \to \text{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}$ such that $\mathfrak{B}, m \circ \mu \vDash \downarrow(\phi \wedge \textit{Support})$ and $\mathfrak{B} \vDash \forall((I \wedge \textit{Has}) \vee m)$. We want $\mathfrak{A}, \mu \vDash \downarrow(\phi \wedge \textit{Support})$. Assume not. By lemma 3.3.16, $\phi \wedge \textit{Support}$ is $(I \wedge \textit{Has}, d_r)$-retractable. Since $d \geqslant d_r$, by lemma 3.3.2 $\mathfrak{A}, \mu \vDash \phi \wedge \textit{Support}$. So there is $\mathfrak{A}' \lhd \mathfrak{A}$ such that $\mathfrak{A}', \mu \vDash \phi \wedge \textit{Support}$. Note that this implies $\text{Im}(\mu) \subseteq \text{dom}(\mathfrak{A}')$.

Since $\mathfrak{A}$ is supported, by lemma 3.3.20, there is $\mathfrak{B}' \lhd \mathfrak{B}$ and $m' : (\mathfrak{A}', \mathcal{B}_{\mathfrak{A}'}^{\Delta,\mu}(d)) \to \mathfrak{B}'$ such that $m \upharpoonright m^{-1}(\mathfrak{B}') = m'$. We show $\mathfrak{B}', m \circ \mu \vDash \phi \wedge \textit{Support}$, which contradicts $\mathfrak{B}, m \circ \mu \vDash \downarrow (\phi \wedge \textit{Support})$:

Take any $\nu : \langle I \wedge \textit{Has} \rangle \to \text{dom}(\mathfrak{B}')$. Since $I \wedge \textit{Has}$ is a pre, quantifier free formula and $\mathfrak{B}' \unlhd \mathfrak{B}$, by lemma 3.3.15 $\mathfrak{B}, \nu \vDash I \wedge \textit{Has}$ implies $\mathfrak{B}', \nu \vDash I \wedge \textit{Has}$. Otherwise, by assumption $\text{Im}(\nu) \subseteq \text{Im}(m)$. Since $\text{Im}(\nu) \subseteq \text{dom}(\mathfrak{B}')$, $\text{Im}(\nu) \subseteq \text{Im}(m \upharpoonright m^{-1}(\mathfrak{B}')) = m'(\mathfrak{A}')$. So $\mathfrak{B}' \vDash \forall((I \wedge \textit{Has}) \vee m')$. By lemma 3.3.16, $\phi \wedge \textit{Support}$ is $(I \wedge \textit{Has}, d_c)$-context-free. Since $d \geqslant d_c$, by lemma 3.3.1 we get $\mathfrak{B}', m' \circ \mu \vDash \phi \wedge \textit{Support}$.

Since $m \upharpoonright m^{-1}(\mathfrak{B}') = m'$ and $\text{Im}(\mu) \subseteq \text{dom}(\mathfrak{A}')$, $m \circ \mu = m' \circ \mu$, so $\mathfrak{B}', m \circ \mu \vDash \phi$.

$\square$

3.3.3. Correctness of the $\vdash_*$ rules

Theorem 3.3.22 is a direct consequence of the lemmas above.

Theorem 3.3.22. If $\Gamma \vdash_* \phi$ then ϕ is $(\bigwedge \Gamma \wedge \textit{Has}, \|\phi\|)$-context-free and retractable.

Proof. By induction on the derivation.

- LITERAL$_*$: By lemmas 3.3.3 and 3.3.1 L is $(\textit{Has}, 0)$-context-free, by lemmas 3.3.4 and 3.3.2 L is $(\textit{Has}, 0)$-retractable.

- CLAUSE$_*$: Lemma 3.3.14 together with the proviso on CLAUSE$_*$ implies that $\forall \mathcal{C}$ is $(\mathcal{C}^\circ, |\langle \mathcal{C} \rangle| - 1)$-context-free. Since $\|\forall \mathcal{C}\| \geqslant |\langle \mathcal{C} \rangle| - 1$ (trivial), by lemma 3.3.1 $\forall \mathcal{C}$ is $(\mathcal{C}^\circ \wedge \mathsf{Has}, \|\forall \mathcal{C}\|)$-context-free. $\forall \mathcal{C}$ is $(\mathcal{C}^\circ \wedge \mathsf{Has}, \|\forall \mathcal{C}\|)$-retractable by direct application of lemmas 3.3.4, 3.3.8 and 3.3.2.

- DISJ$_*$: By induction hypothesis each ϕ_i is $(\wedge \Gamma_i \wedge \mathsf{Has}, \|\phi_i\|)$-context-free and retractable, and $\|\phi_1 \vee \phi_2\| = \max(\|\phi_1\|, \|\phi_2\|)$ so by lemmas 3.3.1, 3.3.2 and 3.3.6, $\phi_1 \vee \phi_2$ is $(\wedge \Gamma_1, \Gamma_2 \wedge \mathsf{Has}, \|\phi_1 \vee \phi_2\|)$-context-free and retractable.

- CONJ$_*$: By induction hypothesis each ϕ_i is $(\wedge \Gamma_i \wedge \mathsf{Has}, \|\phi_i\|)$-context-free and retractable, and $\|\phi_1 \wedge \phi_2\| = \max(\|\phi_1\|, \|\phi_2\|)$ so by lemmas 3.3.1, 3.3.2 and 3.3.5, $\phi_1 \vee \phi_2$ is $(\wedge \Gamma_1, \Gamma_2 \wedge \mathsf{Has}, \|\phi_1 \wedge \phi_2\|)$-context-free and retractable.

- $\forall$GUARD$_*$: Let $\psi := \forall x.\ \alpha(x, y) \rightarrow \phi$. By induction hypothesis ϕ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\phi\|)$-context-free and retractable.

  For context-free: By lemma 3.3.7, and since $\|\psi\| = \|\phi\| + 1$, ψ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\psi\|)$-context-free.

  For retractable: Since $\neg\alpha(x, y)$ is quantifier-free, by lemmas 3.3.4, 3.3.2 and 3.3.6, $\alpha(x, y) \rightarrow \phi$ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\psi\|)$-retractable, so by lemma 3.3.8 so is ψ.

- $\exists$GUARD$_*$: Let $\psi := \exists x.\ \alpha(x, y) \wedge \phi$. By induction hypothesis ϕ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\phi\|)$-context-free and retractable.

  For retractable: By lemma 3.3.10, and since $\|\psi\| = \|\phi\| + 1$, ψ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\psi\|)$-retractable.

  For context-free: Since $\alpha(x, y)$ is quantifier-free, by lemmas 3.3.3, 3.3.1 and 3.3.5, $\alpha(x, y) \wedge \phi$ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\psi\|)$-context-free, so by lemma 3.3.9 so is ψ.

- CIRCUM$_*$: Let $\psi := \downarrow (\phi \wedge \mathit{Support})$. By induction hypothesis, ϕ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\phi\|)$-context-free and retractable. By lemma 3.3.21, ψ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\phi\|)$-retractable (lemma 3.3.21 was applied to $\wedge \Gamma \wedge \mathsf{Has}$ and we then rewrote $\wedge \Gamma \wedge \mathsf{Has} \wedge \mathsf{Has}$ to $\wedge \Gamma \wedge \mathsf{Has}$); since $\|\psi\| \geqslant \|\phi\|$, by lemma 3.3.2 ψ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\psi\|)$-retractable. Similarly, by lemmas 3.3.19 and 3.3.1, ψ is $(\wedge \Gamma \wedge \mathsf{Has}, \|\psi\|)$-context-free.

$\square$

We have defined a class of formulas which survive going through matches modulo some side conditions. This is a good first step, but we need more: we need the side conditions themselves to survive throughout entire runs. That will be the focus of the next section.

3.4. Runs of context-free and retractable formulas

In this section, we refine the deduction rules $\vdash_*$ into $\vdash_\rho$, which lets us prove the main theorem of this chapter: if ϕ is $\vdash_\rho$-provable and the starting state is "nice", then the runs of the models of ϕ are exactly the executions of the rules of the canonical patterns of ϕ. Recall that, by the explanations in section 3.3 (page 39), we already have a correspondence between the execution on a set of models and the denotational runs on a set of models. So it now suffices to show that the runs of the instances of the canonical patterns of ϕ are exactly the executions of the rules of of the canonical patterns of ϕ.

Context-freeness and retractability are defined modulo some formula I on the match target. But this formula may not stay true at every step of the run. We start by refining good clauses to great clauses, and we show why they help maintain that invariant I. Then we introduce the

deduction rules $\vdash_{\bar{\rho}}$, explain how $\vdash_{\bar{\rho}}$ is a restriction of $\vdash_{\bar{*}}$ and motivate the differences. Then we state the main theorem before we go about proving it.

First, we give an overview of how that proof will go. Section 3.4.4 contains two key lemmas.

Runs on instances included in run on models   Lemma 3.4.6 will use context-freeness, and asserts that runs on pattern instances are included in runs on models of φ. The idea is: if φ is context-free modulo some I which only requires checking a small enough area, and φ ensures I on the postcondition (meaning φ ensures I*), then from a "nice" starting point $\mathfrak{M}_0$ we:

1. Take the first transition $\mathfrak{A}_1$ in a run on the *instances* of the canonical patterns of φ (so a run in $\rho_{\mathfrak{M}_0}(\mathsf{Insts}(\mathsf{P}))$).

2. Because $\mathfrak{M}_0$ is "nice", the precondition of $\mathfrak{A}_1$ satisfies I (modulo details).

3. The transition is an instance of a pattern $(\mathfrak{B}_1, K_1)$ such that $\mathfrak{B}$ satisfies φ (by definition). So by context-freeness of φ, $\mathfrak{A}_1$ satisfies φ as well.

4. Now that $\mathfrak{A}_1$ satisfies φ, a run on the models of φ (so a run in $\rho_{\mathfrak{M}_0}(\phi)$) can begin with $\mathfrak{A}_1$.

5. Also, $\mathfrak{A}_1$ satisfying φ means that its postcondition satisfies I (by assumption).

6. So the precondition of $\mathfrak{A}_2$, the next transition in the run, also satisfies I (by definition of a run).

7. Moreover, that next transition is also an instance of a pattern $(\mathfrak{B}_2, K_2)$ such that (...)

8. And so on.

Runs on models included in run on instances   Lemma 3.4.7 uses retractability, and asserts that the runs on models of φ are included in runs on pattern instances. The idea is: if φ is retractable modulo some I which only requires checking a small enough area, and φ ensures I on the postcondition, then from a "nice" starting $\mathfrak{M}_0$ point we:

1. Take the first transition $\mathfrak{A}_1$ in a run on the *models* of φ (so a run in $\rho_{\mathfrak{M}_0}(\phi)$.

2. Because $\mathfrak{M}_0$ is "nice", the precondition of that transition satisfies I (modulo details).

3. We define a special pattern p_1 which has a match into the $\mathfrak{A}_1$. Since φ is retractable, the first projection of p_1 satisfies φ as well. Because p_1 has a special shape, this means p_1 is one of the canonical patterns of φ.

4. Since $\mathfrak{A}_1$ is an instance of p_1, a run on the instances of canonical patterns of φ (so a run in $\rho_{\mathfrak{M}_0}(\mathsf{Insts}(\mathsf{P}))$) can begin with $\mathfrak{A}_1$.

5. We knew that $\mathfrak{A}_1$ satisfies φ, so its postcondition satisfies I (by assumption).

6. So the precondition of $\mathfrak{A}_2$, the next transition in the run, also satisfies I (by definition of a run).

7. Moreover, another special pattern p_2 with a match into $\mathfrak{A}_2$ can be constructed, and (...)

8. And so on.

Figure 3.8.: $\mathfrak{B}_1$ satisfies F_1 while $\mathfrak{B}_2$ does not. Changes highlighted in green.

3.4.1. Overview of the $\vdash_\rho$ rules and main theorem

Example. We show a $\vdash_*$-provable formula which does *not* remain satisfied throughout a run based on its canonical patterns.

Recall the formula $\mathsf{F}_1 := \forall x.\, \mathsf{On}(x) \vee \mathsf{Active}^\star(x)$ from figure 3.4. With $\mathcal{C} = \mathsf{On}(x) \vee \mathsf{Active}^\star(x)$, we have $\mathcal{C}^\circ \vdash_* \mathsf{F}_1$. So F_1 is $(\mathcal{C}^\circ, 0)$-context-free and retractable.

Consider figure 3.8. It represents two transitions, $\mathfrak{B}_1, \mathfrak{B}_2$, where the interpretation in $\mathfrak{B}_1$ of **Dyn*** is equal to the interpretation in $\mathfrak{B}_2$ of **Dyn** – that is, the middle column is the postcondition of $\mathfrak{B}_1$ and the precondition of $\mathfrak{B}_2$. In addition, the identity is a match from $\mathfrak{A}_1$ to $\mathfrak{B}_1$ and from $\mathfrak{A}_2$ to $\mathfrak{B}_2$. Each $\mathfrak{A}_i$ is defined by a grey, dashed outline. Now consider:

- In $\mathfrak{B}_1$, a ceases to be On. Shown in a dotted box is the transition $\mathfrak{A}_1$. There is a match $\mathrm{id}_{\mathfrak{A}_1} : (\mathfrak{A}_1, \mathcal{B}_{\mathfrak{A}_1}^{\Delta, \emptyset}(0)) \to \mathfrak{B}_1$. Since $\mathfrak{B}_1, x \mapsto b \vDash \mathcal{C}^\circ$ we obtain $\mathfrak{B}_1 \vDash \mathsf{F}_1$ (by context-freeness).

- In $\mathfrak{B}_2$, b goes from being On to being Active. Shown in a dotted box is the transition $\mathfrak{A}_2$. There is a match $\mathrm{id}_{\mathfrak{A}_2} : (\mathfrak{A}_2, \mathcal{B}_{\mathfrak{A}_2}^{\Delta, \emptyset}(0)) \to \mathfrak{B}_2$. But $\mathfrak{B}_1, x \mapsto a \nvDash \mathcal{C}^\circ$. And indeed: $\mathfrak{B}_1 \nvDash \mathsf{F}_1$.

The issue here is that $\mathfrak{B}_2, x \mapsto a \nvDash \mathcal{C}^\circ$. Or, said differently, the issue is that $\mathfrak{B}_1, x \mapsto a \nvDash \mathcal{C}^\star$. With a an element of $\mathfrak{B}_1$, $a \notin [\![\mathsf{On}^\star]\!]_{\mathfrak{B}_1}$ is fine. But the postcondition of $\mathfrak{B}_1$ becomes the precondition of $\mathfrak{B}_2$, and $a \notin [\![\mathsf{On}]\!]_{\mathfrak{B}_2}$ is *not* fine.

As a result, while the run $\mathfrak{B}_1, \mathfrak{B}_2$ is *not* in $\rho(\mathsf{F}_1)$, it is in $\rho(\mathsf{Insts}(\mathcal{P}(\mathsf{F}_1, 1)))$.

Informally, our goal is now to find for which formulas we could deduce $\mathfrak{B}_1 \vDash \forall \mathcal{C}^\star$ from $\mathfrak{B}_1 \vDash \mathsf{F}_1$. That way, the proviso $\mathcal{C}^\circ$ will propagate the run $\mathfrak{B}_1, \mathfrak{B}_2$. ∎

Great clauses

Great clauses are a restriction on good clauses.

Definition (Great clause). A clause $\mathcal{C}$ is *great* if $\mathcal{L}_{\mathrm{Dyn}}(\mathcal{C})$ nonempty implies that there is some $\mathsf{L} \in \mathcal{L}_{\mathrm{Dyn}}(\mathcal{C})$ such that $\overline{\mathsf{L}^\star} \in \mathcal{L}_{\mathrm{Dyn}}^\star(\mathcal{C})$. and $\mathsf{G}_\mathcal{C}$ is connected. ∎

Of course, any great clause is also good.

The first condition ($\mathcal{L}_{\mathrm{Dyn}}(\mathcal{C})$ nonempty implies some $\mathsf{L} \in \mathcal{L}_{\mathrm{Dyn}}(\mathcal{C})$ such that $\overline{\mathsf{L}^\star} \in \mathcal{L}_{\mathrm{Dyn}}^\star(\mathcal{C})$) solves the requirement outlined in the introduction of this section: if $\mathcal{C}$ depends on the precon-

dition at all, then $\mathcal{C}^\circ$ is equivalent to $\top$. So it becomes very easy for a precondition to satisfy $\mathcal{C}^\circ$.

The second condition ($G_\mathcal{C}$ connected) is a technical requirement to fulfill the requirements in the definition of context-free and retractable: the $\mathfrak{B} \vDash \forall(I \vee f)$ part of the definition needs $I(x_1, \ldots, x_k)$ to be true by default when the interpretations for $x_1, \ldots, x_k$ are too far apart.

Example. The clause $A(x) \vee \mathrm{Link}(x, y)$ is good but not great. Same for $A(x) \vee \neg\mathrm{Link}(x, y)$: good but not great. $\neg A(x) \vee B(x) \vee A^\star(x)$ is great. $\neg\mathrm{Link}^\star(x, y)$ is great. ∎

Rules $\vdash_\rho$

The rules $\vdash_\rho$ are a restriction of $\vdash_*$. The small ρ replaces the training wheel $*$ and should evoke runs, because $\vdash_\rho$-provable formulas will be well-behaved on runs.

$$\frac{}{\vdash_\rho L} \quad \text{LITERAL}_\rho \qquad\qquad \mathcal{C} \text{ great} \ \frac{}{\mathcal{C}^\circ \vdash_\rho \forall\mathcal{C}} \quad \text{CLAUSE}_\rho$$

$$\frac{\Gamma \vdash_\rho \phi_1 \qquad \Gamma \vdash_\rho \phi_2}{\Gamma \vdash_\rho \phi_1 \vee \phi_2} \quad \text{DISJ}_\rho \qquad\qquad \frac{\Gamma_1 \vdash_\rho \phi_1 \qquad \Gamma_2 \vdash_\rho \phi_2}{\Gamma_1, \Gamma_2 \vdash_\rho \phi_1 \wedge \phi_2} \quad \text{CONJ}_\rho$$

$$\frac{\vdash_\rho \phi}{\vdash_\rho \forall x. \ \alpha(x, y) \to \phi} \quad \forall\text{GUARD}_\rho \qquad\qquad \frac{\vdash_\rho \phi}{\vdash_\rho \exists x. \ \alpha(x, y) \wedge \phi} \quad \exists\text{GUARD}_\rho$$

$$\frac{\Gamma \vdash_\rho \phi}{\Gamma \vdash_\rho \downarrow(\phi \wedge \textit{Support})} \quad \text{CIRCUM}_\rho$$

Here are the differences between $\vdash_*$ and $\vdash_\rho$:

- The proviso of CLAUSE$_\rho$ is "$\mathcal{C}$ great" (it is "$\mathcal{C}$ good" in CLAUSE$_*$)

- DISJ$_\rho$ is additive (DISJ$_*$ is multiplicative)

- The context in the premise of $\forall$GUARD$_\rho$ and $\exists$GUARD$_\rho$ must be $\emptyset$ (it can be any Γ in $\forall$GUARD$_*$ and $\exists$GUARD$_*$)

The move from good to great was explained earlier.

We give some intuition behind the move from a multiplicative DISJ$_*$ to an additive DISJ$_\rho$: essentially, weakening does not hold in $\vdash_\rho$ (while it does in $\vdash_*$). One way to see it is that in $\vdash_*$ the context Γ describes a sufficient condition for a match to carry satisfaction (forward and backward). So strengthening Γ was free. Now, Γ should be an *invariant* in the sense that it should always carry over to the next transition in a run. So it can't be strengthened arbitrarily – or it becomes that much harder to prove in the next step. A multiplicative version of CLAUSE$_\rho$ would require proving too strong a conclusion (based on Γ_1, Γ_2) from too weak a premise (based on $\phi_1 \vee \phi_2$). In the case of CONJ$_\rho$, Γ_1, Γ_2 happens to be exactly be the weakest context that can prove $\phi_1 \wedge \phi_2$.

An empty context in $\forall$GUARD$_\rho$ and $\exists$GUARD$_\rho$ is required for similar reasons. For instance, in the case of $\forall$GUARD$_\rho$, keeping the same context Γ makes it too strong because the formula $\forall x. \ \alpha(x, y) \to \phi$ might be true while ϕ is not (for some interpretation). It may be possible to find richer contexts that carry enough information to expand the conditions where $\forall$GUARD$_\rho$ and $\exists$GUARD$_\rho$ apply, but we would lose an important feature of the current rules: since DISJ$_\rho$ is

additive, we can prove formulas of the form $\downarrow(F_1 \wedge G) \vee \downarrow(F_2 \wedge G)$ as long as F_1, F_2 are proved without CLAUSE$_\rho$. If $\forall$GUARD$_\rho$ and $\exists$GUARD$_\rho$ added context information depending on $\alpha(x,y)$, they would not be allowed in the proofs of F_1 and F_2. This feature has practical applications in chapter 5.

Of course, $\Gamma \vdash_{\bar{\rho}} \phi$ implies $\Gamma \vdash_{\bar{*}} \phi$.

Example. The formula $\forall x.\ \mathsf{On}(x)$ is $\vdash_{\bar{*}}$-provable but not $\vdash_{\bar{\rho}}$-provable. ∎

Example. The formula F_1 shown at the beginning of this section is $\vdash_{\bar{*}}$-provable (which means it is context-free and retractable for some radii, but not $\vdash_{\bar{\rho}}$-provable. As shown on figure 3.8, it should not be provable: the goal of $\vdash_{\bar{\rho}}$ is to build formulas that remain satisfied throughout executions; figure 3.8 shows two sucessive pattern matches where each pattern satisfies F_1 but the target of the 2nd match does not satisfy F_1. ∎

Overview of main theorem

To state the theorem, we need the following notions:

Definition (Initial states)**.** If I a pre, quantifier-free formula, a structure $\mathfrak{S} : \Theta^\circ$ is in $\mathsf{Init}(I)$ if $\mathfrak{S}$ is the empty structure or $\mathfrak{S} \vDash \forall I$. ∎

The theorem to prove is:

Theorem 3.5.1. If $\Gamma \vdash_{\bar{\rho}} \phi$ and $\mathfrak{M} \in \mathsf{Init}(\wedge \Gamma \wedge \mathsf{Has})$, then

$$\rho_{\mathfrak{M}}(\phi_{S0}) = \mathtt{exec}_{\mathfrak{M}}(\mathtt{rules}(\mathcal{P}(\phi_{S0}, \|\phi_{S0}\|)))$$

The proof is in section 3.5.

3.4.2. Run inclusion

Lemma 3.4.2 neatly packs up some sufficient conditions for a set of runs to be included in another. We start with an obvious fact:

Lemma 3.4.1. If A, B are sets of transitions such that $A \subseteq B$ then $\rho(A) \subseteq \rho(B)$

Proof. Let $r \in \rho(A)$. For all $\mathfrak{A}_i, \mathfrak{A}_{i+1} \in r$, $\mathsf{Post}(\mathfrak{A}_i) = \mathsf{Pre}(\mathfrak{A}_{i+1})$ by assumption on r and $\mathfrak{A}_i, \mathfrak{A}_{i+1} \in B$ by assumption on A and B. So $r \in \rho(B)$ by definition of $\rho(B)$. □

The idea behind this lemma is to explicitly state that some invariant is needed; here the set S is a set of states, and run inclusion happens with the invariant shifting from one state to the next.

Lemma 3.4.2. If $\mathfrak{S} : \Theta^\circ$, M, N are two sets of transitions, S a set of states, $\mathfrak{S} \in S$ and for all supported $\mathfrak{A} \in M$, $\mathsf{Pre}(\mathfrak{A}) \in S$ implies $\mathfrak{A} \in N$ and $\mathsf{Post}(\mathfrak{A}) \in S$, then $\rho_{\mathfrak{S}}(M) \subseteq \rho_{\mathfrak{S}}(N)$.

Proof. Let $r := \mathfrak{A}_1, \ldots \in \rho_{\mathfrak{S}}(M)$. By induction on i, we show $\mathsf{Pre}(\mathfrak{A}_i) \in S$ and $\mathsf{Post}(\mathfrak{A}_i) \in S$ for all $A_i \in r$.

If $i = 1$ we have $\mathsf{Pre}(\mathfrak{A}_1) = \mathfrak{S}$ by definition of $\rho_{\mathfrak{S}}(M)$ so by assumption on $\mathfrak{S}$ and S, $\mathsf{Pre}(\mathfrak{A}_1) \in S$. So $\mathsf{Post}(\mathfrak{A}_1) \in S$.

Otherwise there is $\mathfrak{A}_{i-1} \in M$ such that $\mathsf{Post}(\mathfrak{A}_{i-1}) = \mathsf{Pre}(\mathfrak{A}_i)$, and by induction hypothesis $\mathsf{Post}(\mathfrak{A}_{i-1}) \in S$, so $\mathsf{Pre}(\mathfrak{A}_i) \in S$, and $\mathsf{Post}(\mathfrak{A}_i) \in S$.

By definition of $\rho_{\mathfrak{S}}(N)$, it suffices to show $\mathfrak{A}_i \in N$ for all $i > 0$, $\mathfrak{A}_i \in r$. Take any such $\mathfrak{A}_i$. By assumption $\mathfrak{A}_i \in M$ and by the above $\mathsf{Pre}(\mathfrak{A}_i) \in S$. So $\mathfrak{A}_i \in N$. □

3.4.3. d-**validity**

Definition (d-validity). For $d \geqslant 0$, a formula ψ is d-*valid* whenever $\mathfrak{A}, \mu \vDash \psi$ for all $\mathfrak{A}, \mu : \langle \psi \rangle \to \mathrm{dom}(\mathfrak{A})$ such that for some $a, b \in \mathrm{Im}(\mu)$, $\delta_{\mathfrak{A}}(a, b) > d$. ∎

This technical requirement of d-validity is related to the proviso that $G_{\mathcal{C}}$ is connected in a great clause $\mathcal{C}$. The lemmas below show its usefulness: given a match $m : p \to \mathfrak{B}$, a requirement of both context-freeness and retractability is that for some I, we must have $\mathfrak{B} \vDash \forall (I \vee m)$. During a run, we will see that $\mathrm{Pre}(\mathfrak{B}) \vDash \forall I$ is given by the previous run step (or by the initial state of the run). However, some elements may *appear*, i.e. not be in $\mathrm{Pre}(\mathfrak{B})$. For those, we do know that they have an m-preimage in p, so they are mostly taken care of by the "m" disjunct of $\mathfrak{B} \vDash \forall (I \vee m)$.

However there is a technical corner case, which we now discuss informally: suppose I has 2 free variables, $a, b \in \mathfrak{B}$, $a \notin \mathrm{Pre}(\mathfrak{B})$, and b does not have an m-preimage. So the "m" disjunct cannot help us. We need $\mathfrak{B} \vDash I(a, b)$, but that is not given to us by $\mathrm{Pre}(\mathfrak{B}) \vDash \forall I$ since $a \notin \mathrm{Pre}(\mathfrak{B})$. We will see in the lemma below that by their characterisation, a and b must be far apart in $\mathfrak{B}$, in fact, far enough apart that d-validity ensures $\mathfrak{B} \vDash I(a, b)$!

Lemma 3.4.3. If $\mathfrak{A}, \mathfrak{B} : \Theta$, I is d-valid, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}(\llbracket \Delta \rrbracket_{\mathfrak{A}}, d)) \to \mathfrak{B}$, $\mathrm{Pre}(\mathfrak{B}) \in \mathrm{Init}(I)$, and $\mathfrak{B}$ is supported, then $\mathfrak{B} \vDash \forall (I \vee m)$

Proof. Take any $\nu : \langle I \rangle \to \mathrm{dom}(\mathfrak{B})$. If $\mathrm{Im}(\nu) \subseteq \mathrm{dom}(\mathrm{Pre}(\mathfrak{B}))$, we are done since $\mathrm{Pre}(\mathfrak{B}) \in \mathrm{Init}(I)$.

Otherwise, there is $a \in \mathrm{Im}(\nu) \setminus \mathrm{dom}(\mathrm{Pre}(\mathfrak{B}))$, so $\delta_{\mathfrak{B}}(a, \llbracket \neg \mathsf{Has} \rrbracket_{\mathfrak{B}}) = 0$.

Since $\mathfrak{B}$ is supported, $\llbracket \neg \mathsf{Has} \rrbracket_{\mathfrak{B}} \subseteq \llbracket \Delta \rrbracket_{\mathfrak{B}}$, so $\delta_{\mathfrak{B}}(a, \llbracket \Delta \rrbracket_{\mathfrak{B}}) = 0$. We get the following inclusions by the lemmas indicated below:

$$\mathcal{B}_{\mathfrak{B}}(\llbracket \Delta \rrbracket_{\mathfrak{B}}, d) \underset{\text{lemma 3.3.13}}{\subseteq} \mathcal{B}_{\mathfrak{B}}(m(\llbracket \Delta \rrbracket_{\mathfrak{A}}), d) \underset{\text{lemma 2.4.5}}{\subseteq} m(\mathfrak{A})$$

I is d-valid so wlog $\delta_{\mathfrak{B}}(a, b) \leqslant d$ for every $a, b \in \mathrm{Im}(\nu)$. We have $\mathrm{Im}(\nu) \subseteq \mathcal{B}_{\mathfrak{B}}(\llbracket \Delta \rrbracket_{\mathfrak{B}}, d) \subseteq m(\mathfrak{A})$, so we are done. □

The next lemma is key to taking care of the corner case mentioned above. It explains why great clauses must have a connected graph.

Lemma 3.4.4. If $\Gamma \vdash_{\rho} \phi$ then $\wedge \Gamma$ is $\|\phi\|$-valid.

Proof. First, note that if a formula is d-valid, it is d'-valid for all $d' \geqslant d$. We proceed by induction on the derivation.

- LITERAL$_{\rho}$: Trivial.

- CLAUSE$_{\rho}$: $\mathcal{C}$ is great so $G_{\mathcal{C}}$ is connected; trivially, so is $G_{\mathcal{C}^{\circ}}$. By lemma 3.3.12, if $\mathfrak{A}, \mu \nvDash \mathcal{C}^{\circ}$ then $\delta_{\mathfrak{A}}(a, b) \leqslant d$ for all $a, b \in \mathrm{Im}(\mu)$; this is the contrapositive phrasing of $|\mathrm{Im}(\mu)|$-validity. Since $|\mathrm{Im}(\mu)| \leqslant |\langle \mathcal{C}^{\circ} \rangle| = \|\forall \mathcal{C}\|$, we are done.

- DISJ$_{\rho}$, $\forall$GUARD$_{\rho}$, $\exists$GUARD$_{\rho}$, CIRCUM$_{\rho}$: Direct application of induction hypothesis: $\wedge \Gamma$ is $\|\phi\|$-valid and the formula in the conclusion has quantifier rank at least that of ϕ.

- CONJ$_{\rho}$: We have a $\|\phi_i\|$-valid $\wedge \Gamma_i$ for $i \in \{1, 2\}$. It is easy to see that $\wedge \Gamma_1, \Gamma_2$ is $\max(\|\phi_1\|, \|\phi_2\|)$-valid; this happens to be the quantifier rank of $\phi_1 \wedge \phi_2$.

□

3.4.4. Run equality modulo context-free, retractable, d-validity

In this subsection we state the main lemmas which lead to the theorem. The next lemma simply says that support plays nice with validity in image complement.

Lemma 3.4.5. If $\mathfrak{m} : (\mathfrak{A}, K) \to \mathfrak{B}$ and $\mathfrak{B}$ is supported then $\mathfrak{B} \vDash \forall(\mathsf{Has} \vee \mathfrak{m})$.

Proof. Let $\mathfrak{a} \notin \mathfrak{m}(\mathfrak{A})$ so by lemma 3.3.13 $\mathfrak{a} \notin \llbracket\Delta\rrbracket_\mathfrak{B}$. Since $\mathfrak{B}$ is supported, $\mathfrak{a} \in \llbracket\mathsf{Has}\rrbracket_\mathfrak{B}$. □

The two lemmas after that are dual of each other and correspond to the reasoning outlined in the introduction of this chapter.

Lemma 3.4.6. If ϕ is $(I \wedge \mathsf{Has}, d)$-context-free, I is d-valid, $\phi \vDash \forall I^\star$, and $\mathfrak{S} \in \mathsf{Init}(I)$, then $\rho_\mathfrak{S}(\mathsf{Inst}(\mathcal{P}(\phi, d))) \subseteq \rho_\mathfrak{S}(\phi)$

Proof. By lemma 3.4.2 we need $\mathfrak{B} \in \llbracket\exists\phi\rrbracket$ and $\mathsf{Post}(\mathfrak{B}) \in \mathsf{Init}(I)$ for every supported $\mathfrak{B} \in \mathsf{Insts}(\mathcal{P}(\phi, d))$ such that $\mathsf{Pre}(\mathfrak{B}) \in \mathsf{Init}(I)$. Since $\phi \vDash \forall I^\star$, it will suffice to show $\mathfrak{B}, \nu \vDash \phi$ for some $\nu : \langle\phi\rangle \to \mathsf{dom}(\mathfrak{B})$; this will imply both $\mathfrak{B} \in \llbracket\exists\phi\rrbracket$ and $\mathsf{Post}(\mathfrak{B}) \in \mathsf{Init}(I)$.

By definition of $\mathcal{P}(\phi, d)$, there is $(\mathfrak{A}, \mathcal{B}_\mathfrak{A}^{\Delta,\mu}(d)) \in \mathcal{P}(\phi, d)$ and $\mathfrak{m} : (\mathfrak{A}, \mathcal{B}_\mathfrak{A}^{\Delta,\mu}(d)) \to \mathfrak{B}$ with $\mu : \langle\phi\rangle \to \mathsf{dom}(\mathfrak{A})$ and $\mathfrak{A}, \mu \vDash \phi$. We will show $\mathfrak{B}, \mathfrak{m} \circ \mu \vDash \phi$. Since ϕ is context-free modulo $I \wedge \mathsf{Has}$, it suffices to show $\mathfrak{B} \vDash \forall((I \wedge \mathsf{Has}) \vee \mathfrak{m})$. We get $\mathfrak{B} \vDash \forall(\mathsf{Has} \vee \mathfrak{m})$ by lemma 3.4.5. Since $\mathcal{B}_\mathfrak{A}(\llbracket\Delta\rrbracket_\mathfrak{A}, d) \subseteq \mathcal{B}_\mathfrak{A}^{\Delta,\mu}(d)$, by lemma 2.4.4 $\mathfrak{m} : (\mathfrak{A}, \mathcal{B}_\mathfrak{A}\llbracket\Delta\rrbracket_\mathfrak{A}(d)) \to \mathfrak{B}$, so we can use lemma 3.4.3 to get $\mathfrak{B} \vDash \forall(I \vee \mathfrak{m})$. □

Lemma 3.4.7. If ϕ is $(I \wedge \mathsf{Has}, d)$-retractable, I is d-valid, $\phi \vDash \forall I^\star$, and $\mathfrak{S} \in \mathsf{Init}(I)$ then $\rho_\mathfrak{S}(\phi) \subseteq \rho_\mathfrak{S}(\mathsf{Inst}(\mathcal{P}(\phi, d)))$.

Proof. By lemma 3.4.2 we need $\mathfrak{B} \in \mathsf{Inst}(\mathcal{P}(\phi, d))$ and $\mathsf{Post}(\mathfrak{B}) \in \mathsf{Init}(I)$ for every supported $\mathfrak{B} \in \llbracket\exists\phi\rrbracket$ such that $\mathsf{Pre}(\mathfrak{B}) \in \mathsf{Init}(I)$. Since $\phi \vDash \forall I^\star$, $\mathsf{Post}(\mathfrak{B}) \in \mathsf{Init}(I)$; so it suffices to show $\mathfrak{B} \in \mathsf{Inst}(\mathcal{P}(\phi, d))$.

By assumption there is $\mu : \langle\phi\rangle \to \mathsf{dom}(\mathfrak{B})$ such that $\mathfrak{B}, \mu \vDash \phi$. Let $\mathfrak{A} := \mathfrak{B} \upharpoonright \mathcal{B}_\mathfrak{B}^{\Delta,\mu}(d+1)$. We trivially have $\mathsf{id}_\mathfrak{A} : (\mathfrak{A}, \mathcal{B}_\mathfrak{A}^{\Delta,\mu}(d)) \to \mathfrak{B}$.

By lemma 2.4.4, $\mathsf{id}_\mathfrak{A} : (\mathfrak{A}, \mathcal{B}_\mathfrak{A}(\llbracket\Delta\rrbracket_\mathfrak{A}, d)) \to \mathfrak{B}$ so by lemma 3.4.3, $\mathfrak{B} \vDash \forall(I \vee \mathsf{id}_\mathfrak{A})$. By lemma 3.4.5, $\mathfrak{B} \vDash \forall(\mathsf{Has} \vee \mathfrak{m})$. So $\mathfrak{B} \vDash \forall((I \wedge \mathsf{Has}) \vee \mathfrak{m})$. Since ϕ is retractable with d a retraction radius, $\mathfrak{A}, \mu \vDash \phi$. So $(\mathfrak{A}, \mathcal{B}_\mathfrak{A}^{\Delta,\mu}(d)) \in \mathcal{P}(\phi, d)$, so $\mathfrak{B} \in \mathsf{Inst}(\mathcal{P}(\phi, d))$. □

The next lemma is key to maintaining the invariant I in (I, d)-context-freeness and retractability throughout runs. It explains why DISJ_ρ is additive, why $\vdash_\rho$ does not admit weakening, why the context of $\forall\mathsf{GUARD}_\rho$ and $\exists\mathsf{GUARD}_\rho$ is empty, and why great clauses must have dual, pre-post literals when they mention at least one literal with a symbol in **Dyn**.

Lemma 3.4.8. If $\Gamma \vdash_\rho \phi$ then $\phi \vDash \forall(\wedge\Gamma)^\star$

Proof. By induction on the derivation.

- $\mathsf{LITERAL}_\rho$, DISJ_ρ, CONJ_ρ, CIRCUM_ρ, $\forall\mathsf{GUARD}_\rho$, $\exists\mathsf{GUARD}_\rho$: Trivial.

- CLAUSE_ρ $\mathcal{C}$ is great, so either $\mathcal{L}_{\mathbf{Dyn}}(\mathcal{C})$ is empty, in which case $\mathcal{C} = \mathcal{C}^\star$, or it is not, in which case $\mathcal{C}^\star \equiv \top$.

□

We can now say when $\rho_{\mathfrak{S}}(\psi) = \rho_{\mathfrak{S}}(\mathsf{Insts}(\mathcal{P}(\psi, \|\psi\|)))$: when it is the conjunction of a $\vdash_{\bar{\rho}}$-provable formula and *Support*, and when $\mathfrak{S}$ is nice enough:

Theorem 3.4.9. Let $\psi := \phi \wedge Support$. If $\Gamma \vdash_{\bar{\rho}} \phi$ and $\mathfrak{S} \in \mathsf{Init}(\wedge\Gamma)$ then

$$\rho_{\mathfrak{S}}(\psi) = \rho_{\mathfrak{S}}(\mathsf{Inst}(\mathcal{P}(\psi, \|\psi\|)))$$

Proof. To apply lemmas 3.4.6 and 3.4.7 we need $\phi \wedge Support$ to be $(\wedge\Gamma \wedge \mathsf{Has}, \|\phi \wedge Support\|)$-context-free and retractable, $\wedge\Gamma$ to be $\|\phi \wedge Support\|$-valid, and $\phi \wedge Support \vDash \forall(\wedge\Gamma)^{\star}$.

By lemma 3.3.22, ϕ is $(\wedge\Gamma \wedge \mathsf{Has}, \|\phi\|)$-context-free and retractable. By lemma 3.3.16, $\phi \wedge Support$ is as well. By lemma 3.4.4, $\wedge\Gamma$ is $\|\phi\|$-valid, so it is $\|\phi \wedge Support\|$-valid. By lemma 3.4.8, $\phi \vDash \forall(\wedge\Gamma)^{\star}$, so $\phi \wedge Support \vDash \forall(\wedge\Gamma)^{\star}$.

□

3.5. Main theorem

This section summarises the last two chapters by saying when the runs of a formula are equal to the execution of the rules induced by the canonical patterns of that formula. It is simply a matter of applying theorems 3.4.9 and 2.6.4.

Theorem 3.5.1. If $\Gamma \vdash_{\bar{\rho}} \phi$ and $\mathfrak{M} \in \mathsf{Init}(\wedge\Gamma \wedge \mathsf{Has})$, then

$$\rho_{\mathfrak{M}}(\phi_{S0}) = \mathtt{exec}_{\mathfrak{M}}(\mathtt{rules}(\mathcal{P}(\phi_{S0}, \|\phi_{S0}\|)))$$

Proof. Note that $Support_0 = Support \wedge \mathsf{Has}_0 \wedge \mathsf{Has}_0^{\star}$. By lemma A.0.1 (given in appendix), there is $\psi \equiv \mathsf{Has}_0 \wedge \mathsf{Has}_0^{\star}$ and Γ' such that $\Gamma' \vdash_{\bar{\rho}} \psi$, so $\Gamma, \Gamma' \vdash_{\bar{\rho}} \phi \wedge \psi$. Since $\mathfrak{M} \in \mathsf{Init}(\wedge\Gamma)$, theorem 3.4.9 is applicable to $\phi \wedge \psi$, so

$$\rho_{\mathfrak{M}}(\phi_{S0}) = \rho_{\mathfrak{M}}(\mathsf{Inst}(\mathcal{P}(\phi_{S0}, \|\phi_{S0}\|)))$$

By definition any pattern in $\mathcal{P}(\phi_{S0}, \|\phi_{S0}\|)$ is pure and strongly supported. $\mathfrak{M} \in \mathsf{Init}(\mathsf{Has})$, so theorem 2.6.4 is applicable, so

$$\rho_{\mathfrak{M}}(\mathsf{Inst}(\mathcal{P}(\phi_{S0}, \|\phi_{S0}\|))) = \mathtt{exec}_{\mathfrak{M}}(\mathtt{rules}(\mathcal{P}(\phi_{S0}, \|\phi_{S0}\|)))$$

□

We complete figure 2.9 and show in figure 3.9 the connection between the runs on the instances of $\mathcal{P}(\phi, \|\phi\|)$ and the runs on the models of ϕ. The satisfaction relation on the right is missing its interpretation component.


Figure 3.9.: Left: a run in $\rho_{\mathfrak{M}_0}(\mathsf{Insts}(P))$. Right: a run in $\rho_{\mathfrak{M}_0}(\phi)$

4. Bounded change-minimal specifications

In the previous chapter we showed when runs on formulas can be interpreted as the execution of a set of rules. But is that set of rules finite? In this chapter, we show when that set of rules can be finite. This will be done by defining another set of deduction rules (the name clash is unfortunate), $\vdash_\lhd$, such that $\vdash_\lhd$-provable formulas have a finite set of canonical patterns.

The proof uses a property called preservation. Essentially, models of a preserved formula ϕ can lose changes in a very controlled way and still satisfy ϕ. Preservation is useful as a property because it plays well with logical operations.

A first step is to restrict the class of transitions further. We introduce *Bounded Forests with Links* in section 4.1. One crucial property of BFLs is that the graph of functions in $\mathbf{Stat}_{fun}$ is bounded by some parameter n. In section 4.2 we define preservation. First, we introduce a parameterized operation on BFLs called *taking a sub*. A formula ϕ is preserved (modulo some parameters) if the sub of a model of ϕ is still a model of ϕ. In section 4.3 we show why preservation leads to a finite set of canonical patterns, and finally in section 4.4 we introduce deduction rules to build preserved formulas.

We start by going back to the issue of an infinite set of rules. If $\mathcal{P}(\phi, \|\phi\|)$ is infinite, theorem 3.5.1 is not very useful. If there are infinitely many patterns to sift through, an execution engine for $\twoheadrightarrow$ cannot do much good. Where does the infinity come from? First, there is no bound on the size of a ball (even of radius 0) around a element of a transition. We will restrict the class of transitions so that for every element, its d-ball is bounded for all $d \geqslant 0$.

Moreover, the domains and application conditions of the canonical patterns are of the form $\mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)$, i.e. balls around two sets: 1) the image of an interpretation (μ), and 2) changes (Δ). The image of μ is bounded by the formula, but the number of changes is not. The min operator $\downarrow$ can help, but there is still work to do: a formula may very well be of the form $\downarrow\phi$ yet have infinitely many canonical patterns.
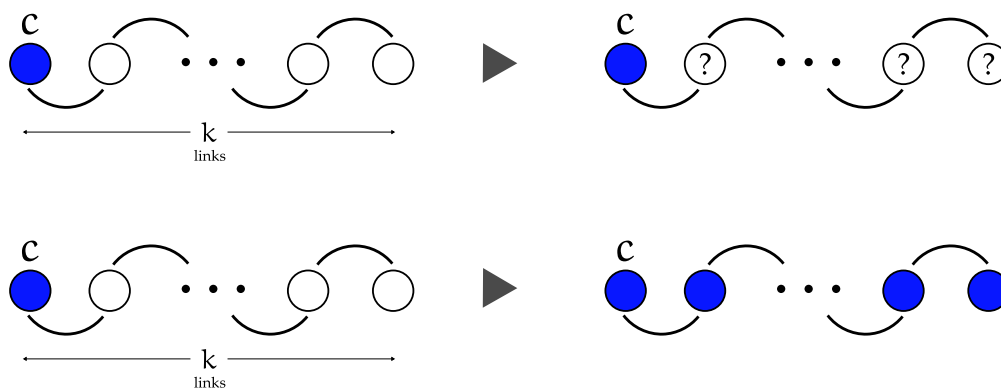

Figure 4.1.: Precondition of $\mathfrak{A}_k$. $\mathfrak{A}_k \vDash \downarrow CC$, the entire connected component of c must become Blue.

Example. Suppose Link, Link* or of outdegree 2 at most. Consider:

$$\downarrow CC := \downarrow (\mathsf{Blue}(c) \wedge$$
$$\forall x. \; (x \neq c \to \neg\mathsf{Blue}(x)) \wedge$$
$$\forall y. \; \mathsf{Blue}^\star(x) \to \mathsf{Link}^\star(x, y) \to \mathsf{Blue}^\star(y))$$

with $\mathsf{Blue} \in \mathbf{Dyn}$. $CC(c)$ specifies *a)* that the interpretation of c is Blue and no other element is blue, and *b)* an element with a Blue* neighbor must be Blue*.

For every $k \geqslant 0$ consider the transition $\mathfrak{A}_k$, with a precondition as in figure 4.1. There is a connected component (for the relation Link) of size k. Since Blue* is true for $[\![c]\!]_{\mathfrak{A}_k}$, Blue* must be true of $[\![c]\!]_{\mathfrak{A}_k}$'s neighbour, and so on for the entire connected component. For every $k \geqslant 0$, the transition of figure 4.1 is a model of $\downarrow CC$. So the number of canonical patterns for $\downarrow CC$ is infinite even up to isomorphism. ∎

4.1. Bounded Forests with Links

In this section we define bounded forests with links and show the theories that characterise them. The proofs are in the appendix.

We are moving to a more restricted class of transitions, so that a transition can always be described as a pair of forests with trees of bounded height, linked by a functional and symmetric relation. The bounded height of the trees is necessary to get a finite number of canonical patterns.

Definition (BFL signature). Θ is a *BFL signature* if $\mathbf{Stat}_{\mathsf{fun}}$ contains a distinguished function symbol parent. ∎

Definition (n-BFL). If $n \geqslant 0$, Θ is a BFL signature and $\mathfrak{A} : \Theta$, $\mathfrak{A}$ is an *n-BFL* (notation: $\mathfrak{A} : \Theta_n$) whenever:

- The loop-free union of the graphs of functions $[\![f]\!]_{\mathfrak{A}}$ for $f \in \mathbf{Stat}_{\mathsf{fun}} \setminus \{\mathsf{parent}\}$ is a forest.

- The trees of that forest have height at most n.

- The loop-free graph of $[\![\mathsf{parent}]\!]_{\mathfrak{A}}$ is the parent relation in that forest.

- The relations $[\![\mathsf{Link}]\!]_{\mathfrak{A}}$ and $[\![\mathsf{Link}^\star]\!]_{\mathfrak{A}}$ are symmetric and functional.

∎

The parent function makes the proofs simpler since any embedding is now a 0-embedding (since accessibility through the function graph is now symmetric).

Note that if $\mathfrak{A} : \Theta_n$, the restriction of $G_{\mathfrak{A}}$ to 0-edges, $G_{\mathfrak{A}}^0$, is the undirected version of the forest in the definition of n-BFL.

Implicit parameter [BFL height] $n \geqslant 0$ is an arbitrary maximum height.

Definition (Tree). A tree of t of $G_{\mathfrak{A}}$ is the restriction of $G_{\mathfrak{A}}^0$ to one of its connected components. $V_t \subseteq \mathsf{dom}(\mathfrak{A})$ is the vertex set of that tree. If $a \in \mathsf{dom}(\mathfrak{A})$, t_a is the tree of $G_{\mathfrak{A}}$ that contains a. ∎

It is easy to see that $V_{t_a} = \mathcal{B}_{\mathfrak{A}}(a, 0)$ for all $a \in \mathsf{dom}(\mathfrak{A})$ and that for any $b \in V_{t_a}$, $b \in \mathfrak{A} \upharpoonright \{a\}$.

Example. In figure 4.2, the transition $\mathfrak{A}$ has two trees. We have $\mathbf{Stat}_{\mathsf{fun}} = \{f, g, \mathsf{parent}\}$, and $\mathfrak{A}$ is a 1-BFL. The interpretations of Link and Link* are symmetric and functional (no change from earlier graphical representations). The parent function is indeed the inverse of all other functions everywhere except on tree leafs. The tree containing a is t_a, and $V_{t_a} = \mathcal{B}_{\mathfrak{A}}(\{a\}, 0) = \{a, b, c\}$. ∎

Figure 4.2.: A 1-BFL $\mathfrak{A}$

It is possible to characterise n-BFLs with universal, first-order formulas:

Definition. The theory of n-BFLs is:

$$\mathcal{T}_n := Symlink \wedge Funlink \wedge \mathcal{T}_n'$$

where

$$\mathcal{T}_n' := \forall x. PSpec(x) \wedge H^n(x)$$

$$PSpec(x) := \Big(\bigwedge_{f \in \mathbf{Stat}_{fun}} f(x) \neq x \to \mathrm{parent}(f(x)) = x \Big)$$

$$\wedge \Big(\mathrm{parent}(x) \neq x \to \bigvee_{f \in \mathbf{Stat}_{fun}} f(\mathrm{parent}(x)) = x \Big)$$

$$H^i(x) := \bigwedge_{f \in \mathbf{Stat}_{fun}} f(x) \neq x \to H^{i-1}(f(x)) \qquad \text{(for } 0 \leqslant i < n)$$

$$H^{-1}(x) := \bot$$

$$FunLink' := \forall x, y, z.\ \mathrm{Link}(x,y) \to \mathrm{Link}(x,z) \to y = z$$

$$FunLink := FunLink \wedge FunLink^\star$$

$$Symlink' := \forall x, y.\ \mathrm{Link}(x,y) \to \mathrm{Link}(y,x)$$

$$Symlink := Symlink \wedge Symlink^\star$$

$\blacksquare$

Lemma 4.1.1. If $\mathfrak{A} : \Theta$, $\mathfrak{A} : \Theta_n$ iff $\mathfrak{A} \vDash \mathcal{T}_n$.

See proof in Appendix A.

Definition. $\mathcal{T}_{ns} := \mathcal{T}_n \wedge Support$ $\blacksquare$

4.2. Preservation

In this section we define preservation and give some examples to help build intuition. The definition is a bit involved and goes through intermediary steps. First, some elements in an n-BFL are considered active if they are involved in a particular kind of change. Next, a tree t in an n-BFL $\mathfrak{A}$ can be *cleared*, producing a *sub*transition $\mathfrak{A}'$. This clearing operation removes changes involving t in a very controlled way. For instance, it does not remove changes that touch trees with active elements. Subtransitions are parameterised by a set of elements and a distance. We then define preservation on formulas by saying that for a formula ϕ, if satisfaction is closed under the operation of clearing a tree (thus producing a sub) modulo those parameters, then ϕ is preserved modulo those parameters.

We conclude with more examples to help understand subs and preservation.

Notation (Useful formula macros)**.** If $A \in \mathbf{Dyn}$ and x is a tuple of variables of size the arity of A, then:

$$\oplus A(x) := \neg A(x) \wedge A^\star(x) \qquad \text{(``}A(x) \text{ becomes true'')}$$
$$\ominus A(x) := A(x) \wedge \neg A^\star(x) \qquad \text{(``}A(x) \text{ becomes false'')}$$

∎

Definition (Operations on relations and sets)**.** If R is a relation and X is a set,

$$R\,/\,X := \{\, r \in R \mid r \cap X \neq \emptyset \,\} \qquad \text{(tuples that mention X)}$$
$$R \setminus X := R \setminus (R\,/\,X) \qquad \text{(tuples that do not mention X)}$$

∎

For instance, $\{(a,b),(a,c)\}\,/\,\{b\} = \{(a,b)\}$ and $\{(a,b),(a,c)\}\setminus\{b\} = \{(a,c)\}$. The operation $\setminus$ generalizes the usual difference in that the definition reduces to set or relation difference one when R is a unary relation and either R is understood as a set, or when X is understood as a set of 1-tuples.

We move on to preservation. First, some preliminary definitions:

Definition (Active elements and trees)**.** If $\mathfrak{A} : \Theta_n$, $a \in \mathfrak{A}$ is *active* whenever at least one of the following is true:

- $a \in [\![\Delta A]\!]_{\mathfrak{A}}$, for some $A \in \mathbf{Dyn}$ unary

- There is $b \in V_{t_a}$ such that $(a,b) \in [\![\ominus\mathsf{Link}]\!]_{\mathfrak{A}}$

- There is $b \in \mathrm{dom}(\mathfrak{A})$ such that $(a,b) \in [\![\oplus\mathsf{Link}]\!]_{\mathfrak{A}}$

A tree t of $G_{\mathfrak{A}}$ is active if it contains at least one active element. The set of elements in active trees is $\mathcal{A}(\mathfrak{A})$, and for t a tree of $G_{\mathfrak{A}}$, the set of elements in active trees different from t is $\mathcal{A}_{\bar{t}}(\mathfrak{A}) := \mathcal{A}(\mathfrak{A}) \setminus V_t$. ∎

In plain english, a elements is active if it is involved in any change other than an edge deletion involving a different tree.

Example. In figure 4.3, the transition $\mathfrak{A}$ contains 2 trees, one with the elements a, b, c and one with the single element d. The active elements are in red. a is active because it is in $[\![\Delta A]\!]_{\mathfrak{A}}$, and c is active because it loses an internal edge. However, b and d both lose an edge connected to another tree (the orientation of the edge is irrelevant), so they are not active. As a result, t_a is active because it has at least one active element, but t_d is not active. ∎

Figure 4.3.: active elements in red.

Figure 4.4.: $\mathfrak{B}$, after clearing t_a in $\mathfrak{A}$.

Definition (Subs). If $\mathfrak{A} : \Theta_n$, $\mathfrak{B} \trianglelefteq \mathfrak{A}$, $K \subseteq \mathrm{dom}(\mathfrak{B})$, $d \geqslant 0$ and t is a tree of $G_{\mathfrak{A}}$, $\mathfrak{B}$ is a (K, d)-*sub* of $\mathfrak{A}$ with cleared tree t whenever there is $K' \subseteq \mathrm{dom}(\mathfrak{B})$ such that:

$$V_t \cap K' = \emptyset$$
$$\mathcal{A}_{\bar{t}} \cup \mathcal{B}_{\mathfrak{A}}(K, d) \subseteq K'$$
$$[\![\Delta A]\!]_{\mathfrak{B}} = [\![\Delta A]\!]_{\mathfrak{A}} \setminus V_t \qquad \text{(for all unary } A \in \mathbf{Dyn})$$
$$[\![\oplus \mathrm{Link}]\!]_{\mathfrak{B}} = [\![\oplus \mathrm{Link}]\!]_{\mathfrak{A}} \setminus V_t$$
$$[\![\ominus \mathrm{Link}]\!]_{\mathfrak{B}} = ([\![\ominus \mathrm{Link}]\!]_{\mathfrak{A}} \setminus V_t) \cup ([\![\ominus \mathrm{Link}]\!]_{\mathfrak{A}} \,/\, \mathcal{B}_{\mathfrak{B}}(K', 0))$$

If t is not specified, we say that $\mathfrak{B}$ is *a* (K, d)-sub of $\mathfrak{A}$, and if K, d are not specified, we say that $\mathfrak{B}$ is a sub of $\mathfrak{A}$. ∎

In plain english, in a (K, d)-sub with cleared tree t, we

1. Pick a 0-ball around a protected set. It must contain at least a certain ball and active elements outside of t

2. Make sure t does not touch the protected set

3. Make t inactive

4. Also remove edge deletions that involve t and any edge deletion with one leg in t and the other unprotected.

5. Possibly remove any unprotected element of $\mathrm{dom}(\mathfrak{A})$ in a way compatible with $\trianglelefteq$.

The idea is that gradually removing changes this way lets us keep satisfying a class of formulas on the way down to minimal changes.

If a change (unary predicate change, or edge deletion, or edge addition) is present in $\mathfrak{A}$ but not present in $\mathfrak{B}$, we say that the change has been *cleared*.

Example. Clearing t_a in $\mathfrak{A}$ from figure 4.3 (and picking a minimal K') result in $\mathfrak{B}$, a $(\emptyset, 0)$-sub of $\mathfrak{A}$, seen figure 4.4. All changes from t_a are removed, *including the edge deletion between* b *and* d. The edge deletion is cleared because t_d is not active.

Contrast with transitions $\mathfrak{A}'$ (figure 4.5) and $\mathfrak{B}'$ (figure 4.6). The difference between $\mathfrak{A}$ and $\mathfrak{A}'$ is that d is active in $\mathfrak{A}'$ (because $d \in [\![\Delta A]\!]_{\mathfrak{A}'}$). Just as $\mathfrak{B}$ is a $(\emptyset, 0)$-sub of $\mathfrak{A}$ with cleared tree t_a, $\mathfrak{B}'$ is a $(\emptyset, 0)$-sub of $\mathfrak{A}'$ with cleared tree t_a. Note the difference: since t_d is active in $\mathfrak{A}'$, the link deletion between b and d *has not been cleared*. However, clearing t_d in $\mathfrak{B}'$ would finally clear that link deletion. ∎

Figure 4.5.: $\mathfrak{A}$ with an active d.

Figure 4.6.: $\mathfrak{B}'$, after clearing t_a in $\mathfrak{A}'$.

Definition (Preservation). If $d \geqslant 0$ and $\mathcal{V}$ is a set of variables a formula ϕ is $(\mathcal{V}, d)$-*preserved* whenever for all $\mathfrak{A} : \Theta_n$, $\mu : \mathcal{V} \to \mathrm{dom}(\mathfrak{A})$, $(\mu(\mathcal{V}), d)$-sub $\mathfrak{B}$ of $\mathfrak{A}$ and $\nu : \langle \phi \rangle \setminus \mathcal{V} \to \mathrm{dom}(\mathfrak{B})$, if $\mathfrak{A}, \mu, \nu \models \phi$ then $\mathfrak{B}, \mu, \nu \models \phi$.

∎

Informally, ϕ is $(\mathcal{V}, d)$-preserved if ϕ-satisfaction is closed under $(\mathcal{V}, d)$-sub.

Examples.

- Consider the formula $A^\star(x)$. If $a \in [\![A^\star]\!]_{\mathfrak{A}} \setminus [\![A]\!]_{\mathfrak{A}}$, that remains true for any sub of a as long as a is not cleared. So $A^\star(x)$ is $(\{x\}, 0)$-preserved. In general, postcondition constraints increase the level of protection necessary for a formula to remain preserved.

- Consider $\forall x.\, A(x)$. Since only the precondition is mentioned, taking a sub cannot break satisfaction of the formula. In general, precondition constraints cannot impact preservation.

- Consider $\mathrm{Link}^\star(x, y)$. If a link "appears" between the interpretation of x and y and either of them is cleared, the link creation will be cleared as well. So they must both be protected: the formula is $(\{x, y\}, 0)$-preserved, as well as $(\{x\}, 1)$-preserved and $(\{y\}, 1)$-preserved. In general, positive link constraints should be protected on both sides.

- Consider $\exists x.\, N(x)$, with $N \in \mathbf{Stat}$ a predicate symbol. With the model $\mathfrak{A}$ with $\{a, b\} = \mathrm{dom}(\mathfrak{A})$ and $\{a\} = [\![N]\!]_{\mathfrak{A}}$, removing a yields a $(\emptyset, 0)$-sub $\mathfrak{B}$ which is not a model of $\exists x.\, N(x)$. In general, without precautions, existential quantification can easily make formulas be not-preserved, even if the formula below the $\exists$ does not talk about dynamic properties at all.

∎

4.3. Preservation leads to finitely many canonical patterns

In this section, we show that preserved formula of a certain shape have finitely many canonical patterns.

Overall, preservation leads to finitely many canonical patterns (for minimised formulas) because as one takes subs, one goes down the change order $\lhd$. Eventually all changes are inside a protected area defined as a ball around the interpretation of variables. Therefore there is a bounded number of changes in the minimal models of a preserved formula. Since function graphs in n-BFLs are bounded, balls around elements have bounded size; and therefore corresponding canonical patterns have bounded size.

The next two lemmas are weakening lemmas:

Lemma 4.3.1. If $\mathfrak{B}$ is a (K, d)-sub of $\mathfrak{A}$ and $\mathcal{B}_{\mathfrak{A}}(K', d') \subseteq \mathcal{B}_{\mathfrak{A}}(K, d)$ then $\mathfrak{B}$ is a (K', d')-sub of $\mathfrak{A}$.

Proof. Trivial. $\qquad\qquad\square$

Lemma 4.3.2. If ϕ is $(\mathcal{V}, d)$-preserved, $\mathcal{V} \subseteq \mathcal{V}'$ and $d \leqslant d'$, ϕ is $(\mathcal{V}', d)$-preserved.

Proof. Let X be the variables of $\mathcal{V}$, Y be the variables of $\mathcal{V}' \setminus \mathcal{V}$.

Let $\mathfrak{A} : \Theta_n$. Let $\mu_Y : X \uplus Y \to \mathrm{dom}(\mathfrak{A})$ be any interpretation, and $\mu : X \to \mathrm{dom}(\mathfrak{A}) := \mu_Y \upharpoonright X$. Let $\mathfrak{B}$ be a $(\mathrm{Im}(\mu_Y), d)$-sub of $\mathfrak{A}$, $\nu_Y : \langle \phi \rangle \setminus (X \uplus Y) \to \mathrm{dom}(\mathfrak{B})$ be any interpretation, and $\nu : \langle \phi \rangle \setminus Y \to \mathrm{dom}(\mathfrak{B}) := \nu_Y \uplus (\mu_Y \upharpoonright Y)$. We have $\mu_Y \uplus \nu_Y = \mu \uplus \nu$.

By lemma 4.3.1 and assumption on ϕ, $\mathfrak{B}, \mu, \nu \models \phi$. So $\mathfrak{B}, \mu_Y, \nu_Y \models \phi$. $\qquad\square$

Here we see that taking subs does not break out of the class of n-BFLs:

Lemma 4.3.3. If $\mathfrak{A} : \Theta_n$ and $\mathfrak{B}$ is a sub of $\mathfrak{A}$, then $\mathfrak{B} : \Theta_n$.

Proof. First, we show that Link and Link* are symmetric and functional.

By definition of $\trianglelefteq$, $[\![\mathsf{Link}]\!]_{\mathfrak{B}} = [\![\mathsf{Link}]\!]_{\mathfrak{A}}$, so $[\![\mathsf{Link}]\!]_{\mathfrak{B}}$ is symmetric and functional.

So we need $[\![\Delta\mathsf{Link}]\!]_{\mathfrak{B}}$ symmetric and functional. It is defined by a restriction of $[\![\Delta\mathsf{Link}]\!]_{\mathfrak{A}}$, which is symmetric and functional, to subrelations of $[\![\mathsf{Link}]\!]_{\mathfrak{A}}$ and $[\![\mathsf{Link}^\star]\!]_{\mathfrak{A}}$, which are symmetric and functional. These subrelations are defined by restriction to sets or deletion of sets, so they are symmetric and functional as well. Simple set-theoretic considerations yield the result.

Now we want the loop-free union of $[\![h]\!]_{\mathfrak{B}}$ for $h \in \mathbf{Stat}_{\mathrm{fun}} \setminus \{\mathsf{parent}\}$ to be a tree and $[\![\mathsf{parent}]\!]_{\mathfrak{B}}$ to be the parent function of that tree. By definition of $\trianglelefteq$, $\mathrm{dom}(\mathfrak{B}) \subseteq \mathrm{dom}(\mathfrak{A})$ and $[\![\mathbf{Stat}]\!]_{\mathfrak{B}} = [\![\mathbf{Stat}]\!]_{\mathfrak{A}} \upharpoonright \mathrm{dom}(\mathfrak{B})$, so it suffices to show that for $e \in V_{G_{\mathfrak{A}}}$, with t_e the tree of e in $G_{\mathfrak{A}}$, $e \notin \mathrm{dom}(\mathfrak{B})$ implies $V_{t_e} \cap \mathrm{dom}(\mathfrak{B}) = \emptyset$ (that is, that removing any element removes their entire tree).

If e has no parent or child in t_e, we are done. If there is v parent of e in $G_{\mathfrak{A}}$, there is $f \in \mathfrak{f}$ such that $[\![f]\!]_{\mathfrak{A}}(v) = e$. So $v \in \mathrm{dom}(\mathfrak{B})$ implies $e \in \mathrm{dom}(\mathfrak{B})$. If there is v child of e, then $[\![\mathsf{parent}]\!]_{\mathfrak{A}}(v) = e$, and again $v \in \mathrm{dom}(\mathfrak{B})$ implies $e \in \mathrm{dom}(\mathfrak{B})$. The proof goes by induction on the distance to e. $\qquad\square$

Definition. Let

$$Supp_D := \forall xy. \ (\mathsf{Link}(x, y) \vee \mathsf{Link}(y, x) \vee \bigvee_{A \in \mathbf{Dyn \ unary}} A(x)) \to \mathsf{Has}(x)$$

$$Support_D := Support \wedge Supp_D \wedge Supp_D^\star$$

$\blacksquare$

Definition. $\mathcal{T}_{ns_D} := \mathcal{T}_n \wedge Support_D$ $\qquad\qquad\blacksquare$

This strengthening of support requires that an element with at least some true dynamic property is deemed to be present.

A transition is *dynamically supported* when it satisfies $Support_D$, and a formula is *dynamically supported* when it implies $Support_D$. $Support_D$ is this chapter's version of $Support_0$: a formula stronger than $Support$ necessary for some inductions to go through or preservation properties to hold.

The following lemma shows that by repeatedly taking subs from a dynamically supported transition, we reach a point where all changes are within a small boundary:

Lemma 4.3.4. If $\mathfrak{A} : \Theta_n$ is dynamically supported, $d \geqslant 0$, $K \subseteq \mathrm{dom}(\mathfrak{A})$, there exists a dynamically supported (K, d)-sub $\mathfrak{B}$ of $\mathfrak{A}$ with $\mathfrak{B} \neq \mathfrak{A}$ whenever either

- There is $a \notin \mathcal{B}_\mathfrak{A}(K, d)$ such that $a \in [\![\Delta A]\!]_\mathfrak{A}$ for some $A \in \mathbf{Dyn}$, or such that $(a, b) \in [\![\oplus\mathsf{Link}]\!]_\mathfrak{A}$ for some b.

- There is $a \notin \mathcal{B}_\mathfrak{A}(K, d+1)$ such that $(a, b) \in [\![\ominus\mathsf{Link}]\!]_\mathfrak{A}$ for some b.

Proof. We start by defining the cleared tree t. If t_a (the tree of $G_\mathfrak{A}$ that contains a) is active, let $t := t_a$. Otherwise, by definition of active, there is $b \in \mathrm{dom}(\mathfrak{A})$ with tree $t_b \neq t_a$ and $(a, b) \in [\![\ominus\mathsf{Link}]\!]_\mathfrak{A}$. Let $t := t_b$.

Let $\mathfrak{B}$ be as follows, for $A \in \mathbf{Dyn}$ unary:

$$[\![\oplus\mathsf{Link}]\!]_\mathfrak{B} := [\![\oplus\mathsf{Link}]\!]_\mathfrak{A} \setminus V_t$$

$$[\![\ominus\mathsf{Link}]\!]_\mathfrak{B} := ([\![\ominus\mathsf{Link}]\!]_\mathfrak{A} \setminus V_t) \cup \left([\![\ominus\mathsf{Link}]\!]_\mathfrak{A} / \left(\mathcal{A}_{\bar{t}}(\mathfrak{A}) \cup \mathcal{B}_\mathfrak{A}(K, d) \right) \right)$$

$$[\![\Delta A]\!]_\mathfrak{B} := [\![\Delta A]\!]_\mathfrak{A} \setminus V_t \qquad\qquad [\![\mathbf{Dyn}]\!]_\mathfrak{B} := [\![\mathbf{Dyn}]\!]_\mathfrak{A}$$

which induces $[\![A^\star]\!]_\mathfrak{B}$ for all $A \in \mathbf{Dyn}$. Now let $\mathrm{dom}(\mathfrak{B}) := [\![\mathsf{Has}]\!]_\mathfrak{B} \cup [\![\mathsf{Has}^\star]\!]_\mathfrak{B}$, and $[\![\mathbf{Stat}]\!]_\mathfrak{B} := [\![\mathbf{Stat}]\!]_\mathfrak{A} \restriction \mathrm{dom}(\mathfrak{B})$.

We show that $\mathfrak{B}$ well-defined: it is trivial for the static part; if $b \in \mathrm{dom}(\mathfrak{A})$ is in the underlying set of $[\![A]\!]_\mathfrak{A} = [\![A]\!]_\mathfrak{B}$ for some $A \in \mathbf{Dyn}$, since $\mathfrak{A}$ is dynamically supported, $b \in [\![\mathsf{Has}]\!]_\mathfrak{B} \subseteq \mathrm{dom}(\mathfrak{B})$. Otherwise, suppose b is in the underlying set of $[\![\oplus A]\!]_\mathfrak{A}$. Again by dynamic support, $b \in [\![\oplus\mathsf{Has}]\!]_\mathfrak{A}$. If $b \in V_t$ then $b \notin [\![\mathsf{Has}^\star]\!]_\mathfrak{B}$, so $b \notin \mathrm{dom}(\mathfrak{B})$, otherwise $b \in [\![\mathsf{Has}^\star]\!]_\mathfrak{B}$, so $b \in \mathrm{dom}(\mathfrak{B})$. The same reasoning shows that $\mathfrak{B}$ is dynamically supported.

We show that $\mathfrak{B} \trianglelefteq \mathfrak{A}$. $[\![\mathbf{Dyn}]\!]_\mathfrak{B} = [\![\mathbf{Dyn}]\!]_\mathfrak{A}$ and $[\![\mathbf{Stat}]\!]_\mathfrak{B} = [\![\mathbf{Stat}]\!]_\mathfrak{A} \restriction \mathrm{dom}(\mathfrak{B})$ by definition. Since for $A \in \mathbf{Dyn}$, $[\![\Delta A]\!]_\mathfrak{B}$ is defined by restriction, we have $[\![\Delta\mathbf{Dyn}]\!]_\mathfrak{B} \subseteq [\![\Delta\mathbf{Dyn}]\!]_\mathfrak{A}$. Since $[\![\mathsf{Has}]\!]_\mathfrak{B} = [\![\mathsf{Has}]\!]_\mathfrak{A}$ and $[\![\mathsf{Has}^\star]\!]_\mathfrak{B}$ is defined by restriction of $[\![\mathsf{Has}^\star]\!]_\mathfrak{A}$, $\mathrm{dom}(\mathfrak{B}) \subseteq \mathrm{dom}(\mathfrak{A})$.

By construction, $\mathfrak{B}$ is a (K, d)-sub of $\mathfrak{A}$ (with $K' = \mathcal{A}_{\bar{t}} \cup \mathcal{B}_\mathfrak{A}(K, d)$, which is closed by taking the 0-ball around it).

We show $\mathfrak{B} \neq \mathfrak{A}$.

If t_a is active in $\mathfrak{A}$ then by definition of active, there is $c \in V_t$ (remember that here, $t = t_a$) such that either, for some $A \in \mathbf{Dyn}$ unary, $c \in [\![\Delta A]\!]_\mathfrak{A}$, so $[\![\Delta A]\!]_\mathfrak{A} \neq [\![\Delta A]\!]_\mathfrak{A} \setminus V_t = [\![\Delta A]\!]_\mathfrak{B}$; or there is $d \in \mathrm{dom}(\mathfrak{A})$ such that $(c, d) \in [\![\oplus\mathsf{Link}]\!]_\mathfrak{A}$, so by definition of $[\![\oplus\mathsf{Link}]\!]_\mathfrak{B}$ we get $[\![\Delta\mathsf{Link}^\star]\!]_\mathfrak{B} \neq [\![\Delta\mathsf{Link}^\star]\!]_\mathfrak{A}$.

Otherwise, we have $(a, b) \in [\![\ominus\mathsf{Link}]\!]_\mathfrak{A}$. Since t_a is not active, $a \notin \mathcal{A}_{\bar{t}}(\mathfrak{A})$; and by definition of t, $b \notin \mathcal{A}_{\bar{t}}(\mathfrak{A})$. By assumption $a \notin \mathcal{B}_\mathfrak{A}(K, d+1)$ so $a \notin \mathcal{B}_\mathfrak{A}(K, d)$, and b is at distance at most 1 from a, so $b \notin \mathcal{B}_\mathfrak{A}(K, d)$. We get $[\![\ominus\mathsf{Link}]\!]_\mathfrak{B} \neq [\![\ominus\mathsf{Link}]\!]_\mathfrak{A}$.

$\square$

Using the previous lemma, we obtain that dynamic support is preserved by drawing any protected boundary and clearing a change outside of the boundary.

Lemma 4.3.5. If $\mathfrak{A} : \Theta_n$ is dynamically supported, $d \geqslant 0$, $K \subseteq \mathrm{dom}(\mathfrak{A})$, and there is $A \in \mathbf{Dyn}$, $a \in e \in [\![\Delta A]\!]_\mathfrak{A}$, and $a \notin \mathcal{B}_\mathfrak{A}(K, d+1)$, there is a dynamically supported (K, d)-sub $\mathfrak{B}$ of $\mathfrak{A}$ such that $\mathfrak{B} \neq \mathfrak{A}$.

Proof. Immediate application of lemma 4.3.4 since $a \notin \mathcal{B}_\mathfrak{A}(K, d+1)$ implies that $a \notin \mathcal{B}_\mathfrak{A}(K, d)$. $\square$

We are interested in the size of protected areas, specifically of a ball of radius d around an element. The trees of an n-BFL have a bounded size (height bounded by n, degree bounded by $|\mathbf{Stat}_{\mathrm{fun}}|$) and a bounded number of neighbors (by functionality of Link and $\mathsf{Link}^\star$) so we get an upper bound. Out of convenience, the formulation of the lemma names a witness bounding function: in chapter 7 we will enumerate a number of variables bounded by the size of balls in n-BFLs, and having an explicit bound will simplify notation.

Lemma 4.3.6. There is a function $\ell_{\Theta_n}(d) \in \mathcal{O}(|\mathbf{Stat}_{\mathrm{fun}}|^{2dn^2})$ such that for all $\mathfrak{A} : \Theta_n$ if $a \in \mathrm{dom}(\mathfrak{A})$ and $d \geqslant 0$ then $|\mathcal{B}_\mathfrak{A}(a,d)| \leqslant \ell_{\Theta_n}(d)$.

Proof. Trees of $G_\mathfrak{A}$ have degree at most $|\mathbf{Stat}_{\mathrm{fun}}| - 1$ and height at most n, so they are of order $\leqslant |\mathbf{Stat}_{\mathrm{fun}}|^n$. Since each element in a tree is linked to at most two other elements (one in Link and one in $\mathrm{Link}^\star$), we can bound the number of trees at distance d from an element by the order of a graph of diameter $2d$ and degree at most $s := 2|\mathbf{Stat}_{\mathrm{fun}}|^n$. By [Miller and Sirán, 2005], that number is bounded above by $\frac{s(s-1)^{2d}-2}{s-2} + 1$ if $n > 0$ and $4d + 1$ if $n = 0$. So $|\mathcal{B}_\mathfrak{A}(a,d)| \in \mathcal{O}(|\mathbf{Stat}_{\mathrm{fun}}|^{2dn^2})$ (number of trees $\times$ size of trees). $\square$

Direct application of the previous lemmas shows the importance of preservation:

Theorem 4.3.7. If $\downarrow\phi_1 \vee \ldots \vee \downarrow\phi_m$ is $(\mathcal{V},d)$-preserved and implies $\mathcal{T}_{ns}$, then for all $d' \geqslant 0$, $\mathcal{P}(\bigvee \downarrow\phi_i, d')$ is finite up to iso.

Proof. It suffices to show that the sup of $|\mathrm{dom}(\mathfrak{A})|$ for $(\mathfrak{A}, K) \in \mathcal{P}(\bigvee \phi_i, d')$ is finite.
Let $(\mathfrak{A}, K) \in \mathcal{P}(\bigvee \phi_i, d')$. For some μ and $1 \leqslant i \leqslant m$, $\mathfrak{A}, \mu \vDash \downarrow\phi_i$. By definition of $\mathcal{P}$,

$$\mathrm{dom}(\mathfrak{A}) = \mathcal{B}_\mathfrak{A}^{\Delta,\mu}(d'+1) = \mathcal{B}_\mathfrak{A}(\mathrm{Im}(\mu), d'+1) \cup \mathcal{B}_\mathfrak{A}(\llbracket \Delta \rrbracket_\mathfrak{A}, d'+1)$$

Since $\mathfrak{A} : \Theta_n$ is supported, by lemma 4.3.5 $\llbracket \Delta \rrbracket_\mathfrak{A} \subseteq \mathcal{B}_\mathfrak{A}(\mu(\mathcal{V}), d+1)$. Moreover $\mu(\mathcal{V}) \subseteq \mathrm{Im}(\mu)$, so

$$\mathcal{B}_\mathfrak{A}(\llbracket \Delta \rrbracket_\mathfrak{A}, d'+1) \subseteq \mathcal{B}_\mathfrak{A}(\mathrm{Im}(\mu), d'+d+2)$$

therefore

$$\mathrm{dom}(\mathfrak{A}) \subseteq \mathcal{B}_\mathfrak{A}(\mathrm{Im}(\mu), d'+d+2)$$

So by lemma 4.3.6, $|\mathrm{dom}(\mathfrak{A})|$ is bounded by a function of $|\langle\bigvee \phi_i\rangle|$, Θ, d and d'. $\square$

4.4. Building preserved formulas with $\vdash_\lhd$

In this section, we show how to construct preserved formula by introduced a set of "typing rules" we call $\vdash_\lhd$. First, we present the ruls and give the intuition behind them. Then we prove that $\vdash_\lhd$ is correct for preservation. Recall that preservation is a property parameterised by a set of variables $\mathcal{V}$ and some $d \geqslant 0$. Similarly, $\vdash_\lhd$ has a context of the form $\mathcal{V}, d$.

4.4.1. The rule system $\vdash_\lhd$

Note that in the following rules, the context are parameters of the preservation property. In the sequel, we will show that if $\mathcal{V}, d \vdash_\lhd \phi$ then ϕ is $(\mathcal{V},d)$-preserved.

Definition. If x, y are two variables and $\mathcal{V}$ is a set of variables, $\mathcal{V}_{\cdot x}^{+y} := (\mathcal{V} \cup \{y\}) \setminus \{x\}$. ∎

In the rules, $\alpha(x,y)$ is any guard. The preservation property is designed to compose well. The following deduction rules all produce preserved formulas:

$$\frac{\mathrm{L} \in \mathcal{L}_{\mathrm{Dyn}^\star}}{\langle \mathrm{L}\rangle\,;0 \vdash_\lhd \mathrm{L}}\ \textsc{Dynamic}_\lhd \qquad\qquad \frac{\mathrm{L} \in \mathcal{L}_{\overline{\mathrm{Dyn}},\mathrm{Stat}}}{\emptyset\,;0 \vdash_\lhd \mathrm{L}}\ \textsc{Static}_\lhd$$

$$\frac{\mathcal{V} \subseteq \mathcal{V}' \quad \mathcal{V}\,;d \vdash_\lhd \phi}{d \leqslant d' \quad \mathcal{V}'\,;d' \vdash_\lhd \phi}\ \textsc{Weak}_\lhd \qquad\qquad \oplus \in \{\wedge,\vee\}\ \frac{\mathcal{V}\,;d \vdash_\lhd \phi_1 \quad \mathcal{V}\,;d \vdash_\lhd \phi_2}{\mathcal{V}\,;d \vdash_\lhd \phi \oplus \phi_1}\ \textsc{Bool}_\lhd$$

$$\frac{\mathcal{V}\,;d \vdash_{\lhd} \phi}{\langle\phi\rangle\,;d+1 \vdash_{\lhd}\downarrow(\phi \wedge Support_D \wedge Symlink)}\ \text{Circum}_{\lhd} \qquad \frac{\{x\}\,;0 \vdash_{\lhd} \phi^{\star}}{\emptyset\,;0 \vdash_{\lhd} \phi^{\circ} \wedge \phi^{\star}}\ \text{Invariant}_{\lhd}$$

$$x \notin \mathcal{V}\ \frac{\mathcal{V}\,;d \vdash_{\lhd} \phi}{\mathcal{V}\,;d \vdash_{\lhd} \forall x.\ \phi}\ \forall_{\lhd} \qquad \frac{\mathcal{V}\,;d \vdash_{\lhd} \phi}{\mathcal{V}^{+y}_{-x}\,;d+|\{x\}\cap\mathcal{V}| \vdash_{\lhd} \alpha(x,y) \rightarrow \phi}\ \text{Guard}_{\lhd}$$

$$\frac{\mathcal{V}\,;d \vdash_{\lhd} \phi}{\mathcal{V}^{+y}_{-x}\,;d+1 \vdash_{\lhd} \exists x.\ \alpha(x,y) \wedge \phi}\ \exists\text{-Guard}_{\lhd}$$

Discussion

The rules $\text{Static}_{\lhd}$ and $\text{Dynamic}_{\lhd}$ introduce literals. Equality literals, literals on the precondition and static literals do not need any protected area to remain satisfied under sub since a sub only removes changes by altering the *postcondition*. On the other hand, any variable mentioned in a literal from $\mathcal{L}_{\mathbf{Dyn}^{\star}}$ should a priori be protected (the rule can actually be slightly relaxed).

The rule $\text{Weak}_{\lhd}$ directly implements the weakening of lemma 4.3.2: if ϕ-satisfaction is closed under taking subs for some protected area, it's always safe to grow that area.

The two rules packed in $\text{Bool}_{\lhd}$ show that positive boolean operations are well-behaved with regards to preservation.

At a high level, the rule $\text{Circum}_{\lhd}$ is intuitive: if a set is closed for subs, its $\unlhd$-minimal set of $\unlhd$ is as well (remember that the sub relation is finer than $\unlhd$). The formulas $Support_D$ and $Symlink$ are added in because they are not provable (unlike the rest of $\mathcal{T}_{ns}$). Also, while not necessary for the lemmas and theorems of this chapter, we grow the context in $\text{Circum}_{\lhd}$ (adding 1 to d, adding all free variables to $\mathcal{V}$) so that we can prove syntactic properties about $\vdash_{\lhd}$-provable formulas in chapter 7.

More discussion about $Support_D$ and $Symlink$ is warranted. $Support_D$ is not provable, and is indeed not preserved. This is simply because any "appearing" element (so an element in some $[\![\mathsf{Has}^{\star}]\!]_{\mathfrak{A}} \setminus [\![\mathsf{Has}]\!]_{\mathfrak{A}}$) ceases to be supported once it is selected for clearing (when taking a sub). However, put under $\downarrow$, $Support_D$ is such that there are no proper subs at all, so the property of preservation becomes vacuously true. As for $Symlink$, that formula actually *is* $(\emptyset, 0)$-preserved, but the deduction rules cannot prove it. In the last chapter we mention some ideas on how to make $Symlink$ provable.

The rule $\text{Invariant}_{\lhd}$ says that a postcondition property that only requires protecting changes local to a single tree requires no protection at all once that property is extended to the precondition.

Example. Take the formula $A^{\star}(x) \vee B^{\star}(x)$. If $\mathfrak{A}$ is an n-BFL and $a \in \text{dom}(\mathfrak{A})$ such that $\mathfrak{A}, x \mapsto a \vDash A^{\star}(x) \vee B^{\star}(x)$ then any sub of $\mathfrak{A}$ which protects a still satisfies the formula. Now with the formula

$$F = (A(x) \vee B(x)) \wedge (A^{\star}(x) \vee B^{\star}(x))$$

no protection is required: if the changes local to a are cleared, they are cleared all at once. Since $A(a) \vee B(a)$ is satisfied in the precondition, a complete absence of changes means $A^{\star}(a) \vee B^{\star}(a)$ will be satisfied in the postcondition.

We illustrate the effect in figure 4.7. On the top left, a transition satisfies $\forall A^{\star}(x) \vee B^{\star}(x)$. But after clearing a (bottom left), the new transition does not satisfy $\forall A^{\star}(x) \vee B^{\star}(x)$. Compare with the right side: on the top right, a transition satisfies $\forall F$, and the transition that results from clearing a still satisfies $\forall F$. ∎

$$\forall A^{\star} \vee B^{\star} \dashv \quad a \quad \blacktriangleright \quad a^{A} \qquad \qquad a^{B} \quad \blacktriangleright \quad a^{AB} \qquad \vDash \forall (A \vee B)$$
$$\wedge (A^{\star} \vee B^{\star})$$

$$\underline{\quad clear\ a \quad} \qquad \qquad \underline{\quad clear\ a \quad}$$

$$\textbf{✗} \ \forall A^{\star} \vee B^{\star} \nvdash \quad a \quad \blacktriangleright \quad a \qquad \qquad a^{B} \quad \blacktriangleright \quad a^{B} \qquad \vDash \forall (A \vee B) \quad \textbf{✔}$$
$$\wedge (A^{\star} \vee B^{\star})$$

Figure 4.7.: Illustrating the effect of INVARIANT$_{\lhd}$

In tandem with INVARIANT$_{\lhd}$, the rule $\forall_{\lhd}$ can build universal sentences: once a variable x has been removed from the context $\mathcal{V}$ using INVARIANT$_{\lhd}$, x can be universally quantify using $\forall_{\lhd}$ (note the proviso $x \notin \mathcal{V}$).

The rule GUARD$_{\lhd}$ combines two ideas. The first idea is that if we know that the interpretation of two variables x and y are linked, instead of having both x and y in the context, we can keep only y in the context and extend the protected area by increasing d. Again, this means we can universally quantify on x. The second idea is that if $x \notin \mathcal{V}$, we actually only need to protect y and do not need include the interpretation of x in the protection radius. This is thanks to the asymmetry in the definition of subs:

Example. Take the formula $\neg \mathsf{Link}^{\star}(x, y)$. It is equivalent to a formula of the form $\alpha(x, y) \to \phi$ (where ϕ is $x \neq x$, for instance). Going just by DYNAMIC$_{\lhd}$, the context should be $\{x, y\}, 0$. But by GUARD$_{\lhd}$, the context can be just $\{y\}, 0$; that is, $\neg \mathsf{Link}^{\star}(x, y)$ is $(\{y\}, 0)$-preserved. This is due to the asymmetry in the definition of subs: link deletions are not cleared when they touch an element of the protected set. However, note that the positive literal $\mathsf{Link}^{\star}(x, y)$ is *not* $(\{y\}, 0)$-preserved, following the definition of sub which removes any link addition which touches a cleared tree. One may wonder if the same trick used for link deletions could also be used for link additions. We were not able to construct such a symmetric notion of sub; in particular the definition of *active* becomes difficult. In addition we think there is a more fundamental problem, not apparent in the current version of $\vdash_{\lhd}$. As mentioned earlier, *Symlink* is not provable even though it is preserved. With more powerful rules capable of proving formulas like *Symlink*, we foresee that a symmetric version of sub would quickly lead to unbounded changes. ∎

Finally, the rule $\exists$-GUARD$_{\lhd}$ acts both as a guard and as a way to add existential quantifiers. The guard part works as in GUARD$_{\lhd}$, except the distance is always increased by 1 because taking a sub can remove an element.

Design note. We think it is possible to build deduction rules with a context of the form $\mathcal{V}, q$, with $q \geqslant 0$, based on a different definition of preservation – instead of taking a ball of radius d around the interpretation of $\mathcal{V}$, one simply asks that it suffices to protect q trees at most in addition those given by $\mathcal{V}$ (which q trees may depend on the interpretation of $\mathcal{V}$). Under this definition, it should be possible to have a rule with unguarded existential quantification. ∎

4.4.2. $\vdash_{\lhd}$-provable formulas are preserved

The following lemmas prepare for theorem 4.4.5.

The first one shows that binary dynamic atoms can be treated as unary when the link loops back:

Lemma 4.4.1. If $\mathcal{V}$ is a set of variables, $d \geqslant 0$, and t is a term such that $\langle t \rangle \notin \mathcal{V}$ then the formula $\mathsf{Link}^{\star}(t, t)$ is not $(\mathcal{V}, d)$-preserved.

Proof. Take a model $\mathfrak{M}$ such that $\mathrm{dom}(\mathfrak{M}) := \{e\}$, an interpretation μ such that $[\![t]\!]_{\mathfrak{M}, \mu} = e$, no $[\![\mathsf{Link}]\!]_{\mathfrak{M}} := \emptyset$ and $[\![\mathsf{Link}^{\star}]\!]_{\mathfrak{M}} := \{(e, e)\}$.

We have $\mathfrak{M}, \mu \models \mathsf{Link}^{\star}(t, t)$, but a $(\mathcal{V}, d)$-sub of $\mathfrak{M}$ with cleared tree t_e (possible since there is no other variable than $\langle t \rangle$ in $\mathsf{Link}^{\star}(t, t)$) does not satisfy $\mathsf{Link}^{\star}(t, t)$. $\square$

This lemma shows that for negative literals on links, it is actually enough to protect only one of the ends of the link. Note that $\vdash_{\lhd}$ cannot prove that directly; but the rule $\mathrm{G{\small UARD}}_{\lhd}$ exploits that fact.

Lemma 4.4.2. If t, u are terms then $\neg\mathsf{Link}^{\star}(t, u)$ and $\neg\mathsf{Link}^{\star}(u, t)$ are $(\{u\}; 0)$-preserved.

Proof. Take $\mathfrak{A}$ and appropriate μ and $\mathfrak{B}$ a $(\{u\}, 0)$-sub of $\mathfrak{A}$ named $\mathfrak{B}$. By definition of a sub, no edge deletion touching the tree of u may be cleared since $[\![u]\!]_{\mathfrak{A}, \mu} \in \mathcal{B}_{\mathfrak{A}}([\![\langle u \rangle]\!]_{\mathfrak{A}, \mu}, 0)$. $\square$

The next lemma shows that pre formula do not look at $\mathsf{Dyn}^{\star}$ at all.

Lemma 4.4.3. If ϕ is a pre formula, $\mathfrak{A}, \mu, \nu \models \phi$, $\mathfrak{B} : \Theta$, $\mathrm{Im}(\mu) \cup \mathrm{Im}(\nu) \subseteq \mathrm{dom}(\mathfrak{B})$, $[\![\mathbf{Dyn}]\!]_{\mathfrak{B}} = [\![\mathbf{Dyn}]\!]_{\mathfrak{A}}$ and $[\![\mathbf{Stat}]\!]_{\mathfrak{B}} = [\![\mathbf{Stat}]\!]_{\mathfrak{A}} \restriction \mathrm{dom}(\mathfrak{B})$ then $\mathfrak{B}, \mu, \nu \models \phi$.

Proof. Trivial by induction on the structure of ϕ. $\square$

This lemma shows that $\phi^{\star}$ is the syntactic counterpart of the mirror operation on pre formulas.

Lemma 4.4.4. If ϕ is a pre formula and $\mathfrak{A}, \mu, \nu, \models \phi$ then $\mathfrak{A}^{-1}, \mu, \nu \models \phi^{\star}$

Proof. Trivial by induction on the structure of ϕ. $\square$

We now show the correctness of $\vdash_{\lhd}$. The proof is by case analysis; no induction on the derivation is required. The interesting cases are $\mathrm{G{\small UARD}}_{\lhd}$ and $\mathrm{I{\small NVARIANT}}_{\lhd}$.

For $\mathrm{G{\small UARD}}_{\lhd}$, the trick is lemma 4.4.2, which shows that a negative guard $\alpha(x, y)$ only needs to protect one of its variables. Note that this is thanks to the asymmetry between link creation and link deletion in the definition of a sub. As a result, a formula of the form $\alpha(x, y) \to \phi$ may not need to protect both x and y if ϕ does not require it.

The case of $\mathrm{I{\small NVARIANT}}_{\lhd}$ is illustrated to help the reader. The asymmetry in the definition of a sub is again on display: we see no element needs to be protected from clearance but first clearing the only tree that *could* need protection (the tree of the interpretation of x) and then taking the mirror of the structure. Any remaining change involving the tree of interpretation of x is a link addition, and due to the way a sub is defined, any remaining link creation can be cleared *without* explicitly clearing that tree, but just by clearing its neighbors.

Theorem 4.4.5. If $\mathcal{V}; d \vdash_{\lhd} \phi$, then ϕ is $(\mathcal{V}, d)$-preserved.

Proof. By case analysis.

W{\small EAK}$_{\lhd}$ Immediate by lemma 4.3.2.

BOOL$_\lhd$ Let $\mathfrak{A} : \Theta_n$, $\mu : \langle \mathcal{V} \rangle \to \mathrm{dom}(\mathfrak{A})$. Let $\mathfrak{B}$ be a $(\mu(\mathcal{V}), d)$-sub of $\mathfrak{A}$. Let $\nu : \langle \phi \rangle \setminus \langle \mathcal{V} \rangle \to \mathrm{dom}(\mathfrak{B})$ be such that $A, \mu, \nu \models \phi_1 \oplus \phi_2$.

If $\oplus = \wedge$, by assumption on ϕ_1 and ϕ_2, we have $\mathfrak{B}, \mu, \nu \models \phi_1 \wedge \phi_2$.

If $\oplus = \vee$, for $i = 1, 2$, if $\mathfrak{A}, \mu, \nu \models \phi_i$ then by assumption on ϕ_i, $\mathfrak{B}, \mu, \nu \models \phi_i$. Either way $\mathfrak{B}, \mu, \nu \models \phi_1 \vee \phi_2$.

GUARD$_\lhd$ Let $d' := d + |\{y\} \cap \mathcal{V}|$. Let $\mu : \mathcal{V}_{-x}^{+y} \to \mathrm{dom}(\mathfrak{A})$, and $\mathfrak{B}$ be a $(\llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu}, d')$-sub of $\mathfrak{A}$, and $\nu : \langle \alpha(x, y) \to \phi \rangle \setminus \mathcal{V}_{-x}^{+y} \to \mathrm{dom}(\mathfrak{B})$ be such that $\mathfrak{A}, \mu, \nu \models \alpha(x, y) \to \phi$.

Suppose $\mathfrak{B}, \mu, \nu \models \alpha(x, y)$. We show $\mathfrak{A}, \mu, \nu \models \alpha(x, y)$ by showing that $\neg\alpha(x, y)$ is $(\mathcal{V}_{-x}^{+y}, d')$-preserved: if α is equality or Link, this is trivial by definition of a sub. Otherwise, $\alpha(x, y)$ is of the form $\mathrm{Link}^\star(t, u)$ with $\{x\} = \langle t \rangle$ or $\{x\} = \langle u \rangle$; so by lemma 4.4.2, $\neg\alpha(x, y)$ is preserved under $\{y\}; 0$, so by lemma 4.3.2, it is also preserved under $\mathcal{V}_{-x}^{+y}; d'$. Now we know $\mathfrak{A}, \mu, \nu \vDash \alpha(x, y)$, so $\mathfrak{A}, \mu, \nu \vDash \phi$. To get $\mathfrak{B}, \mu, \nu \vDash \phi$, we only have to show that $\mathfrak{B}$ is a $((\mu \uplus \nu)(\mathcal{V}), d)$-sub of $\mathfrak{A}$. We consider whether $y \in \mathcal{V}$:

If $y \in \mathcal{V}$ then $d' = d + 1$. By lemma 4.3.1, we only need $\mathcal{B}_\mathfrak{A}((\mu \uplus \nu)(\mathcal{V}), d) \subseteq \mathcal{B}_\mathfrak{A}(\llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu}, d + 1)$. Since $\mathcal{V} \setminus \mathcal{V}_{-x}^{+y} \subseteq \{x\}$, it suffices to show $\mathcal{B}_\mathfrak{A}(\nu(x), d) \subseteq \mathcal{B}_\mathfrak{A}(\llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu, \nu}, d + 1)$. Since $\mathfrak{A}, \mu, \nu \vDash \alpha(x, y)$, $\delta_\mathfrak{A}(\nu(x), \mu(y)) \leqslant 1$, which gives us $\mathcal{B}_\mathfrak{A}(\nu(x), d) \subseteq \mathcal{B}_\mathfrak{A}(\mu(y), d + 1)$. Since $\mu(y) \in \llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu}$, we are done.

If $y \notin \mathcal{V}$ then $\mathcal{V} \subseteq \mathcal{V}_{-x}^{+y}$ and $d' = d$, so we can immediately apply lemma 4.3.2, to get $\mathfrak{B}, \mu, \nu \models \phi$.

$\forall_\lhd$ Let $\mathfrak{A} : \Theta_n$ and $\mu : \langle \mathcal{V} \rangle \to \mathrm{dom}(\mathfrak{A})$. Let $\mathfrak{B}$ be a $(\mu(\mathcal{V}), d)$-sub of $\mathfrak{A}$, and $\nu : \langle \phi \rangle \setminus \langle \mathcal{V} \rangle \to \mathrm{dom}(\mathfrak{B})$ be such that $\mathfrak{A}, \mu, \nu \models \forall y. \phi$,

For all $a \in \mathrm{dom}(\mathfrak{A})$, $\mathfrak{A}, \mu, \nu, y \mapsto a \models \phi$. Since $\mathrm{dom}(\mathfrak{B}) \subseteq \mathrm{dom}(\mathfrak{A})$, for all $b \in \mathrm{dom}(\mathfrak{B})$, by assumption on ϕ we have $\mathfrak{B}, \mu, \nu, y \mapsto b \models \phi$. So $\mathfrak{B}, \mu, \nu \models \forall y. \phi$.

$\exists$-GUARD$_\lhd$ Let $\mu : \mathcal{V}_{-x}^{+y} \to \mathrm{dom}(\mathfrak{A})$, and $\mathfrak{B}$ be a $(\llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu}, d + 1)$-sub of $\mathfrak{A}$, and $\nu : \langle \alpha(x, y) \wedge \phi \rangle \setminus \mathcal{V}_{-x}^{+y} \setminus \{x\} \to \mathrm{dom}(\mathfrak{B})$, $a \in \mathrm{dom}(\mathfrak{A})$ be such that $\mathfrak{A}, \mu, \nu, x \mapsto a \models \alpha(x, y) \wedge \phi$. We show $a \in \mathrm{dom}(\mathfrak{B})$ and $\mathfrak{B}, \mu, \nu, x \mapsto a \models \alpha(x, y) \wedge \phi$. Note how we cannot assume $a \in \mathrm{dom}(\mathfrak{B})$ because of the existential quantifier we are about to add, and that since $x \neq y$ we have $y \in \mathcal{V}_{-x}^{+y}$.

Let $b := \mu(y)$. First, note that since $\delta_\mathfrak{A}(a, b) \leqslant 1$ and $y \in \mathcal{V}_{-x}^{+y}$, we have $\mathcal{B}_\mathfrak{A}(a, d) \subseteq \mathcal{B}_\mathfrak{A}(b, d + 1) \subseteq \mathcal{B}_\mathfrak{A}(\llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu}, d + 1)$. So, by definition of subs, $a \in \mathrm{dom}(\mathfrak{B})$, and any link between t_a and t_b is preserved in $\mathfrak{B}$, so $\mathfrak{B}, \mu, \nu, x \mapsto a \models \alpha(x, y)$.
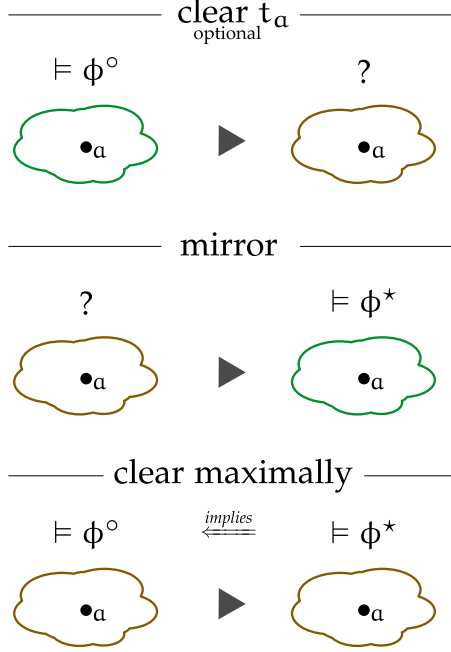
For ϕ, as in the case of GUARD$_\lhd$, we need $\mathcal{B}_\mathfrak{A}((\mu \uplus x \mapsto a)(\mathcal{V}), d) \subseteq \mathcal{B}_\mathfrak{A}(\llbracket \mathcal{V}_{-x}^{+y} \rrbracket_{\mathfrak{A}, \mu}, d + 1)$ to apply lemma 4.3.1. Since $y \in \mathcal{V}_{-x}^{+y}$, it would suffice to show $\mathcal{B}_\mathfrak{A}(a, d) \subseteq \mathcal{B}_\mathfrak{A}(b, d + 1)$, which we just established. So $\mathfrak{B}, \mu, \nu, x \mapsto a \models \phi$.

It immediately follows that $\mathfrak{B}, \mu, \nu \models \exists x. \alpha(x, y) \wedge \phi$ whenever $\mathfrak{A}, \mu, \nu$ does.

INVARIANT$_\lhd$ We know $\{x\}, 0 \vdash_\lhd \phi^\star$. Let $\nu : \langle \phi \rangle \setminus \{x\} \to \mathrm{dom}(\mathfrak{A})$, $\mathfrak{A}, x \mapsto a, \nu \models \phi^\circ \wedge \phi^\star$ and $\mathfrak{B}$ be a sub of $\mathfrak{A}$ with cleared tree t.

Since ϕ is a pre formula, by lemma 4.4.3, $\mathfrak{B}, x \mapsto a, \nu \models \phi^\circ$.

Case 1 : If $a \in V_t$ or if t_a is not active in $\mathfrak{A}$, we have that no element of V_{t_a} occurs in any $[\![\Delta A]\!]_\mathfrak{B}$, except possibly in edge deletions involving t_a and another tree adjacent to t_a. In the first case it is because $V_{t_a} = V_t$ has just been made inactive by going from $\mathfrak{A}$ to $\mathfrak{B}$, in the second case it is by definition of active. We want $\mathfrak{B}, x \mapsto a, \nu \models \phi^\star$. The proof will work by flipping $\mathbf{Dyn}$ and $\mathbf{Dyn}^\star$ in $\mathfrak{B}$. Any remaining edge deletion that touches t_a will become an edge addition and so will become clearable by taking subs. Then, once all changes have been removed, we can deduce satisfaction of $\phi^\star$ from the satisfaction of ϕ°. Consider $\mathfrak{B}^{-1}$. Since $\mathfrak{B}, x \mapsto a, \nu \models \phi^\circ$, by lemma 4.4.4. $\mathfrak{B}^{-1}, x \mapsto a, \nu \models \phi^\star$.

Now, for any $\mathfrak{C} \trianglelefteq \mathfrak{B}^{-1}$ such that $\mathfrak{C}, x \mapsto a, \nu \models \phi^\star$, we show that if $[\![\Delta]\!]_\mathfrak{C} \neq \emptyset$ then there is $\mathfrak{D} \triangleleft \mathfrak{C}$ such that $\mathfrak{D}, x \mapsto a, \nu \models \phi^\star$ and $\mathfrak{D}$ has strictly fewer changes than $\mathfrak{C}$.

- If there is an active tree t different from t_a, clear it (that is, let $\mathfrak{D}$ be a sub of $\mathfrak{C}$ with cleared tree t and minimal K'). By assumption on ϕ, $\mathfrak{D}, x \mapsto a, \nu \models \phi^\star$, and clearing an active tree always results in strictly fewer changes by definition.

- Otherwise, note that since t_a is not active in $\mathfrak{B}$, the only changes in $\mathfrak{C}$ that can involve an element $e \in V_{t_a}$ are an edge addition between e and an element $e' \notin V_{t_a}$. So if there is no tree $t \neq t_a$ active in $\mathfrak{C}$, then t_a is not active in $\mathfrak{C}$.

  Therefore no tree is active in $\mathfrak{C}$. So the only possible changes in $\mathfrak{C}$ are removal of some (b, c) edge, with t_a, t_b, t_c pairwise distinct. Let $\mathfrak{D}$ be a sub of $\mathfrak{C}$ with cleared tree t_b and minimal K'. Since t_c is not active and is different from t_a, $\mathfrak{D}$ contains strictly fewer changes than $\mathfrak{C}$. Again by assumption on ϕ, $\mathfrak{D}, x \mapsto a, \nu \models \phi^\star$.

Using the above statement, by induction on the product ordering on each $[\![\Delta A]\!]_{\mathfrak{B}^{-1}}$, there is $\mathfrak{C} \trianglelefteq \mathfrak{B}^{-1}$ such that $[\![\Delta]\!]_\mathfrak{C} = \emptyset$ and $\mathfrak{C}, x \mapsto a, \nu \models \phi^\star$.
Since $[\![A^\star]\!]_\mathfrak{C} = [\![A]\!]_\mathfrak{C}$ for every $A \in \mathbf{Dyn}$, $\mathfrak{C}, x \mapsto a, \nu \models \phi^\circ$.
Since $\mathfrak{C} \trianglelefteq \mathfrak{B}^{-1}$, $[\![A]\!]_\mathfrak{C} = [\![A]\!]_{\mathfrak{B}^{-1}}$ for every $A \in \mathbf{Dyn}$. So $\mathfrak{B}^{-1}, x \mapsto a, \nu \models \phi^\circ$.
Since $[\![A]\!]_{\mathfrak{B}^{-1}} = [\![A^\star]\!]_\mathfrak{B}$ for every $A \in \mathbf{Dyn}$, $\mathfrak{B}, x \mapsto a, \nu \models \phi^\star$.
The side figure illustrates case 1.
Case 2 : Otherwise, $a \notin V_t$ and t_a is active. If t does not share an edge with t_a, $\mathfrak{B}$ is a $(\{a\}, 0)$-sub of $\mathfrak{A}$. If it does, since t_a is active and different from t, it is in $\mathcal{A}_{\bar{t}}(\mathfrak{A})$, so, any edge deletion between t and t_a is preserved, so $\mathfrak{B}$ is still a $(\{a\}, 0)$-sub of $\mathfrak{A}$. Since $\mathfrak{A}, x \mapsto a, \nu \models \phi^\star$, we have $\mathfrak{B}, x \mapsto a, \nu \models \phi^\star$.

STATIC$_\triangleleft$ Trivial since equality, preconditions and static predicates do not change by taking a sub.

DYNAMIC$_{\lhd}$ Unary case: let $A \in \mathbf{Dyn}$ and take a well-defined $[\![t]\!]_{\mathfrak{A},\mu}$ for some $t,\mathfrak{A},\mu$. By definition of preservation, its may not be cleared. So in any $([\![\langle t \rangle]\!]_{\mathfrak{A},\mu},0)$-sub $\mathfrak{B}$ of $\mathfrak{A}$, $[\![t]\!]_{\mathfrak{A},\mu} \in [\![A^\star]\!]_{\mathfrak{A}}$ iff $[\![t]\!]_{\mathfrak{B},\mu} \in [\![A^\star]\!]_{\mathfrak{B}}$.

Binary case works the same way: take well defined $[\![t]\!]_{\mathfrak{A},\mu}$ and $[\![u]\!]_{\mathfrak{A},\mu}$. No edge with both ends in $\mathcal{B}_{\mathfrak{A}}([\![\langle t \rangle \cup \langle u \rangle]\!],0)$ may be cleared, so $[\![(t,u)]\!]_{\mathfrak{A},\mu} \in [\![Link^\star]\!]_{\mathfrak{A}}$ iff $[\![(t,u)]\!]_{\mathfrak{B},\mu} \in [\![Link^\star]\!]_{\mathfrak{B}}$.

CIRCUM$_{\lhd}$ We show that $\downarrow(\phi \wedge Support_{\mathrm{D}} \wedge Symlink)$ is $(\mathcal{V},d)$-preserved, the rest is by WEAK$_{\lhd}$.

Take some $\mathfrak{A} : \Theta_n$ such that $\mathfrak{A},\mu \vDash \phi$, let $K := \mu(\mathcal{V})$, and take a $\mathcal{B}_{\mathfrak{A}}(K,d)$-sub $\mathfrak{B}$ of $\mathfrak{A}$. By assumption $\mathfrak{B},\mu \vDash \phi$. By lemma 4.3.3, $\mathfrak{B} : \Theta_n$, so $\mathfrak{B} \vDash Symlink$. Now suppose $\mathfrak{B}$ is not dynamically supported. In that case, the cleared tree t contains an element $a \in [\![\oplus Has]\!]_{\mathfrak{A}}$, and that element is outside of $\mathcal{B}_{\mathfrak{A}}(K,d)$, by definition of $\mathfrak{B}$. So by lemma 4.3.4, there is $\mathfrak{B}'$, a (K,d)-sub of $\mathfrak{A}$ which is dynamically supported and such that $\mathfrak{B}' \neq \mathfrak{A}$. Again by 4.3.3 $\mathfrak{B}' \vDash Symlink$, and by assumption, $\mathfrak{B}',\mu \vDash \phi$. By definition of a sub, $\mathfrak{B}' \unlhd \mathfrak{A}$, so $\mathfrak{A}$ was not minimal. Absurd. $\qquad \square$

5. A safe program syntax

In this chapter we introduce a syntax for logic programs. Interpreted as formulas, these programs are guaranteed to have their run semantics coincide with the execution of a finite set of rules.

The shape of a program is a sequence of local "reactions" followed by global "invariants". The reactions should be seen as the successor to graph-rewriting rules: local pieces of knowledge that say what mechanistic interactions can happen in the cell. The invariants are universal formulas. They enforce coherence or encoding constraints that should be true at all times during a run.

To show that all programs can be seen as a finite set of rules, we show that the formula induced by a program is both $\vdash_{\!\rho}$ and $\vdash_{\!\lhd}$-provable (up to logical equivalence). In the case of invariants, this requires a small amount of syntactic gymnastics to reach a clausal form where every clause is great.

Notation. In the sequel, if ϕ is a formula:

> $\cdot \vdash_{\!\rho} \phi$ stands for "there is a context Γ and a formula $\psi \equiv \phi$ such that $\Gamma \vdash_{\!\rho} \psi$"
>
> $\cdot \vdash_{\!\lhd} \phi$ stands for "there is a context $\mathcal{V}, d$ and a formula $\psi \equiv \phi$ such that $\mathcal{V}, d \vdash_{\!\lhd} \psi$"

■

5.1. Programs P and formulas ϕ_P

In this section we introduce the syntax of programs. Here is an example program :

Definition. A *program* is of the following form:

$$\theta_1 \, ; \, \ldots \, ; \, \theta_k \quad \texttt{ensure } I_1,\ldots,I_{k'} \quad \texttt{maintain } I_{k'+1},\ldots,I_{k''}$$

where the (`ensure ...`) and (`maintain ...`) parts are optional, and the θ_i are formulas generated by the following grammar:

$$\theta ::= L \mid \forall x.\ \alpha(x,y) \to \theta \mid \exists x.\ \alpha(x,y) \wedge \theta \mid \theta \wedge \theta \mid \theta \vee \theta \qquad \text{for all } L \in \mathcal{L}, x, y, \alpha(x,y)$$

and the I_j are pre, quantifier-free formulas of the form

$$I ::= \alpha_1(x,y_1) \to \ldots \alpha_q(x,y_q) \to \psi \qquad\qquad q \geqslant 0$$

such that $\langle \psi \rangle \subseteq \{x, y_1, \ldots, y_q\}$ and for every dynamic literal L in ψ, $\langle L \rangle = \{x\}$.

■

Before we explain how a program corresponds to a formula on transitions, we give an example of a simple program:

Example.

$$\oplus\mathsf{Link}(x.s, y.t) \wedge \mathsf{E}(x) \; ; \; \mathsf{Link}(x.s, y.t) \wedge \oplus\mathsf{Active}(x)$$

$$\texttt{ensure}$$

$$\mathsf{Link}(z.s, w.t) \to \mathsf{E}(z) \to (\mathsf{E}(w) \wedge \mathsf{Open}(z.s))$$

where $x.s, y.t, \ldots$ are syntactic sugar for terms like $s(x)$. This program has two rules. The first can add a Link between two s and t sites if the parent of the s site is an Enzyme. The second can make an element Active if its s site is linked to a t site.

In addition, there is an invariant which says that for all elements z, w linked through their respective sites s and t, if z is an Enzyme then w must also be an Enzyme and the site s of z must be Open. ∎

We will now define the formulas induced by programs. First, we add some program P as an implicit paramter:

Implicit parameter $\mathsf{P} := \theta_1; \ldots; \theta_k \; \texttt{ensure} \; I_1, \ldots, I_{k'} \; \texttt{maintain} \; I_{k'+1}, \ldots, I_{k''}$ is a program.

Definition.

$$Support_{0,D} := Support_D \wedge Support_0$$

∎

As a reminder, $Support_D$ means that a transition is supported and that if an element is in any dynamic property from **Dyn** then it is in Has, and if an element is in any dynamic property from **Dyn*** then it is in Has*. $Support_0$ means that a transition is supported at the level of trees, that each, each tree is at least fully in Has or fully in Has*.

Definition. P induces a formula, ϕ_P:

$$Invs := \bigwedge_{1 \leqslant j \leqslant k'} \forall (I_j \wedge I_j^\star) \wedge \bigwedge_{k'+1 \leqslant j \leqslant k''} \forall (I_j \to I_j^\star) \qquad Context := \mathcal{T}_n \wedge Support_{0,D} \wedge Invs$$

$$\phi_P := \bigvee_{1 \leqslant i \leqslant k} \downarrow (\theta_i \wedge Context)$$

∎

In P, each θ_i is a local reaction: every quantification is guarded, so the formula cannot say anything beyond a neighbhood of the interpretation of the variables. The I's are both "state" formulas (in the sense that they either not use **Dyn** or not use **Dyn***) and purely universal. They are separated in two sequences.

Members of the first are interpreted as $(I \wedge I^\star)$, i.e. a constraint that should be true at every step of a run. For instance, suppose an element must at least be phosphorylated to be considered "active". We can write $(\mathsf{Active}(x) \to \mathsf{Phos}(x)) \wedge (\mathsf{Active}(x) \to \mathsf{Phos}(x))^\star$.

Members of the second sequence are interpreted as $(I \to I^\star)$, a constraint that should not vary throughout a run. For instance, suppose that some existing program is not supposed to influence the phosphorylation state of the elements throughout the run. We can write $\mathsf{P} \wedge (\mathsf{Phos} \to \mathsf{Phos}^\star) \wedge (\neg\mathsf{Phos} \to \neg\mathsf{Phos}^\star)$ to make sure that it is the case.

5.2. $\cdot \vdash_\rho \phi_P$

In this section, we show that program formulas are $\vdash_\rho$-provable. The θ_i part is trivial.

Lemma 5.2.1. For each θ_i, $\vdash_{\bar{p}} \theta_i$.

Proof. By induction on the structure: if a literal, use LITERALρ, if a disjunction use DISJρ, if a conjunction use CONJρ. If a guarded universal quantification use ∀GUARDρ, if a guarded existential quantification use ∃GUARDρ. □

The following lemma shows that invariants a $\vdash_{\bar{p}}$-provable. Note that the conditions are weaker than those given in the definition of program: the free variables in dynamic literals of ψ do *not* have to all be x.

The difficulty here is that we want unguarded universal quantification, and all we have are universally quantified clauses thanks to the rule CLAUSEρ. We present the idea with an example. Consider the formula

$$F(x,y) := \mathsf{Link}(x,y) \to (A(x) \wedge K(y))$$

We show $\cdot \vdash_{\bar{p}} F \to F^\star$. First, note that

$$F(x,y) \equiv \mathcal{C}_1(x,y) \wedge \mathcal{C}_2(x,y)$$
$$\mathcal{C}_1(x,y) := \neg\mathsf{Link}(x,y) \vee A(x)$$
$$\mathcal{C}_2(x,y) := \neg\mathsf{Link}(x,y) \vee K(y)$$

And both $G_{\mathcal{C}_1}$ and $G_{\mathcal{C}_2}$ are connected. Now we have

$$F(x,y) \to F^\star(x,y)$$
$$\equiv (\mathcal{C}_1(x,y) \wedge \mathcal{C}_2(x,y)) \to (\mathcal{C}_1^\star(x,y) \wedge \mathcal{C}_2^\star(x,y))$$
$$\equiv \left(\mathcal{C}_1^\star(x,y) \vee \mathsf{Link}(x,y) \vee \mathsf{Link}(x,y)\right) \wedge \left(\mathcal{C}_1^\star(x,y) \vee \mathsf{Link}(x,y) \vee \neg K(y)\right) \tag{1}$$
$$\wedge \left(\mathcal{C}_1^\star(x,y) \vee \neg A(x) \vee \mathsf{Link}(x,y)\right) \wedge \left(\mathcal{C}_1^\star(x,y) \vee \neg A(x) \vee \neg K(y)\right) \tag{2}$$
$$\wedge \left(\mathcal{C}_2^\star(x,y) \vee \mathsf{Link}(x,y) \vee \mathsf{Link}(x,y)\right) \wedge \left(\mathcal{C}_2^\star(x,y) \vee \mathsf{Link}(x,y) \vee \neg K(y)\right) \tag{3}$$
$$\wedge \left(\mathcal{C}_2^\star(x,y) \vee \neg A(x) \vee \mathsf{Link}(x,y)\right) \wedge \left(\mathcal{C}_2^\star(x,y) \vee \neg A(x) \vee \neg K(y)\right) \tag{4}$$

We have avoided making boolean simplification. Each of the clauses above also has a connected graph. For the second clauses lines (3) and (4), $\mathcal{C}_i^\star(x,y)$ contains $K(y)$, so the clause is valid. All the other clauses are great:for those on lines (1) and for the firsts of lines (2),(3) and (4), it is because they contain both $\mathsf{Link}(x,y)$ and $\neg\mathsf{Link}^\star(x,y)$. For the second clauses of lines (2) and (3), it is because they contain both $\neg A(x)$ and $A^\star(x)$. Since they are great, they are provable and the rest is by CONJρ and the rest is because ∀ distributes over ∧.

We now give the proof for the general case:

Lemma 5.2.2. If $q \geqslant 0$ and ϕ is a pre, quantifier-free formula of the form

$$\alpha_1(x,y_1) \to \ldots \to \alpha_q(x,y_q) \to \psi$$

such that $\langle\psi\rangle \subseteq \{x, y_1, \ldots, y_q\}$, then $\cdot \vdash_{\bar{p}} \forall\phi \to \phi^\star$, $\cdot \vdash_{\bar{p}} \forall\phi^\star \to \phi$ and $\cdot \vdash_{\bar{p}} \forall\phi \wedge \phi^\star$.

Proof. We first show $\cdot \vdash_{\bar{p}} \phi \to \phi^\star$. Let $\mathcal{C}_1', \ldots, \mathcal{C}_k'$ be clauses such that $\psi \equiv \bigwedge_j \mathcal{C}_j'$.

ϕ is equivalent to $\bigwedge_{1 \leqslant j \leqslant k} \mathcal{C}_j$ where $\mathcal{C}_j := \mathcal{C}_j' \vee \bigvee_i \neg\alpha_i(x,y_i)$. Note that due to the proviso on $\langle\psi\rangle$, each $G_{\mathcal{C}_j}$ is connected We can also write each $\mathcal{C}_j$ as $L_{j.1} \vee \ldots \vee Lj.m_j$, with each $L_{j.i}$ a literal

and m_j the size of C_j. Now:

$$\phi \to \phi^\star \equiv \left(\bigwedge_{1\leqslant j\leqslant k} C_j \right) \to \left(\bigwedge_{1\leqslant j\leqslant k} C_j^\star \right)$$

$$\equiv \bigwedge_{1\leqslant j\leqslant k} C_j^\star \vee \bigvee_{1\leqslant j\leqslant k} \neg C_j$$

$$\equiv \bigwedge_{1\leqslant j\leqslant k} \bigwedge_{1\leqslant i_1\leqslant m_1} \cdots \bigwedge_{1\leqslant i_k\leqslant m_k} C_j^\star \vee \overline{L_{1.i_1}} \vee \ldots \vee \overline{L_{k.i_k}} \qquad (5.1)$$

Take any clause $\mathcal{D}_{j.i_1.\cdots.i_k} := C_j^\star \vee \overline{L_{1.i_1}} \vee \ldots \vee \overline{L_{k.i_k}}$ in (5.1). Each contains $C_j^\star$, so $G_{\mathcal{D}_{j.i_1.\cdots.i_k}}$ is connected. Now consider $\overline{L_{j.i_j}}$. By definition $L_{j.i_j}^\star \in C_j^\star$. If $L_{j.i_j} = L_{j.i_j}^\star$, $\mathcal{D}_{j.i_1.\cdots.i_k} \equiv \top$ and we remove it. Otherwise $\overline{L_{j.i_j}} \in \mathsf{Dyn}(\mathcal{D}_{j.i_1.\cdots.i_k})$ and $\overline{L_{j.i_j}^\star} \in \mathsf{Dyn}^\star(\mathcal{D}_{j.i_1.\cdots.i_k})$, so $\mathcal{D}_{j.i_1.\cdots.i_k}$ is great and we keep it. Let $\mathcal{D}_{k_1}, \ldots, \mathcal{D}_{k_s}$ be the great clauses we kept.

$$\cfrac{\cfrac{\mathcal{D}_{k_1}^\circ \vdash_{\!\rho} \forall\mathcal{D}_{k_1}}{}\ \text{CLAUSE}_\rho \qquad \cfrac{\mathcal{D}_{k_2}^\circ \vdash_{\!\rho} \forall\mathcal{D}_{k_2}}{}\ \text{CLAUSE}_\rho}{\mathcal{D}_{k_1}^\circ, \mathcal{D}_{k_2}^\circ \vdash_{\!\rho} (\forall\mathcal{D}_{k_1}) \wedge (\forall\mathcal{D}_{k_2})}\ \text{CONJ}_\rho$$

$$\cfrac{\ddots \qquad\qquad \cfrac{\mathcal{D}_{k_s}^\circ \vdash_{\!\rho} \forall\mathcal{D}_{k_s}}{}\ \text{CLAUSE}_\rho}{\mathcal{D}_{k_1}^\circ, \ldots, \mathcal{D}_{k_s}^\circ \vdash_{\!\rho} (\forall\mathcal{D}_{k_1}) \wedge \ldots \wedge (\forall\mathcal{D}_{k_s})}\ \text{CONJ}_\rho$$

By the above proof, we get $\cdot \vdash_{\!\rho} \phi \to \phi^\star$. For $\phi^\star \to \phi$, swap the stars in (5.1): again using $L_{j.i_j}$, each clause is either valid or great and the rest follows.

Finally, each $C_j^\star$ is great, so (as above) by repeated application of CLAUSE$_\rho$ and CONJ$_\rho$, $\vdash_{\!\rho} \forall\phi^\star$. Since we know $\cdot \vdash_{\!\rho} \forall\phi^\star \to \phi$ and since $\phi \wedge \phi^\star \equiv \phi^\star \wedge (\phi^\star \to \phi)$, we get $\cdot \vdash_{\!\rho} \forall\phi \wedge \phi^\star$ using CONJ$_\rho$. $\qquad\square$

Lemma 5.2.3. $\cdot \vdash_{\!\rho} \mathcal{T}_n$

Proof. $\mathcal{T}_n$ contains no dynamic literal and a single bound variable, so when in CNF, each clause is great; repeated application of CLAUSE$_\rho$ and CONJ$_\rho$ yields $\cdot \vdash_{\!\rho} \mathcal{T}_n$, up to equivalence.

For *Symlink* and *Funlink*, we apply lemma 5.2.2 as both are the form $\forall\phi \wedge \phi^\star$ with ϕ as the lemma requires. $\qquad\square$

Lemma 5.2.4. $\cdot \vdash_{\!\rho} \phi_P$

Proof. First, $Support_{0,D} \equiv Support \wedge \forall\mathsf{Has}_0 \wedge \mathsf{Has}_0^\star \wedge Supp_D \wedge Supp_D^\star$ and by lemma 5.2.2, $\vdash_{\!\rho} \forall\mathsf{Has}_0 \wedge \mathsf{Has}_0^\star$ and $\vdash_{\!\rho} \forall Supp_D \wedge Supp_D^\star$. Moreover, by lemma 5.2.3, $\cdot \vdash_{\!\rho} \mathcal{T}_n$. By lemma 5.2.1, $\vdash_{\!\rho} \theta_i$ and for each I_j, again by lemma 5.2.2, $\cdot \vdash_{\!\rho} \forall(I_j \wedge I_j^\star)$ and $\cdot \vdash_{\!\rho} \forall(I_j \to I_j^\star)$. Note that (unlike in lemma 5.3.4), context matters since $\vdash_{\!\rho}$ has no weakening and DISJ$_\rho$ is additive. Modulo equivalence to the right of $\vdash_{\!\rho}$, we get:

$$\cfrac{\cfrac{\vdots}{\vdash_{\!\rho} \theta_i} \quad \cfrac{\vdots}{J_1 \vdash_{\!\rho} Invs} \quad \cfrac{\vdots}{J_2 \vdash_{\!\rho} \mathcal{T}_n} \quad \cfrac{\vdots}{J_3 \vdash_{\!\rho} \forall\mathsf{Has}_0 \wedge \mathsf{Has}_0^\star} \quad \cfrac{\vdots}{J_4 \vdash_{\!\rho} \forall Supp_D \wedge Supp_D^\star}}{\cfrac{J_1, J_2, J_3, J_4 \vdash_{\!\rho} \theta_i \wedge Invs \wedge \mathcal{T}_n \wedge \forall\mathsf{Has}_0 \wedge \mathsf{Has}_0^\star \wedge Supp_D \wedge Supp_D^\star}{J_1, J_2, J_3, J_4 \vdash_{\!\rho} \downarrow\left(\theta_i \wedge Invs \wedge \mathcal{T}_n \wedge Support_{0,D}\right)}\ \text{CIRCUM}_\rho}\ \text{CONJ}_\rho^{(5)}$$

Now let $\phi_i := \downarrow \left(\theta_i \wedge \textit{Invs} \wedge \mathfrak{T}_n \wedge \textit{Support}_{0,D} \right)$,

$$\frac{J_1, J_2, J_3 \vdash_{\rho} \phi_1 \quad \cdots \quad J_1, J_2, J_3 \vdash_{\rho} \phi_k}{J_1, J_2, J_3 \vdash_{\rho} \phi_P} \text{ DISJ}_{\rho}{}^{(k-1)}$$

So by $\textsc{Bool}_{\lhd}$, $\cdot \vdash_{\rho} \phi_P$. $\qquad\square$

5.3. $\cdot \vdash_{\lhd} \phi_P$

Here we show that program formulas are $\vdash_{\lhd}$-provable.

Lemma 5.3.1. For each θ_i, $\cdot \vdash_{\lhd} \theta_i$.

Proof. By induction on the structure: if a literal, use $\textsc{Dynamic}_{\lhd}$ or $\textsc{Static}_{\lhd}$, if a conjunction or a disjunction use $\textsc{Bool}_{\lhd}$ together with $\textsc{Weak}_{\lhd}$ if needed. If a guarded existential quantification use $\exists\text{-}\textsc{Guard}_{\lhd}$. Finally, if a guarded universal quantification:

$$\frac{\dfrac{\mathcal{V}, d \vdash_{\lhd} \theta_i'}{\mathcal{V} \setminus \{x\}, d' \vdash_{\lhd} \alpha(x,y) \to \theta_i'} \text{ Guard}_{\lhd}}{\mathcal{V} \setminus \{x\}, d' \vdash_{\lhd} \forall x.\ \alpha(x,y) \to \theta_i'} \forall_{\lhd}$$

Where $d' \geqslant d$. $\qquad\square$

The following lemma shows that invariants a $\vdash_{\rho}$-provable. Note that the conditions are weaker than those given in the definition of program: the free variables of ψ do not have to be included in $\{x, y_1, \ldots, y_q\}$.

Lemma 5.3.2. If $q \geqslant 0$ and ϕ is a pre, quantifier-free formula of the form

$$\alpha_1(x, y_1) \to \ldots \to \alpha_q(x, y_q) \to \psi$$

such that for every dynamic literal L in ψ, $\langle L \rangle = \{x\}$, then $\cdot \vdash_{\lhd} \forall \phi \to \phi^{\star}$ and $\cdot \vdash_{\lhd} \forall \phi \wedge \phi^{\star}$.

Proof. We first show $\cdot \vdash_{\lhd} \forall \phi \wedge \phi^{\star}$.

We start with $\{x\}, 0 \vdash_{\lhd} \psi^{\star}$. Wlog assume ψ is in negation normal form. By induction on the structure of $\psi^{\star}$: If a conjunction or disjunction, apply $\textsc{Weak}_{\lhd}$ as needed, then $\textsc{Bool}_{\lhd}$. If a literal, apply $\textsc{Dynamic}_{\lhd}$ or $\textsc{Static}_{\lhd}$. By assumption all dynamic literals have $\{x\}$ as variables.

$$\frac{\dfrac{\vdots}{\dfrac{\{x\}, 0 \vdash_{\lhd} \psi^{\star}}{\{x\}, 0 \vdash_{\lhd} \alpha_1(x, y_1)^{\star} \to \psi^{\star}} \text{ Guard}_{\lhd}}}{\vdots \text{ Guard}_{\lhd}}$$

$$\frac{\{x\}, 0 \vdash_{\lhd} \phi^{\star}}{\emptyset, 0 \vdash_{\lhd} \phi \wedge \phi^{\star}} \text{ Invariant}_{\lhd}$$

$$\frac{\vdots\ \forall_{\lhd}}{\emptyset, 0 \vdash_{\lhd} \forall \phi \wedge \phi^{\star}} \forall_{\lhd}$$

By the above proof we are done for $\forall \phi \wedge \phi^\star$.

Now, since ϕ is pre and quantifier-free, repeated applications of $\textsc{Static}_\lhd$ and $\textsc{Bool}_\lhd$ yield $\emptyset, 0 \vdash_\lhd \neg\phi$ (up to equivalence). We get:

$$
\dfrac{
  \dfrac{\vdots \quad\quad\quad\quad \vdots}{\dfrac{\emptyset, 0 \vdash_\lhd \neg\phi \quad\quad \emptyset, 0 \vdash_\lhd \phi \wedge \phi^\star}{\emptyset, 0 \vdash_\lhd \neg\phi \vee (\phi \wedge \phi^\star)}} \ \textsc{Bool}_\lhd
  \quad\vdots\ \forall_\lhd
}{\emptyset, 0 \vdash_\lhd \forall \neg\phi \vee (\phi \wedge \phi^\star)} \ \forall_\lhd
$$

Since $\phi \to \phi^\star \equiv \neg\phi \vee (\phi \wedge \phi^\star)$, we are done. $\qquad\square$

Lemma 5.3.3. $\cdot \vdash_\lhd \mathcal{T}'_n \wedge \textit{Funlink}$

Proof. Using $\textsc{Bool}_\lhd$, it suffices to show that each conjunct is $\vdash_\lhd$-provable.

$\mathcal{T}'_n$ is purely universal with no dynamic literal, so there is a quantifier-free formula F such that $\emptyset, 0 \vdash_\lhd F$ (using $\textsc{Static}_\lhd$, $\textsc{Bool}_\lhd$ and $\textsc{Weak}_\lhd$) and $\forall F \equiv \mathcal{T}'_n$. The rest is by $\forall_\lhd$.

For *Funlink*, we can apply lemma 5.3.2 as it has the form

$$
\forall (\alpha_1(x, y) \to \alpha_2(x, z) \to y = z) \wedge (\alpha_1(x, y) \to \alpha_2(x, z) \to y = z)^\star
$$

and $y = z$ is not a dynamic literal. $\qquad\square$

Note that *Symlink* is not $\vdash_\lhd$-provable.

Lemma 5.3.4. $\cdot \vdash_\lhd \phi_P$

Proof. First, $\textit{Support}_{0,D} := \textit{Support} \wedge \forall \mathsf{Has}_0 \wedge \mathsf{Has}_0^\star \wedge \textit{Supp}_D \wedge \textit{Supp}_D^\star$ and by lemma 5.3.2, $\vdash_\lhd \forall \mathsf{Has}_0 \wedge \mathsf{Has}_0^\star$. By lemma 5.3.1, $\vdash_\lhd \theta_i$ and for each I_j, again by lemma 5.3.2, $e \vdash_\lhd \forall (I_j \wedge I_j^\star)$ and $\cdot \vdash_\lhd \forall (I_j \to I_j^\star)$.

Second, $\mathcal{T}_n := \mathcal{T}'_n \wedge \textit{Funlink} \wedge \textit{Symlink}$ and $\vdash_\lhd \mathcal{T}'_n \wedge \textit{Funlink}$ by lemma 5.3.3. Modulo equivalence and ignoring contexts (there are no context restrictions on $\textsc{Bool}_\lhd$ and $\textsc{Circum}_\lhd$), we get:

$$
\dfrac{
  \dfrac{
    \dfrac{\vdots}{\vdash_\lhd \theta_i} \quad \dfrac{\vdots}{\vdash_\lhd \textit{Invs}} \quad \dfrac{\vdots}{\vdash_\lhd \forall \mathsf{Has}_0 \wedge \mathsf{Has}_0^\star} \quad \dfrac{\vdots}{\vdash_\lhd \mathcal{T}'_n \wedge \textit{Funlink}}
  }{\vdash_\lhd \theta_i \wedge \textit{Invs} \wedge \forall \mathsf{Has}_0 \wedge \mathsf{Has}_0^\star \wedge \mathcal{T}'_n \wedge \textit{Funlink}} \ \textsc{Bool}_\lhd{}^{(4)}
}{\vdash_\lhd \downarrow \left(\theta_i \wedge \textit{Invs} \wedge \mathcal{T}_n \wedge \textit{Support}_{0,D} \right)} \ \textsc{Circum}_\lhd
$$

So by $\textsc{Bool}_\lhd$, $\vdash_\lhd \phi_P$. $\qquad\square$

5.4. Finite compilation of ϕ_P

Here we show how $\vdash_\lhd$ and $\vdash_\rho$-provability of a program formula means that there is a finite set of rules with the same execution semantics as the run semantics of the formula.

This preliminary lemmas shows that at the level of rules, reasoning up to iso is enough: since all operations on rules are defined in terms of morphisms, identities stop being important.

Lemma 5.4.1. If R if a set of rules and $R' = R$ up to iso, then $\text{exec}_{\mathfrak{M}}(R) = \text{exec}_{\mathfrak{M}}(R')$ for any $\mathfrak{M} : \Theta$.

Proof. It suffices to show in one direction, for each $\twoheadrightarrow_R$ step. Suppose $R \ni (\mathfrak{S}_1, K, \mathfrak{S}_2) \overset{h}{\simeq} (\mathfrak{S}'_1, K', \mathfrak{S}'_2) \in R'$, $f : (\mathfrak{S}_1, K) \to \mathfrak{M}'$ and $\mathfrak{M}'$ is the pushout of $(f, g, \mathfrak{S}_1, \mathfrak{S}_2, \mathfrak{M})$. Let $f' = f \circ h$ and $g' = g \circ h$. By lemma 2.4.3 and 2.4.1, $f : (\mathfrak{S}'_1, K') \to \mathfrak{M}$. It is routine to check that $g' : \text{dom}(\mathfrak{S}'_2) \setminus \text{dom}(\mathfrak{S}_2) \to \mathcal{U} \setminus \text{dom}(\mathfrak{M})$ is injective and that $\mathfrak{M}'$ is the pushout of $(f', g', \mathfrak{S}_2, \mathfrak{S}'_2, \mathfrak{M})$ (just transport the construction of $\mathfrak{M}'$ as the pushout of $(f, g, \mathfrak{S}_1, \mathfrak{S}'_1, \mathfrak{M})$ through h). $\square$

Definition. If ϕ is a formula and there is a set of rules R such that $\rho_\emptyset(\phi) = \text{exec}_\emptyset(R)$, ϕ can be *compiled* to R. ■

This is the main theorem of the chapter, and the most important of the thesis. It shows that there is a finite set of rules with execution semantics that correspond exactly to the run semantics of a program. Since the rules are of bounded size and model checking is decidable, the proof is constructive (albeit not very practical).

Theorem 5.4.2. ϕ_P can be compiled to a finite set of rules.

Proof. Let $R := \text{rules}(\mathcal{P}(\phi_P, \|\phi_P\|))$. By lemma 5.4.1 it suffices to show that $R/\simeq$ is finite and that $\rho_\emptyset(\phi_P) = \text{exec}_\emptyset(R)$.

By lemma 5.2.4, $\cdot \vdash_\rho \phi_P$, and $\emptyset \in \text{Init}(\psi)$ for any pre,ψ. Moreover $\phi_P \equiv \phi_P \wedge \textit{Support}_0$, so by theorem 3.5.1 $\rho_\emptyset(\phi_P) = \text{exec}_\emptyset(R)$.

By lemma 5.3.4, $\cdot \vdash_\lhd \phi_P$ so by theorem 4.4.5, ϕ_P is $(\mathcal{V}, d)$-preserved for some $\mathcal{V}, d$. By definition ϕ_P implies $\mathcal{T}_{ns}$ and and is a disjunction of minimised formulas, so by lemma 4.3.7, $\mathcal{P}(\phi_P, \|\phi_P\|)$ is finite up to iso, so by definition of rules, $R/\simeq$ is finite. $\square$

6. (Un)decidability modulo n-BFL

This chapter covers the (un)decidability of satisfability for FO Θ-formulas and the elimination of $\downarrow$ from FO[$\downarrow$].

First, in section 6.1, we show that satisfiability modulo the theory of n-BFL is not decidable. The proof goes through a standard domino reduction.

Second, in section 6.2, we show that satisfiability modulo that theory is decidable in $\exists^*\forall^*$ fragment. Without function symbols, this fragment is called the Bernays-Schönfinkel-Ramsay class and decidability in this class is a classic theorem. All we need to do is show that n-BFLs functions are so restricted that the original proof still through.

The *prenex fragments* are used here. As a quick reminder, $\phi \in \exists^*\forall^*$ means that ϕ (not necessarily in FO) is equivalent to some FO formula ϕ' in prenex normal form such that its quantifier prefix is in the language $\exists^*\forall^*$.

Modulo the theory of n-BFLs, theorem 6.1.6 gives a negative result:

Theorem 6.1.6. FO satisfiability modulo $\mathcal{T}_{ns}$ is undecidable for $n \geqslant 1$.

However, there is a fragment where satisfiability is decidable, given by theorem 6.2.3:

Theorem 6.2.3. For $n \geqslant 0$, satisfiability of formulas in $\exists^*\forall^*$ is decidable modulo $\mathcal{T}_n$.

Which is nice from a querying perspective, since all programs are $\vdash_\lhd$-provable and theorem 7.3.8 says:

Theorem 7.3.8. If $\mathcal{V} ; d \vdash_\lhd \phi$, then $\phi \wedge \mathcal{T}_n$ is in both $\exists^*\forall^*$ and $\forall^*\exists^*$.

6.1. Undecidability

We show that FO-satisfiability (finite or not) is undecidable for n-BFLs by a reduction from domino problems. The specific approach is inspired by [Grädel, 1999].

We recall the definitions of domino systems and their tilings below, but first give some intuition: a domino system is given by a set of colored, square dominoes D. Dominoes are available in infinite supply, but there are finitely many colors. They must be tiled over the infinite plane following *horizontal* and *vertical* constraints. Horizontal constraints define which colors can touch horizontally; vertical constraints define which colors can touch vertically. A solution to a domino system is a tiling of the plane with the dominoes of D which respects both horizontal and vertical constraints. The problem of finding tilings for domino systems is undecidable.

Definition. A *domino system* $\mathcal{D}$ is a triple (D, H, V) with D finite and H, V two binary relations on D. $\mathcal{D}$ *tiles* $\mathbb{N}^2$ when there exists a *tiling* $\tau : \mathbb{N}^2 \to D$ such that for all $x, y \in \mathbb{N}$:

1. $(\tau(x, y), \tau(x + 1, y)) \in H$

2. $(\tau(x, y), \tau(x, y + 1)) \in V$

Moreover, τ is *periodic* with horizontal period $h > 0$ and vertical period $v > 0$ if for all $x, y \in \mathbb{N}$,

$$\tau(x, y) = \tau(x + h, y) = \tau(x, y + v)$$

<p align="right">■</p>

Definition (Conservative reduction class)**.** A recursive fragment of FO X is a *conservative reduction class* if there is a recursive function g from FO to X such that for all FO-formulas ψ, ψ is (finitely) satisfiable iff $g(\psi)$ is (finitely) satisfiable. ■

For more on reduction classes, see [Börger et al., 1997]. Since satisfiability is undecidable in full FO, a conservative reduction class has undecidable satisfiability.

Theorem 6.1.1 (Semi-conservative reduction from the domino problem)**.** Let X be a recursive fragment of FO. X is a conservative reduction class if there exists a recursive function that associates with every domino system $\mathcal{D}$ a formula $\psi_{\mathcal{D}} \in X$ such that

1. If $\mathcal{D}$ admits a periodic tiling of $\mathbb{N}^2$, then $\psi_{\mathcal{D}}$ has a finite model.

2. If $\mathcal{D}$ does not tile $\mathbb{N}^2$, then $\psi_{\mathcal{D}}$ is unsatisfiable.

Proof. This is Corollary 3.1.8 in [Börger et al., 1997], chapter 8, p.91. □

Our class X is the set of FO formulas of the form $\phi \wedge \mathcal{T}_{ns}$, so it is recursive.

Domino theory

Let $\mathcal{D} := (D, H, V)$ be a domino system. We will define a BFL signature that can encode domino systems. Then, we will define a formula which is finitely satisfiable when $\mathcal{D}$ has a periodic tiling, and which is unsatisfiable when $\mathcal{D}$ has *no* tiling.

Definition. The *domino signature of $\mathcal{D}$* is the BFL signature $\Sigma_{\mathcal{D}} := (\mathbf{Dyn}, \mathbf{Dyn}^{\star}, \mathbf{Stat})$ where

$$\mathbf{Dyn} := \mathsf{Has}, \mathsf{Link}$$
$$\mathbf{Dyn}^{\star} := \mathsf{Has}^{\star}, \mathsf{Link}^{\star}$$
$$\mathbf{Stat} := \mathbf{D}, \blacktriangle, \blacktriangleright, \blacktriangledown, \blacktriangleleft, parent$$

$\mathbf{D}$ is a tuple of unary predicate symbols with underlying set D (the ordering is not important), and $\blacktriangle, \blacktriangleright, \blacktriangledown, \blacktriangleleft$ are unary function symbols. ■

$\mathbf{D}$ names the roots in $\Sigma_{\mathcal{D}}$-structures. Each domino will be represented by a tree. The root will have its color (some element of $\mathbf{D}$) as a property. To represent adjacency, roots will have 4 children (one for each direction). The unary functions $\blacktriangle, \blacktriangleright, \blacktriangledown, \blacktriangleleft$ are for the directions *up, right, down, left*, respectively. We define the set of direction symbols $\Delta := \{\blacktriangle, \blacktriangleright, \blacktriangledown, \blacktriangleleft\}$.

We also will need to express "the inverse direction". Let

$$\blacktriangle^{\perp} := \blacktriangledown \qquad\qquad\qquad \blacktriangleright^{\perp} := \blacktriangleleft$$
$$\blacktriangledown^{\perp} := \blacktriangle \qquad\qquad\qquad \blacktriangleleft^{\perp} := \blacktriangleright$$

We sometimes omit parentheses after directions for readability. For instance, we may write $\blacktriangle x$ instead of $\blacktriangle(x)$.

Grid theory A grid is a graph of trees connected only through their ⚹ or their ►◄ vertices, and where the ▲ ► diagonal leads to the same tree as the ► ▲ diagonal.

Grids give us a canvas to set the dominos:

$$Grid := \forall x.\ parent(x) = x \rightarrow$$

$$\wedge\ \exists u, r, d.\ \mathsf{Link}(\blacktriangle x, \blacktriangledown u) \wedge \mathsf{Link}(\blacktriangleright u, \blacktriangleleft d) \tag{6.1}$$

$$\wedge\ \mathsf{Link}(\blacktriangleright x, \blacktriangleleft r) \wedge \mathsf{Link}(\blacktriangle r, \blacktriangledown d) \tag{6.2}$$

The first line restricts the condition to tree roots (recall that an n-BFL can be seen as a forest of bounded-height trees. Tree vertices are structure elements). For every root a, line (6.1) forces the existence of a link between the up-child of a and the down-child of some other root u (for *up-root*). It also forces the existence of a link between the right-child of a and the left-child of some other root r (for *right-root*). Finally, it forces the existence of a root d (for *diagonal-root*) which completes the pattern. See figure 6.1 for an illustration of the effect of *Grid*.

Figure 6.1.: The effect of *Grid*

Definition. For any n-BFL $\mathfrak{A} : \Sigma_{\mathcal{D}}$ let $\rho_{\mathfrak{A}} \subseteq \mathrm{dom}(\mathfrak{A})$ be the tree roots of $G_{\mathfrak{A}}$. ∎

The next lemma provides function in the metalanguage which correspond to going through a diagonal in a domino tiling. Since the lemma shows that the functions traverse any 1-BFL satisfying *Grid* and the way we take the diagonal (right-up or up-right) is irrelevant, it implies that grids are well-formed.

Lemma 6.1.2. For any $\mathfrak{A}$ satisfying $\mathcal{T}_{1S}$ and *Grid*, there are functions $\mathrm{up}_{\mathfrak{A}}, \mathrm{right}_{\mathfrak{A}} : \rho_{\mathfrak{A}} \rightarrow \rho_{\mathfrak{A}}$ such that:

- $\mathrm{up}_{\mathfrak{A}} \circ \mathrm{right}_{\mathfrak{A}} = \mathrm{right}_{\mathfrak{A}} \circ \mathrm{up}_{\mathfrak{A}}$

- For all $a \in \rho_{\mathfrak{A}}$ $(\llbracket \blacktriangleright \rrbracket_{\mathfrak{A}}(a), \llbracket \blacktriangleleft \rrbracket_{\mathfrak{A}}(\mathrm{right}_{\mathfrak{A}}(a))) \in \llbracket \mathsf{Link} \rrbracket_{\mathfrak{A}}$

- For all $a \in \rho_{\mathfrak{A}}$ $(\llbracket \blacktriangle \rrbracket_{\mathfrak{A}}(a), \llbracket \blacktriangledown \rrbracket_{\mathfrak{A}}(\mathrm{up}_{\mathfrak{A}}(a))) \in \llbracket \mathsf{Link} \rrbracket_{\mathfrak{A}}$

Proof. For $a \in \rho_{\mathfrak{A}}$, by $Grid(6.1), Grid(6.2)$ there is $u_a \in \mathrm{dom}(\mathfrak{A})$ such that $\mathfrak{A} \models \mathsf{Link}(\blacktriangle a, \blacktriangledown u_a)$. Since $\llbracket \mathsf{Link} \rrbracket_{\mathfrak{A}}$ is functional, u_a is uniquely defined.

Similarly, there are unique r_a and d_a such that $\mathfrak{A} \models \mathsf{Link}(\blacktriangleright u_a, \blacktriangleleft d_a)$, $\mathfrak{A} \models \mathsf{Link}(\blacktriangleright a, \blacktriangleleft r_a)$, and $\mathfrak{A} \models \mathsf{Link}(\blacktriangle r_a, \blacktriangledown d_a)$. We define

$$\mathrm{right}_{\mathfrak{A}}(a) := r_a$$
$$\mathrm{up}_{\mathfrak{A}}(a) := u_a$$

Now for the commutativity, we have by definition $\mathrm{right}_{\mathfrak{A}}(u_a) = d_a$ and $\mathrm{up}_{\mathfrak{A}}(r_a) = d_a$, so $\mathrm{right}_{\mathfrak{A}}(\mathrm{up}_{\mathfrak{A}}(a)) = \mathrm{up}_{\mathfrak{A}}(\mathrm{right}_{\mathfrak{A}}(a))$. □

The next lemma uses $\text{right}_{\mathfrak{A}}$ and $\text{up}_{\mathfrak{A}}$ to define a new function which directly connects the plane and 1-BFLs that satisfy *Grid*.

Lemma 6.1.3. *If* $\mathfrak{A} \vDash \mathcal{T}_{1S} \wedge Grid$, *there is* $\mathfrak{h} : \mathbb{N}^2 \to \rho_{\mathfrak{A}}$ *such that for* $i, j \in \mathbb{N}$,

$$\mathfrak{h}(i+1, j) = \text{right}_{\mathfrak{A}}(\mathfrak{h}(i, j))$$
$$\mathfrak{h}(i, j+1) = \text{up}_{\mathfrak{A}}(\mathfrak{h}(i, j))$$

Proof. Take any $a \in \rho_{\mathfrak{A}}$. Define $\mathfrak{h}$ for $i, j \geqslant 0$:

$$\mathfrak{h}(i, j) := \text{right}_{\mathfrak{A}}^{i}(\text{up}_{\mathfrak{A}}^{j}(a))$$

Were the ^{i} and ^{j} exponents denote repeated function application. We have:

$$\mathfrak{h}(i+1, j) = \text{right}_{\mathfrak{A}}^{i+1}(\text{up}_{\mathfrak{A}}^{j}(a)) = \text{right}_{\mathfrak{A}}(\mathfrak{h}(i, j))$$
$$\mathfrak{h}(i, j+1) = \text{right}_{\mathfrak{A}}^{i}(\text{up}_{\mathfrak{A}}^{j+1}(a)) = \text{up}_{\mathfrak{A}}(\text{right}_{\mathfrak{A}}^{i}(\text{up}_{\mathfrak{A}}^{j}(\mathfrak{h}(a)))) = \text{up}_{\mathfrak{A}}(\mathfrak{h}(i, j))$$

With the central equality on line 2 making use of lemma 6.1.2. $\qquad \square$

Domino theory Once a grid is well-formed, it should contain dominos that tile the plane according to the *horizontal constraints* H and the *vertical constraints* V: a coherent grid w.r.t. V and H ensures that nodes connected through $\mathbf{I}$ respect V and those connected through $\blacktriangleright\blacktriangleleft$ respect H.

The theory associated to $\mathcal{D}$ is *Domino*. Note that this is the first time formulas are influenced by $\mathcal{D}$ (however, the signature $\Sigma_{\mathcal{D}}$ already depended on $\mathcal{D}$).

$$Domino := \forall x. \, parent(x) = x \to$$

$$\left(\text{Horizontal}(x) \wedge \text{Vertical}(x) \wedge \bigvee_{N \in D} N(x) \wedge \bigwedge_{N' \neq N} \neg N'(x) \right)$$

$$\text{Horizontal}(x) := \bigwedge_{N \in D} N(x) \to \bigvee_{(N, N') \in H} \exists y. \, \text{Link}(\blacktriangleright x, \blacktriangleleft y) \wedge N'(y)$$

$$\text{Vertical}(x) := \bigwedge_{N \in D} N(x) \to \bigvee_{(N, N') \in V} \exists y. \, \text{Link}(\blacktriangle x, \blacktriangledown y) \wedge N'(y)$$

As before for well-formed grids, we define a function in the metalanguage which maps a well-tiled plane with a well-tiled grid:

Lemma 6.1.4. *If* $\mathfrak{A} : \Sigma_{\mathcal{D}}$ *satisfies* $\mathcal{T}_{1S}$, *Grid and Domino, there is a function* $\tau_{\mathfrak{A}} : \rho_{\mathfrak{A}} \to D$ *such that for all* $a \in \rho_{\mathfrak{A}}$:

1. $(\tau_{\mathfrak{A}}(a), \tau_{\mathfrak{A}}(\text{right}_{\mathfrak{A}}(a))) \in H$

2. $(\tau_{\mathfrak{A}}(a), \tau_{\mathfrak{A}}(\text{up}_{\mathfrak{A}}(a))) \in V$

Proof. Take any $a \in \rho_{\mathfrak{A}}$. By *Domino*, there is a single $N \in D$ such that $a \in [\![N]\!]_{\mathfrak{A}}$. Define $\tau_{\mathfrak{A}}(a) := N$.

For any $a \in \rho_{\mathfrak{A}}$, with $\phi(x) := \text{Link}(\blacktriangleright a, \blacktriangleleft x)$, by lemma 6.1.2 we have $\mathfrak{A} \vDash \phi(\text{right}_{\mathfrak{A}}(a))$ and by functionality of Link that is only the case for $\text{right}_{\mathfrak{A}}(a)$. So by Horizontal_{H}, $\mathfrak{A} \vDash N'(\text{right}_{\mathfrak{A}}(a))$ for some N' such that $(N, N') \in H$. By definition of $\tau_{\mathfrak{A}}$, $\tau_{\mathfrak{A}}(\text{right}_{\mathfrak{A}}(a)) = N'$, so we are done.

The case for V is similar.

$\qquad \square$

Domino model Suppose $\mathcal{D}$ admits a periodic tiling τ with horizontal period h and vertical period v. Let $\mathfrak{T}$ be the $\Sigma_{\mathcal{D}}$-structure with domain of size $5hv$.

$$\text{dom}(\mathfrak{T}) := \{[i,j],\ [i,j].\delta \mid 1 \leqslant i \leqslant h, 1 \leqslant j \leqslant v, \delta \in \Delta\}$$

$$\text{Has}(\mathfrak{T}) := \top$$

Note that "$[i,j]$", "$[i,j].\blacktriangleleft$", and so on are just suggestively-named symbols. For $1 \leqslant i \leqslant h$, $1 \leqslant j \leqslant v$, and $\delta, \delta' \in \Delta$ we specify the function $[\![\delta]\!]_{\mathfrak{T}}$ to be:

$$[\![\delta]\!]_{\mathfrak{T}}([i,j]) := [i,j].\delta$$
$$[\![\delta]\!]_{\mathfrak{T}}([i.j].\delta') := [i,j].\delta'$$

and $[\![\text{Link}]\!]_{\mathfrak{T}}$ to be the smallest symmetric relation such that:

$$([i,j].\,\blacktriangleright, [i+1 \bmod h, j].\,\blacktriangleleft) \in [\![\text{Link}]\!]_{\mathfrak{T}}$$
$$([i,j].\blacktriangle, [i, j+1 \bmod v].\blacktriangledown) \in [\![\text{Link}]\!]_{\mathfrak{T}}$$

and let $[\![parent]\!]_{\mathfrak{T}}$ be the parent relation of $G_{\mathfrak{T}}$ on non-roots and the identity on roots; this is well defined since $G_{\mathfrak{T}}$ is trivially a tree.

This gives us a discrete h by v torus. Next, we define the predicates $[\![N]\!]_{\mathfrak{T}}$ for $N \in \mathbf{D}$ according to τ. Remembering that by definition, there is a $N \in \mathbf{D}$ for every $\mathbf{N} \in \mathbf{D}$ and vice versa, we define

$$[i,j] \in [\![N]\!]_{\mathfrak{T}} \text{ iff } \tau(i \bmod h, j \bmod v) = N$$

Lemma 6.1.5. $\mathfrak{T} \models \mathcal{T}_{1S} \wedge Grid \wedge Domino$

Proof.

$\mathfrak{T} \models \mathcal{T}_{1S}$ It suffices to show that $\mathfrak{T}$ is a supported 1-BFL.

> $\mathfrak{A}$ *is a supported n-BFL with symmetric and functional link relations.*

> Trivially $\mathfrak{T}$ is a 1-BFL with roots $[i,j]$ and leafs $[i,j].\delta$ with all elements in $[\![\text{Has}]\!]_{\mathfrak{T}}$. $[\![\text{Link}]\!]_{\mathfrak{T}}$ is symmetric by definition. Take some a of the form $[i,j].\delta$. If $\delta \in \{\blacktriangleright, \blacktriangle\}$, symmetrisation had no impact and there is a single b such that $(a,b) \in [\![\text{Link}]\!]_{\mathfrak{T}}$. If $\delta = \blacktriangleleft$, there is no $1 \leqslant i' \neq i \leqslant h$ such that $i+1 \bmod h = i'+1 \bmod h$ since i itself ranges from 1 to h, so a has a single image in $[\![\text{Link}]\!]_{\mathfrak{T}}$. The case is similar when $\delta = \blacktriangledown$.

$\mathfrak{T} \models$ **Grid** Any root is of the form $[i,j]$ so by definition of $[\![\text{Link}]\!]_{\mathfrak{T}}$ and the four $[\![\delta]\!]_{\mathfrak{T}}$, the desired witnesses u, r, d are respectively $[i, j+1 \bmod v]$, $[i+1 \bmod h, j]$, and $[i+1 \bmod h, j+1 \bmod v]$.

$\mathfrak{T} \models$ **Domino** First, roots belong to a single $N_{\mathfrak{T}}$ predicate because τ is h, v-periodic, so the last part of *Domino* is satisfied.

> For Horizontal, only an element of the form $[i,j]$ may have a label N from D. The desired witness y is provided by $[i+1 \bmod h, j]$: by definition of $[\![\blacktriangleright]\!]_{\mathfrak{T}}$, $[\![\blacktriangleleft]\!]_{\mathfrak{T}}$ and $[\![\text{Link}]\!]_{\mathfrak{T}}$ we have $([\![\blacktriangleright]\!]_{\mathfrak{T}}[i,j], [\![\blacktriangleleft]\!]_{\mathfrak{T}}[i+1 \bmod h, j]) \in [\![\text{Link}]\!]_{\mathfrak{T}}$ and since τ tiles $\mathcal{D}$ and is h, v-periodic, $[i+1 \bmod h, j] \in [\![N']\!]_{\mathfrak{T}}$ for some N' such that $(N, N') \in H$.

> The case for Vertical is similar.

□

Theorem 6.1.6. FO satisfiability modulo $\mathcal{T}_{ns}$ is undecidable for $n \geqslant 1$.

Proof. Let $\mathcal{D} = (D, H, V)$ be a domino system. Let $\psi := \mathcal{T}_{1S} \wedge \textit{Grid} \wedge \textit{Domino}$ be a formula under signature $\Sigma_{\mathcal{D}}$. To make use of theorem 6.1.1, we first note that the finite structure $\mathfrak{A}(\tau)$ satisfies ψ by lemma 6.1.5.

Then suppose $\mathfrak{A} \models \psi$ for some $\Sigma_{\mathcal{D}}$-structure $\mathfrak{A}$. By lemmas 6.1.3 and 6.1.4 we can define $\tau : \mathbb{N}^2 \to D$ by $\tau(i,j) := \tau_{\mathfrak{A}}(\mathfrak{h}(i,j))$ and obtain:

$$(\tau(i,j), \tau(i+1,j)) = (\tau_{\mathfrak{A}}(\mathfrak{h}(i,j)), \tau_{\mathfrak{A}}(\mathfrak{h}(i+1,j))) = (\tau_{\mathfrak{A}}(\mathfrak{h}(i,j)), \tau_{\mathfrak{A}}(\text{right}_{\mathfrak{A}}(\mathfrak{h}(i,j)))) \in H$$

$$(\tau(i,j), \tau(i,j+1)) = (\tau_{\mathfrak{A}}(\mathfrak{h}(i,j)), \tau_{\mathfrak{A}}(\mathfrak{h}(i,j+1))) = (\tau_{\mathfrak{A}}(\mathfrak{h}(i,j)), \tau_{\mathfrak{A}}(\text{right}_{\mathfrak{A}}(\mathfrak{h}(i,j)))) \in V$$

So $\mathcal{D}$ tiles $\mathbb{N}^2$ with τ. □

6.2. Decidable class

In this section we show that modulo $\mathcal{T}_{ns}$, formulas in $\exists^* \forall^*$ have decidable satisfiability.

Definition. Given a signature Σ with Herbrand universe W and $\equiv$ an equivalence relation on W, a $\equiv$-*Herbrand* structure $\mathfrak{A}$ is a Σ-structure with universe $W/\equiv$ such that for all constants c of Σ, $[\![c]\!]_{\mathfrak{A}} = [c]_{\equiv}$, and for all function symbols f of Σ with arity k and all ground terms $t_1, ..., t_k$, $[\![f]\!]_{\mathfrak{A}}([t_1]_{\equiv}, ..., [t_k]_{\equiv}) = [f(t_1, ..., t_k)]_{\equiv}$. A $\equiv$-Herbrand structure that satisfies a formula is a $\equiv$-*Herbrand* model for it. ∎

Lemma 6.2.1 (Folklore). A sentence with equality in Skolem normal form is satisfiable iff it has a $\equiv$-Herbrand model.

Proof. A sentence satisfied by a $\equiv$-Herbrand structure is satisfiable. Now assume a sentence in Skolem normal form ϕ satisfied by a structure $\mathfrak{A}$. For any ground Σ-term t, $[\![t]\!]_{\mathfrak{A}}$ is its image in $\mathfrak{A}$.

We proceed by induction on the number n of universal quantifiers in ϕ.

Take the equivalence relation on W defined by $t \equiv t'$ whenever $[\![t]\!]_{\mathfrak{A}} = [\![t']\!]_{\mathfrak{A}}$ and let $\mathfrak{A}'$ be the $\equiv$-Herbrand structure where for every relational symbol R of Σ with arity k, and any k-tuple of ground terms $(t_1, ..., t_k)$, $([t_1]_{\equiv}, ..., [t_k]_{\equiv}) \in [\![R]\!]_{\mathfrak{A}'}$ iff $([\![t_1]\!]_{\mathfrak{A}}, ..., [\![t_k]\!]_{\mathfrak{A}}) \in [\![R]\!]_{\mathfrak{A}}$. This is well-defined since $\equiv$ on W respects equality on $\mathfrak{A}$'s universe.

If $n = 0$, note that the truth value of all atoms (equality and predicates) is the same in $\mathfrak{A}'$ and $\mathfrak{A}$ by construction. If $\mathfrak{A} \models \forall x.\ \psi(x)$ with the induction hypothesis assumed on ψ, we have by substitution $\mathfrak{A} \models \psi(t)$ for all ground terms t, by induction $\mathfrak{A}' \models \psi(t)$ for all ground terms t, and thus $\mathfrak{A}' \models \forall x.\ \psi(x)$. □

Lemma 6.2.2. If $\Sigma = (\Sigma', \mathbf{c})$ where Σ' is a BFL signature and $\mathbf{c}$ are constant symbols, if $\mathfrak{A} : \Sigma$ is a $\equiv$-Herbrand model of $\mathcal{T}_n$ then $\mathfrak{A}$ is finite.

Proof. We show: if $\mathfrak{A} \models Tn$ then $\equiv$ is of finite index.

For any constant c and any ground term t of Σ, if t contains c then $[c]_{\equiv}$ and $[t]_{\equiv}$ are in the same tree of $G_{\mathfrak{A}}$. This is easily shown by structural induction on t. There are finitely many c in Σ, so there are finitely many trees in $G_{\mathfrak{A}}$.

Additionally, by lemma 4.1.1, each tree of $G_\mathfrak{A}$ is finite. So $G_\mathfrak{A}$ is finite. $\mathfrak{A}$ has the same domain, so it is as well. □

We show decidability under a triple restriction : syntactic fragment, signature class, and theory.

Theorem 6.2.3. For $n \geqslant 0$, satisfiability of formulas in $\exists^*\forall^*$ is decidable modulo $\mathcal{T}_n$.

Proof. First, note that $\mathcal{T}_n \in \exists^*\forall^*$. For $\phi \in \exists^*\forall$, $\psi := \phi \wedge \mathcal{T}_n$ can equisatisfiably be turned into a Skolem normal form sentence ψ' (replacing variables with constants). By lemmas 6.2.1 and 6.2.2, ψ' is satisfiable iff it is finitely satisfiable. □

7. Change minimality in FO

We then move on to definability questions. In section 7.1, we introduce *unified circumscription*. Generally speaking, circumscription is a 2nd-order schema parameterised by an order o on a class of structures. It takes as input a formula and produces a 2nd-order formula which has, as models, the o-minimal models of φ. There are various forms of circumscription; here we introduce a reasonably general variant which suits our purposes.

Then in section 7.2, we show how the operator ↓ can be expressed by instantiating unified circumscription on the ⊴ order. This provides a syntactic, 2nd-order definition of ↓. The following section (7.3.4) is about rewriting this formula to eliminate its 2nd-order quantifiers under some conditions.

Finally, we show that the theory of n-BFLs (and in particular the functionality of Link and Link⋆) brings even more simplification: ⊢$_\rho$-provable minimised formula are definable in the ∃*∀* and ∀*∃* fragments of first-order logic.

7.1. Unified circumscription

In this section, we define unified circumscription syntactically and prove that for an instanciation of unified circumscription on a given order, the models of the circumscribed formula are indeed the minimal models of the original formula along that order.

We now work in second-order logic.

Generally speaking, the purpose of circumscription is to denote through second-order means the minimal models of some FO formula, along a given order. For instance, McCarthy's domain circumscription [McCarthy, 1980] minimises the size of the domain but maintains every predicate and function interpretation on the remaining elements. Also originated by McCarthy, predicate circumscription [McCarthy, 1986] minimises the extent of some *minimised predicates* – but leaves the other predicates and the domain untouched. There are many forms of circumscription; the one we introduce here combines general domain circumscription [Doherty et al., 1998] and parallel predicate circumscription (predicate circumscription extended to the product of subset ordering for multiple predicate interpretations).

With unified circumscription, a signature Σ is partitioned into tuples of predicates and functions. Some predicates are *minimised*. The domain is *minimised* as well. Some predicates are *fixed*: they may not change at all. Other predicates, as well as some functions, are *restricted*: they may not change on the minimised domain but may reduce their domain of definition if the domain is reduced. Finally, some predicates and functions are *varying*: they may change in any way. As we will see, once a partition of Σ has been established, this gives rise to an order ⪯ on Σ-structures. Then, given a Σ-formula φ, the *circumscription of* φ is simply the set of ⪯-minimal models of φ.

In this section, Σ is assumed to be a signature of the form $(\mathbf{P}_{\min}, \mathbf{P}_{\text{restr}}, \mathbf{f}_{\text{restr}}, \mathbf{P}_{\text{var}}, \mathbf{f}_{\text{var}}, \mathbf{P}_{\text{fix}})$ where the $\mathbf{P}_\bullet$ are predicate tuples and the $\mathbf{f}_\bullet$ are function tuples. We will define a partial order on Σ-models where :

- domains and predicates in $\mathbf{P}_{\min}$ are minimised

Figure 7.1.: Unified circumscription example.

- predicates in $\mathbf{P}_{restr}$ and functions in f_{restr} are restricted

- predicates in $\mathbf{P}_{var}$ and functions in f_{var} can vary

- predicates in $\mathbf{P}_{fix}$ are fixed

Definition. If $\mathfrak{A}, \mathfrak{B} : \Sigma$, $\mathfrak{A} \preceq \mathfrak{B}$ whenever :

(a) $\mathrm{dom}(\mathfrak{A}) \subseteq \mathrm{dom}(\mathfrak{B})$

(b) $[\![\mathbf{P}_{min}]\!]_{\mathfrak{A}} \subseteq [\![\mathbf{P}_{min}]\!]_{\mathfrak{B}}$

(c) $[\![\mathbf{P}_{restr}]\!]_{\mathfrak{A}} = [\![\mathbf{P}_{restr}]\!]_{\mathfrak{B}} \restriction \mathrm{dom}(\mathfrak{A})$

(d) $[\![f_{restr}]\!]_{\mathfrak{A}} = [\![f_{restr}]\!]_{\mathfrak{B}} \restriction \mathrm{dom}(\mathfrak{A})$

(e) $[\![\mathbf{P}_{fix}]\!]_{\mathfrak{A}} = [\![\mathbf{P}_{fix}]\!]_{\mathfrak{B}}$

∎

There is no condition of $\mathbf{P}_{var}$ and f_{var}.

To give some intuition, we try to make the domain and predicates in $\mathbf{P}_{min}$ as small as possible *given that* $\mathbf{P}_{fix}$ must not change at all, and that $\mathbf{P}_{restr}$ and f_{restr} must not change on the remaining domain elements.

Example. Let the tuples that partition Σ each have a single predicate: $\mathbf{P}_{min} = (\mathsf{min}), \mathbf{P}_{restr} = (\mathsf{restr}), \mathbf{P}_{var} = (\mathsf{var}), \mathbf{P}_{fix} = (\mathsf{fix}), f_{restr} = (f)$, and consider the following formula and models:

$$\phi := \forall x.\ \mathsf{min}(x) \lor \mathsf{var}(x)$$
$$\mathfrak{A} := (\{a, b, c, d\}, \mathsf{min} = \{a\}, \mathsf{fix} = \{a\}, \mathsf{var} = \{b, c, d\}, \mathsf{restr} = \{b, c\}$$
$$f = a \mapsto b, b \mapsto b, c \mapsto d, d \mapsto d)$$
$$\mathfrak{B} := (\{a, b\}, \mathsf{min} = \emptyset, \mathsf{fix} = \{a\}, \mathsf{var} = \{a, b\}, \mathsf{restr} = \{b\}, f = a \mapsto b, b \mapsto b)$$

As illustrated in figure 7.1, we have $\mathfrak{B} \prec \mathfrak{A}$, and as a matter of fact $\mathfrak{B}$ is $\preceq$-minimal among models of ϕ. $a \in [\![\mathsf{fix}]\!]_{\mathfrak{A}}$, so circumscription may not lose a, and since $[\![f]\!]_{\mathfrak{A}}(a) = b$ and f is restricted, b may not disappear either. On the other hand, domain minimisation drops both c and d: $c, d \notin \mathrm{dom}(\mathfrak{B})$. We have $b \in [\![\mathsf{min}]\!]_{\mathfrak{A}}$, but by ϕ $b \notin [\![\mathsf{min}]\!]_{\mathfrak{B}}$ is allowed as long as $b \in [\![\mathsf{var}]\!]_{\mathfrak{B}}$. Since min is minimised, this is what happens. Finally, unlike fix, restr does not prevent its members from being removed ($c \notin \mathrm{dom}(\mathfrak{B})$), but must remain unchanged on the circumscribed domain (so $b \in [\![\mathsf{restr}]\!]_{\mathfrak{B}}$ and $a \notin [\![\mathsf{restr}]\!]_{\mathfrak{B}}$). ∎

Definition (Similarity). Two signatures Σ, Σ' are *similar* when they have the same length and symbols in the same position have the same nature (relational or functional) and the same arity. ∎

Definition (Substitution). For any two symbols X, Y, $\{X \mapsto Y\}$ is the syntactic map from X to Y, and for any formula ϕ, $\phi\{X \mapsto Y\}$ is ϕ with every instance of X replaced by a Y. An extension of the map such as $\{X \mapsto Y, W \mapsto Z\}$ for W, Z two other symbols is the simultaneous syntactic replacementof X by Y and W by Z. For any two similar signatures Σ, Σ', $\{\Sigma \mapsto \Sigma'\}$ is the subtitution that simultaneously replaces the ith symbol of Σ' with the ith symbol of Σ for every position i of Σ. ∎

With a first-order formula ϕ, we can characterise the $\preceq$-minimal elements satisfying ϕ with a 2nd-order formula $\mathord{\downharpoonleft}\phi$:

Definition ($\mathord{\downharpoonleft}$ as a 2nd-order formula). If a signature Σ has been partitioned such that $\mathbf{P_{min}}, \mathbf{P_{fix}}, \ldots$ are defined and ϕ is a Σ-formula, then:

$$\mathord{\downharpoonleft}\phi := \phi \wedge$$
$$\forall D, \mathbf{M}, \mathbf{V}, \mathbf{f}.$$
$$\left(\mathrm{dom}(D, \mathbf{f}) \wedge \mathbf{M} \subseteq \mathbf{P_{min}} \wedge \phi[D]\{\mathbf{P_{min}} \mapsto \mathbf{M}, \mathbf{P_{var}} \mapsto \mathbf{V}, \mathbf{f_{var}} \mapsto \mathbf{f}\}\right) \tag{7.1}$$
$$\rightarrow \left(\mathbf{P_{min}} \subseteq \mathbf{M} \wedge \forall x.D(x)\right) \tag{7.2}$$

The right conjunct requires that any structure $\preceq$-smaller than the "current" one is not strictly so – and therefore, equal to the current structure.

It does so by quantifying over possible subdomains (D) as well as minimised and varying predicates & functions ($\mathbf{M}, \mathbf{V}, \mathbf{f}$).

The left-hand side of the implication (line (7.1)) introduces the assumption that D is a subdomain; that $\mathbf{M}$ is indeed smaller than the minimised predicates $\mathbf{P_{min}}$, and that ϕ is true in this subdomain with the new minimised and varying predicates & functions.

The right-hand side of the implication (line (7.2)) says that the subdomain actually covers the current domain, and that $\mathbf{M} = \mathbf{P_{min}}$.

- The quantified symbols are D, unary relational; $\mathbf{M}$ is similar to $\mathbf{P_{min}}$, $\mathbf{V}$ is similar to $\mathbf{P_{var}}$, and $\mathbf{f}$ is similar to $\mathbf{f_{var}}$.

- $\phi[D]$ is ϕ where every quantifier has been simultaneously guarded by D. For any atomic formula A, any first-order formulas ψ, ψ', any $\oplus \in \{\wedge, \vee\}$:

$$A[D] := A \qquad (\neg\psi)[D] := \neg\psi[D] \qquad (\psi \oplus \psi')[D] := \psi[D] \oplus \psi'[D]$$
$$(\forall x.\ \psi)[D] := \forall x.\ D(x) \rightarrow \psi[D] \qquad\qquad (\exists x.\ \psi)[D] := \exists x.\ D(x) \wedge \psi[D]$$

for any variable x, any formula ψ.

- dom specifies that D behaves like a nonempety domain that covers $\mathbf{P_{fix}}$ and the interpretation of the variables:

$$\mathrm{dom}(D, \mathbf{f}) := \mathbf{f_{restr}}(D) \subseteq D \wedge \mathbf{f}(D) \subseteq D \wedge \mathbf{P_{fix}} \subseteq D \wedge \exists x.D(x) \wedge \bigwedge_{x \in \langle\phi\rangle} D(x)$$

  - For any tuple $\mathbf{g}$ of function symbols, and unary predicate D, $\mathbf{g}(D) \subseteq D$ means that D is $\mathbf{g}$-closed:

$$\mathbf{g}(D) \subseteq D := \bigwedge_{f \in \mathbf{g}} \forall \mathbf{x_f}.\ \left(\bigwedge_{x \in \mathbf{x_f}} D(x)\right) \rightarrow D(f(\mathbf{x_f}))$$

where x_f is a tuple of size the arity of f. Note that whenever f is a constant (i.e. of arity 0), the associated conjunct reduces to D(f).

We do not write the proof of correspondence between $g(D) \subseteq D$ as a logical specification and the notion in metalanguage.

– For any tuple **P** of predicate symbols, and unary predicate D, $\mathbf{P} \subseteq D$ means that any element mentioned in some $S \in \mathbf{P}$ is also in D:

$$\mathbf{P} \subseteq D := \bigwedge_{S \in \mathbf{P}} \forall x_S.\ S(x_S) \to \bigwedge_{x \in x_S} D(x)$$

where x_P is a tuple of size the arity of P.

We do not write the proof of correspondence between $\mathbf{P} \subseteq D$ as a logical specification and the notion in metalanguage.

– Finally, for any two tuples of predicate symbols **P, Q** of same size and pointwise arity, $\mathbf{P} \subseteq \mathbf{Q}$ means that **P** is product-wise included in **Q**. First, some notation : for a symbol S at position i in **P**, Q_S is the symbol at position i in **Q**. By assumption it has the same arity.

$$\mathbf{P} \subseteq \mathbf{Q} := \bigwedge_{S \in \mathbf{P}} \forall x_S.\ S(x_S) \to Q_S(x_S)$$

where x_S is a tuple of size the arity of S.

We do not write the proof of correspondence between $\mathbf{P} \subseteq \mathbf{Q}$ as a logical specification and the notion in metalanguage.

∎

Notation. In the sequel, note that if an interpretation maps symbols of the signature, the interpretation overrides the default meaning of the symbols. For instance, in the signature (P) where P is a unary relational symbol and with the model $\mathfrak{A} := (\mathrm{dom}(\mathfrak{A}) := \{a\}, \llbracket P \rrbracket_{\mathfrak{A}} = \{a\})$, consider:

$$\mathfrak{A} \vDash P(a)$$
$$\mathfrak{A}, P \mapsto \emptyset \nvDash P(a)$$

∎

Our goal is to prove that a structure $\mathfrak{A}$ is $\preceq$-minimal among models of ϕ iff $\mathfrak{A} \vDash \rceil \phi$. The following technical lemma is used in *both directions*. It shows that for two structures $\mathfrak{A} \trianglelefteq \mathfrak{B}$, $\mathfrak{A}$ is a model of ϕ iff $\mathfrak{B}$ satisfies a version of ϕ where every quantification is relativised to $\mathrm{dom}(\mathfrak{A})$ and every predicate/function is interpreted as in $\mathfrak{A}$.

Lemma 7.1.1. If $\mathfrak{A}, \mathfrak{B} : \Sigma$, ϕ is a first-order formula, $\mathfrak{A} \preceq \mathfrak{B}$, $\mathbf{M}, \mathbf{V'}, \mathbf{g}$ are respectively similar to $\mathbf{P}_{min}, \mathbf{P}_{var}, f_{var}$, and

$$\llbracket \mathbf{P}_{min} \rrbracket_{\mathfrak{A}} = \mathbf{M'} \restriction \mathrm{dom}(\mathfrak{A}) \qquad \llbracket \mathbf{P}_{var} \rrbracket_{\mathfrak{A}} = \mathbf{V'} \restriction \mathrm{dom}(\mathfrak{A}) \qquad \llbracket f_{var} \rrbracket_{\mathfrak{A}} = \mathbf{g} \restriction \mathrm{dom}(\mathfrak{A})$$
$$\mu := D \mapsto \mathrm{dom}(\mathfrak{A}), \mathbf{P}_{min} \mapsto \mathbf{M'}, \mathbf{P}_{var} \mapsto \mathbf{V'}, f_{var} \mapsto \mathbf{g}$$

then for all $v : \langle \phi \rangle \to \mathrm{dom}(\mathfrak{A})$:
$$\mathfrak{A}, v \vDash \phi \text{ iff } \mathfrak{B}, v, \mu \vDash \phi[D]$$

Proof. First, we show that for any term t, $[\![t]\!]_{\mathfrak{A},\nu} = [\![t]\!]_{\mathfrak{B},\nu,f_{var}\mapsto g} \in dom(\mathfrak{A})$.

- If $t = x$, by $Im(\nu) \subseteq dom(\mathfrak{A})$ and $dom(\mathfrak{A}) \subseteq dom(\mathfrak{B})$ we are done.

- If $t = f(t_1, ..., t_n)$ with $n \geqslant 0$ and $f \in f_{restr}$, using the induction we just need $f(\mathfrak{A}) = f(\mathfrak{B}) \upharpoonright dom(\mathfrak{A})$, which is given by $[\![f_{restr}]\!]_{\mathfrak{B}} \upharpoonright dom(\mathfrak{A}) = [\![f_{restr}]\!]_{\mathfrak{A}}$.

- If $t = f(t_1, ..., t_n)$ with $n \geqslant 0$ and $f \in f_{var}$, using the induction we just need $[\![f]\!]_{\mathfrak{A}} = [\![f]\!]_{f_{var}\mapsto g} \upharpoonright dom(\mathfrak{A})$, which is given by $[\![f_{var}]\!]_{\mathfrak{A}} = g \upharpoonright dom(\mathfrak{A})$.

Now by induction on ϕ:

- (Equality) $t_1 = t_2$, since for any t, $[\![t]\!]_{\mathfrak{A},\nu} = [\![t]\!]_{\mathfrak{B},f_{var}\mapsto g}$ we are done.

- (Other atomics) $S(\mathbf{t})$, with $\mathbf{t}$ a tuple of terms. By the previous fact, we need $\mathbf{P} = \mathbf{Q} \upharpoonright dom(\mathfrak{A})$ for $(\mathbf{P}, \mathbf{Q}) \in \{([\![P_{min}]\!]_{\mathfrak{A}}, M'), ([\![P_{var}]\!]_{\mathfrak{A}}, V'), ([\![P_{restr}]\!]_{\mathfrak{A}}, [\![P_{restr}]\!]_{\mathfrak{B}}), ([\![P_{fix}]\!]_{\mathfrak{A}}, [\![P_{fix}]\!]_{\mathfrak{B}})\}$. This is all by hypothesis.

- (Bool) The case of propositional connectives is trivial.

- ($\forall, \exists$) For quantifiers, thanks to negation we treat $\forall$ only wlog. If $\mathfrak{A}, \nu \models \forall x. \psi$ then for any $a \in dom(\mathfrak{A})$, $\mathfrak{A}, \nu, x \mapsto a \models \psi$ and the induction gives us $\mathfrak{B}, \nu, x \mapsto a, \mu \models \psi[D]$. Since $\mu(D) = dom(\mathfrak{A})$, we get $\mathfrak{B}, \nu, \mu \models \forall x. D(x) \rightarrow \psi[D]$, i.e. $\mathfrak{B}, \nu, \mu \models (\forall x. \psi)[D]$.

  The other direction works the same.

$\square$

The main lemma of this section shows the semantics-syntax correspondence of $\rfloor$ and $\preceq$. The key lemma is lemma 7.1.1 in both directions. From syntax to semantics, given $\mathfrak{A}, \mu \models \rfloor\phi$, the proof uses the key lemma to put the "innards" of a would-be smaller strictly model $\mathfrak{B}$ as instance of the second-order quantifiers of $\rfloor\phi$, which leads to $\mathfrak{A} \preceq \mathfrak{B}$. From semantics to syntax, we do the same inverse, taking all instance sof the second-order quantifications and showing, through minimality and the key lemma, that they verify the inside the $\rfloor\phi$.

Lemma 7.1.2. If ϕ is a first-order formula, the models of $\rfloor\phi$ are the $\preceq$-minimal models of ϕ.

Proof.

($\Rightarrow$) Suppose $\mathfrak{A}, \mu \models \rfloor\phi$ and let $\mathfrak{B}, \mu \models \phi$, $\mathfrak{B} \preceq \mathfrak{A}$. We want $\mathfrak{A} \preceq \mathfrak{B}$. By definition of $\rfloor\phi$, we get:

$$\mathfrak{A}, \mu, D \mapsto dom(\mathfrak{B}), M \mapsto [\![P_{min}]\!]_{\mathfrak{B}}, V \mapsto [\![P_{var}]\!]_{\mathfrak{B}}, f \mapsto [\![f_{var}]\!]_{\mathfrak{B}}$$
$$\models dom(D, f) \wedge M \subseteq P_{min} \wedge \phi[D]\{P_{min} \mapsto M, P_{var} \mapsto V, f_{var} \mapsto f\}$$
$$\rightarrow (P_{min} \subseteq M \wedge \forall x. D(x))$$

We show that the left-hand side of the implication holds.

The second conjunct ($M \subseteq P_{min}$) is given by $\mathfrak{B} \preceq \mathfrak{A}$. The third conjunct ($\phi[D]\{...\}$) is given by lemma 7.1.1, using $\mathfrak{B}, \mu \models \phi$ and $\mathfrak{B} \preceq \mathfrak{A}$ (note that we apply the lemma with $M' := [\![P_{min}]\!]_{\mathfrak{A}}$, $V' := [\![P_{var}]\!]_{\mathfrak{A}}$, and $g := [\![f_{var}]\!]_{\mathfrak{A}}$).

Now for the first conjunct ($dom(D, f)$).

We need $\mathfrak{A}, D \mapsto dom(\mathfrak{B}) \models \exists x. D(x)$. Since $dom(\mathfrak{B})$ is non-empty and $dom(\mathfrak{B}) \subseteq dom(\mathfrak{A})$ by $\preceq$, there is some $a \in dom(\mathfrak{A}) \cap dom(\mathfrak{B})$.

We also need $\text{dom}(\mathfrak{B})$ to be $[\![f_{\text{restr}}]\!]_{\mathfrak{A}}$-closed and $[\![f_{\text{var}}]\!]_{\mathfrak{B}}$-closed. Since $[\![f_{\text{restr}}]\!]_{\mathfrak{B}} = [\![f_{\text{restr}}]\!]_{\mathfrak{A}} \upharpoonright \text{dom}(\mathfrak{B})$ (by $\preceq$), $\text{dom}(\mathfrak{B})$ is $[\![f_{\text{restr}}]\!]_{\mathfrak{A}}$-closed. The other is by definition.

We need $[\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{A}} \subseteq D$. Since $[\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{B}} = [\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{A}}$ (by $\preceq$) and we have $[\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{B}} \subseteq \text{dom}(\mathfrak{B})$ by definition, we get $[\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{A}} \subseteq \text{dom}(\mathfrak{B})$.

The left-hand side of the implication being proven, we use the right-hand side. We want $\mathfrak{A} \preceq \mathfrak{B}$. We have $[\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{A}} \subseteq [\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{B}}$ and $\text{dom}(\mathfrak{A}) \subseteq \text{dom}(\mathfrak{B})$, so $\mathfrak{A} \preceq \mathfrak{B}$ indeed.

($\Leftarrow$) Suppose $\mathfrak{A}$ is $\preceq$-minimal among models of ϕ with interpretation μ for its free variables. We show $\mathfrak{A}, \mu \models \mathcal{J}\phi$. By assumption $\mathfrak{A}, \mu \models \phi$. Now take any predicate tuples D', M', V' (with size and arity matching the symbols in D, M, V) and functions f' (with size and arity matching the symbols in f) such that

$$\mathfrak{A}, \mu, D \mapsto D', M \mapsto M', V \mapsto V', f \mapsto f'$$
$$\models \text{dom}(D, f) \wedge M \subseteq \mathbf{P}_{\text{min}} \wedge \phi[D]\{\mathbf{P}_{\text{min}} \mapsto M, \mathbf{P}_{\text{var}} \mapsto V, f_{\text{var}} \mapsto f\}$$

Let $\mathfrak{B} : \Sigma$ be such that

$$\text{dom}(\mathfrak{B}) = D' \qquad\qquad [\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{B}} = M' \upharpoonright D' \qquad\qquad [\![\mathbf{P}_{\text{restr}}]\!]_{\mathfrak{B}} = [\![\mathbf{P}_{\text{restr}}]\!]_{\mathfrak{A}} \upharpoonright \text{dom}(\mathfrak{B})$$
$$[\![f_{\text{restr}}]\!]_{\mathfrak{B}} = [\![f_{\text{restr}}]\!]_{\mathfrak{A}} \upharpoonright \text{dom}(\mathfrak{B}) \qquad [\![\mathbf{P}_{\text{var}}]\!]_{\mathfrak{B}} = V' \upharpoonright \text{dom}(\mathfrak{B}) \qquad [\![f_{\text{var}}]\!]_{\mathfrak{B}} = f' \upharpoonright \text{dom}(\mathfrak{B})$$
$$[\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{B}} = [\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{A}}$$

We must check that $\mathfrak{B}$ is well-defined.

Using the dom subformula, $\mathfrak{B}$ is not the empty model since there is at least one element in $\text{dom}(\mathfrak{B})$, $\text{dom}(\mathfrak{B})$ is $[\![f_{\text{restr}}]\!]_{\mathfrak{A}}$-closed, $\text{dom}(\mathfrak{B})$ is $[\![f_{\text{var}}]\!]_{\mathfrak{B}}$-closed, and we also get $[\![\mathbf{P}_{\text{fix}}]\!]_{\mathfrak{B}} \subseteq \text{dom}(\mathfrak{B})$.

Since $[\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{B}} \subseteq M' \subseteq [\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{A}}$, we get $\mathfrak{B} \preceq \mathfrak{A}$. Using $\bigwedge_{x \in \langle \phi \rangle} D(x)$, we can now apply lemma 7.1.1 to get $\mathfrak{B}, \mu \models \phi$. By minimality of $\mathfrak{A}$ among models of ϕ, this implies $\mathfrak{A} \preceq \mathfrak{B}$, i.e. $[\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{A}} \subseteq [\![\mathbf{P}_{\text{min}}]\!]_{\mathfrak{B}} = M' \upharpoonright D' \subseteq M'$ and $\text{dom}(\mathfrak{A}) \subseteq \text{dom}(\mathfrak{B}) = D'$. Which gives us the desired conclusion :

$$\mathfrak{A}, \mu, D \mapsto D', M \mapsto M'$$
$$\models \mathbf{P}_{\text{min}} \subseteq M \wedge \forall x. D(x)$$

$\square$

Following [Etherington, 1986] (p.142, Theorem 6.3) and [De Kleer and Konolige, 1989], it is easy to see that, in addition to relational symbols, formulas that use fixed and varying symbols can also be minimised, that is: $\mathbf{P}_{\text{min}}$ can also contain formulas that do not use other symbols from $\mathbf{P}_{\text{min}}$. To sketch the construction: take a signature Σ partitioned into $\mathbf{P}_{\text{min}}, f_{\text{restr}}$, etc, such that an order $\preceq$ is defined. Given an FO Σ-formula θ that uses only symbols in $(=, \mathbf{P}_{\text{fix}}, \mathbf{P}_{\text{var}})$, we define an order $\preceq_\theta$: $\mathfrak{A} \preceq_\theta \mathfrak{B}$ if $\mathfrak{A} \preceq \mathfrak{B}$ and $[\![\theta]\!]_{\mathfrak{A}} \subseteq [\![\theta]\!]_{\mathfrak{B}}$. Extend the signature Σ to $\Sigma' := (\Sigma, P)$, where P is a fresh relational symbol of arity $|\langle \theta \rangle|$. P is added to the tuple $\mathbf{P}_{\text{min}}$ of Σ'. This gives rise to an order $\preceq'$ on Σ'-structures. The function f from Σ-structures to Σ'-structures which just extends any $\mathfrak{A} : \Sigma$ to Σ' by adding $[\![P]\!]_{f(\mathfrak{A})} := [\![\theta]\!]_{\mathfrak{A}}$ is a bijection between Σ-structures and Σ'-models of $(\forall x. P(x) \leftrightarrow \theta(x))$. Moreover *a)* $\mathfrak{A} \preceq_\theta \mathfrak{B}$ iff $f(\mathfrak{A}) \preceq' f(\mathfrak{B})$ and *b)* for any Σ'-formula F and $\mu : \langle F \rangle \to \text{dom}(\mathfrak{A})$, $f(\mathfrak{A}), \mu \models F$ iff $\mathfrak{A}, \mu \models F'$ where F' is F with all atoms of the form $P(x)$ replaced by $\theta(x)$. This applies in particular when $F = \mathcal{J}\phi$ for some FO, Σ'-formula ϕ, so the definition for circumscribed formulas page 93 admits having formulas (and not just symbols) in $\mathbf{P}_{\text{min}}$.

7.2. ⊴-minimality in 2nd-order logic

In this section we make the simple observation (using the previous section) that the minimisation operator $\downarrow$ can be expressed syntactically as an instance of unified circumscription. Let $\Delta\mathbf{Dyn}$ be the tuple of formulas of the form $\Delta A(\boldsymbol{x}) = \neg(A(\boldsymbol{x}) \leftrightarrow A^{\star}(\boldsymbol{x}))$ for $A \in \mathbf{Dyn}$ and $\boldsymbol{x}$ a tuple of variables of size the arity of A. We define an order $\preceq$ using the following mapping:

$$\mathbf{P}_{\min} = \Delta\mathbf{Dyn} \qquad \mathbf{P}_{\mathrm{restr}} = \mathbf{Stat} \setminus \mathbf{Stat}_{\mathrm{fun}} \qquad \mathbf{P}_{\mathrm{var}} = \mathbf{Dyn}^{\star}$$
$$\mathbf{P}_{\mathrm{fix}} = \mathbf{Dyn} \qquad \mathbf{f}_{\mathrm{restr}} = \mathbf{Stat}_{\mathrm{fun}} \qquad \mathbf{f}_{\mathrm{var}} = \emptyset$$

This defines a circumscription order $\preceq$ on Θ-structures.

Lemma 7.2.1. If $\mathfrak{A}, \mathfrak{B} : \Theta$, $\mathfrak{A} \preceq \mathfrak{B}$ iff $\mathfrak{A} \trianglelefteq \mathfrak{B}$

Proof. Cf. definitions for $\trianglelefteq$ (p. 34) and $\preceq$ (p. 92): (a) iff (4), (b) iff (3), (c) and (d) iff (2), and finally (e) iff (1).

$\square$

We restate the formula for $\downdownarrows$, adapted to the signature Θ and the mapping to $\mathbf{P}_{\min}$, etc:

$$\psi_1(D, \mathbf{Dyn}') := g(D) \subseteq D \wedge \mathbf{Dyn} \subseteq D \wedge (\exists x.\ D(x)) \wedge \bigwedge_{x \in \langle \phi \rangle} D(x)$$

$$\wedge\, \Delta\mathbf{Dyn}' \subseteq \Delta\mathbf{Dyn} \wedge \phi[D]\{\mathbf{Dyn}^{\star} \mapsto \mathbf{Dyn}'\}$$
$$\psi_2(D, \mathbf{Dyn}') := \Delta\mathbf{Dyn} \subseteq \Delta\mathbf{Dyn}' \wedge \forall x.\ D(x)$$
$$\downdownarrows \phi = \phi \ \wedge \forall D, \mathbf{Dyn}'.\ \psi_1(D, \mathbf{Dyn}') \rightarrow \psi_2(D, \mathbf{Dyn}')$$

Where $\mathbf{Dyn}'$ is a tuple similar to $\mathbf{Dyn}^{\star}$ and $\Delta\mathbf{Dyn}'$ is $\Delta\mathbf{Dyn}$ with every predicate in $\mathbf{Dyn}^{\star}$ replaced by its counterpart in $\mathbf{Dyn}'$. For instance, $\Delta\mathsf{Has}' = \neg(\mathsf{Has} \leftrightarrow \mathsf{Has}')$ (variables omitted).

We have only performed simple syntactic renaming and removed some second-order quantification since there are no variable functions ($\mathbf{f}_{\mathrm{var}}$).

By lemma 7.2.1 we have the following equivalence:

$$\downarrow \phi \equiv \phi \ \wedge \forall D, \mathbf{Dyn}'.\ \psi_1(D, \mathbf{Dyn}') \rightarrow \psi_2(D, \mathbf{Dyn}')$$

7.3. 2nd-order quantifier elimination

In this section we show when we can eliminate $\downarrow$ from FO[$\downarrow$]-formulas. We start from the observation that quantifying on subdomains is unnecessary when all the models of a formula are supported. Then, we define the *boundedness* condition on formulas, which restrict the amount of changes that can be present in their minimal models. Note that as seen earlier, the minimal models of any $\vdash_{\trianglelefteq}$-provable formula contain a bounded amount of changes. Modulo that boundedness property on a formula ϕ, all 2nd-order quantifiers can be removed from the 2nd-order formula equivalent to $\downarrow \phi$. During this quantifier elimination, we also notice that first-order quantifiers can be simplified modulo the theory of n-BFLs and $\vdash_{\rho}$-provability. As a result, every $\vdash_{\rho}$-provable, minimised formula is equivalent to a formula in $\exists^{*}\forall^{*}$ and to a formula in $\forall^{*}\exists^{*}$.

7.3.1. Removing domain quantification

In the sequel we will substitute relational symbols with formulas of the same arity. For instance, here we pass a formula $\theta(x)$ with 1 free variable as an argument to $\psi_1(\cdot, \cdot)$:

$$\psi_2(\theta, \mathbf{Dyn}') = \Delta\mathbf{Dyn}' \subseteq \Delta\mathbf{Dyn}' \wedge \forall x.\, \theta(x)$$

Definition. To keep things syntactically simple, we define a context-dependent formula. In any context where a symbol named Has' is defined, let $Supp'(x) := \mathsf{Has}(x) \vee \mathsf{Has}'(x)$. ∎

It is possible to simplify a circumscribed formula $\Lsh\phi$ if we assume all models are supported, i.e. that the pre and post-states cover the entire domain. This is what the next lemma shows:

Lemma 7.3.1. If $\phi \vDash Support$ then:

$$\Lsh\phi \equiv \phi \wedge \forall\mathbf{Dyn}'.\, \psi_1(Supp', \mathbf{Dyn}') \to \psi_2(Supp', \mathbf{Dyn}')$$

Proof. First, note the difference with $\Lsh\phi$ as described in the beginning of this section: the second-order quantification $\forall D$ is gone, and all atoms of the form $D(x)$ for x any variable have been replaced by $Supp'(x)$. This is well-defined since the replacement occurs within a quantification on $\mathbf{Dyn}'$ and there is a symbol $\mathsf{Has}' \in \mathbf{Dyn}'$.

($\Rightarrow$) Assume $\mathfrak{A}, \mu \vDash \forall D, \mathbf{Dyn}'.\, \psi_1(D, \mathbf{Dyn}') \to \psi_2(D, \mathbf{Dyn}')$. For any predicate tuple $\mathbf{M}$ such that $\mathfrak{A}, \mu, \mathbf{Dyn}' \mapsto \mathbf{M} \vDash \psi_1(Supp', \mathbf{Dyn}')$, we need $\mathfrak{A}, \mu, \mathbf{Dyn}' \mapsto \mathbf{M} \vDash \Delta\mathbf{Dyn} \subseteq \Delta\mathbf{Dyn}' \wedge Support$.

With H the image of Has' in $\mathbf{Dyn}' \mapsto \mathbf{M}$, it suffices to show $\mathfrak{A}, \mu, D \mapsto (\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \cup H), \mathbf{Dyn}' \mapsto \mathbf{M} \vDash \psi_1(D, \mathbf{Dyn}')$. Indeed, every conjunct of ψ_1 is trivially satisfied by assumption, so we are done.

($\Leftarrow$) For the other direction, we assume $\mathfrak{A}, \mu \vDash \forall\mathbf{Dyn}'.\, \psi_1(Supp', \mathbf{Dyn}') \to \psi_2(Supp', \mathbf{Dyn}')$.

Take any subset T of $\mathrm{dom}(\mathfrak{A})$ and any appropriately-sized tuple $\mathbf{M}$ of tuples of subsets of $\mathrm{dom}(\mathfrak{A})$ such that $\mathfrak{A}, \mu, D \mapsto T, \mathbf{Dyn}' \mapsto \mathbf{M} \vDash \psi_1(D, \mathbf{Dyn}')$.

Wlog, assume Has' is the first symbol of $\mathbf{Dyn}'$, and let H be the first component of $\mathbf{M}$. Let $\mathbf{M}_{HT}$ be $\mathbf{M}$ with its first component replaced with $H \cap T$.

We have $\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \subseteq T$ by ψ_1. So $\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \cup (H \cap T) \subseteq T$. By hypothesis on ϕ, we have $T \subseteq \llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \cup (H \cap T)$. So we get $T = \llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \cup (H \cap T)$.

- We first show $\mathfrak{A}, \mu, \mathsf{Has}' \mapsto \mathbf{M}_{HT} \vDash \psi_1(Supp', \mathbf{Dyn}')$.

  Remember that $T = \llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \cup (H \cap T)$. By assumption we do have $\llbracket g\rrbracket_{\mathfrak{A}}(T) \subseteq T$ and $\llbracket\mathbf{Dyn}\rrbracket_{\mathfrak{A}} \subseteq T$ as well as T nonempty.

  To show $\Delta\mathbf{M}_{HT} \subseteq \llbracket\Delta\mathbf{Dyn}\rrbracket_{\mathfrak{A}}$, remember that $\Delta\mathbf{M} \subseteq \llbracket\Delta\mathbf{Dyn}\rrbracket_{\mathfrak{A}}$. Suppose there is some e which is in $\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}}\Delta(H \cap T)$ but not in $\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}}\Delta H$. Since $H \cap T \subseteq H$, $e \in T$ is impossible. So $e \in \llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}}$. But by hypothesis, $\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \subseteq T$, which is absurd.

  Finally, we need $\mathfrak{A}, \mu, \mathbf{Dyn}' \mapsto \mathbf{M}_{HT} \vDash \phi[Supp'](\mathbf{Dyn}, \mathbf{Dyn}', \mathbf{Stat})$. We will use $\mathfrak{A}, \mu, D \mapsto T, \mathbf{Dyn}' \mapsto \mathbf{M} \vDash \phi[D](\mathbf{Dyn}, \mathbf{Dyn}', \mathbf{Stat})$:

  We show that $\mathfrak{A}, \mu, D \mapsto T, \mathbf{Dyn}' \mapsto \mathbf{M} \vDash \phi[D](\mathbf{Dyn}, \mathbf{Dyn}', \mathbf{Stat})$ iff $\mathfrak{A}, \mu, \mathbf{Dyn}' \mapsto \mathbf{M}_{HT} \vDash \phi[Supp'](\mathbf{Dyn}, \mathbf{Dyn}', \mathbf{Stat})$ by induction on ϕ. Note that, from left to right, we can assume that μ goes into T (by ψ_1), and from right to left, that μ goes into $\llbracket\mathsf{Has}\rrbracket_{\mathfrak{A}} \cup (H \cap T)$ (again by ψ_1).

  The only nontrivial case is $\phi = \mathsf{Has}^\star(t)$. Then, by induction on t:

- – $t = x$ then either $\mu(x) \in [\![\mathsf{Has}]\!]_{\mathfrak{A}} \subseteq T$ so $\mu(x) \in H$ iff $\mu(x) \in (H \cap T)$; or $\mu(x) \in (H \cap T)$ then we are done.

  – $t = f(\mathbf{t}')$, again by $[\![\mathbf{g}]\!]_{\mathfrak{A}}(T) \subseteq T$ and induction hypothesis we are done.

- Now that we know $\mathfrak{A}, \mu, \mathbf{Dyn}' \mapsto M_{\mathsf{HT}} \vDash \psi_1(\mathit{Supp}', \mathbf{Dyn}')$, we can use the main assumption to obtain $\mathfrak{A}, \mu, \mathbf{Dyn}' \mapsto M_{\mathsf{HT}} \vDash \Delta\mathbf{Dyn} \subseteq \Delta\mathbf{Dyn}' \wedge \mathit{Support}$. We will deduce $\mathfrak{A}, \mu, D \mapsto T, \mathbf{Dyn}' \mapsto M \vDash \Delta\mathbf{Dyn} \subseteq \Delta\mathbf{Dyn}' \wedge \forall x.\ D(x)$ from it:

  First, $T = [\![\mathsf{Has}]\!]_{\mathfrak{A}} \cup (H \cap T)$, $[\![\mathsf{Has}]\!]_{\mathfrak{A}} \Delta [\![\mathsf{Has}^\star]\!]_{\mathfrak{A}} \subseteq [\![\mathsf{Has}]\!]_{\mathfrak{A}} \Delta (H \cap T)$ (obtained from $\Delta\mathbf{Dyn} \subseteq \Delta\mathbf{Dyn}'$ with $\mathbf{Dyn}' \mapsto M_{\mathsf{HT}}$) and *Support*, yield $T = \mathrm{dom}(\mathfrak{A})$ through simple set-theoretic considerations.

  Let $\Delta M := [\![\mathbf{Dyn}]\!]_{\mathfrak{A}} \Delta M$ and $\Delta M_{\mathsf{HT}} := [\![\mathbf{Dyn}]\!]_{\mathfrak{A}} \Delta M_{\mathsf{HT}}$.

  By definition of M_{HT}, and since $[\![\Delta\mathbf{Dyn}]\!]_{\mathfrak{A}} \subseteq \Delta M_{\mathsf{HT}}$, we may only get $[\![\Delta\mathbf{Dyn}]\!]_{\mathfrak{A}} \not\subseteq \Delta M$ if there is some e in $[\![\Delta\mathsf{Has}]\!]_{\mathfrak{A}}$ (and thus in $[\![\mathsf{Has}]\!]_{\mathfrak{A}} \Delta (H \cap T)$), but not in $[\![\mathsf{Has}]\!]_{\mathfrak{A}} \Delta H$. By simple set-theoretic considerations, we get $e \in [\![\mathsf{Has}]\!]_{\mathfrak{A}}$ and $e \notin T$; the latter is absurd. So $[\![\Delta\mathbf{Dyn}]\!]_{\mathfrak{A}} \subseteq \Delta M$.

 $\square$

7.3.2. Removing 2nd-order quantification

Figure 7.2.: Precondition of $\mathfrak{A}_k$. What must the postcondition be to have $\mathfrak{A}_k \vDash \downarrow CC$?

Example. As an example of a minimised formula which is not FO-definable, we reprise the example from page 59 (formula reproduced below).

$$\downarrow CC := \downarrow \big(\mathsf{Blue}(c) \wedge$$
$$\forall x.\ (x \neq c \rightarrow \neg\mathsf{Blue}(x)) \wedge$$
$$\forall y.\ \mathsf{Blue}^\star(x) \rightarrow \mathsf{Link}^\star(x, y) \rightarrow \mathsf{Blue}^\star(y) \big)$$

with $\mathsf{Blue} \in \mathbf{Dyn}$. $CC(c)$ specifies *a)* that the interpretation of c is Blue and no other element is blue, and *b)* an element with a $\mathsf{Blue}^\star$ neighbor must be $\mathsf{Blue}^\star$.

Note that the query defined by $\mathsf{Blue}^\star(y) \wedge \downarrow CC$ captures exactly the elements accessible from $[\![c]\!]_{\mathfrak{A}}$.

To sketch a proof of non-definability in FO: all FO-definable queries are *Gaifman-local*, in the sense that for all such queries Q, there is some $d \geqslant 0$ such that (for the case of queries of arity 1)

Figure 7.3.: The d-balls around a and b are isomorphic

for all structures $\mathfrak{A}$ and $a, b \in \mathrm{dom}(\mathfrak{A})$, if the d-balls around a and b are isomorphic (with a the image of b in the isomorphism), then $a \in Q(\mathfrak{A})$ iff $b \in Q(\mathfrak{A})$. For more on Gaifman-locality, see [Libkin, 2013], chapter 4. We skip over the fact that Gaifman locality is defined for relational signatures and our signatures are non-relational since functions are trivial to encode with first-order formulas and relational symbols.

Now for any d, take any n-BFL $\mathfrak{A}$ with two connected components of size $k > d$, as in figure 7.3 (we only see the preconditions in that figure). c is at one end of one of the connected components. Let a be the other extremity of c's connected component, and let b be any extremity of the other. The d-balls around a and b are isomorphic, but $\mathfrak{A}, y \mapsto a \vDash \mathrm{Blue}^\star(y) \wedge {\downarrow} CC$ and $\mathfrak{A}, y \mapsto b \nvDash \mathrm{Blue}^\star(y) \wedge {\downarrow} CC$ (cf. figure 7.2). So $\mathrm{Blue}^\star(y) \wedge {\downarrow} CC$ is not FO-definable.

We have allowed the interpretation of $\mathrm{Link}^\star$ to be of degree 2 to make things easier; the same construction works on n-BFLs by using trees of height 1.

■

Example. We will soon show that bounding the amount of changes makes any minimised formula FO-definable. However, there are minimised formulas with unbounded changes that are FO-definable. A very simple example :

$$\mathit{Switch} := \forall x.\ \neg A(x) \wedge A^\star(x)$$

Any model $\mathfrak{A}$ of ${\downarrow}\mathit{Switch}$ contains $|\mathrm{dom}(\mathfrak{A})|$ changes.

■

We give a semantic constraint under which 2nd-order quantification can be removed. The idea is that if, for a formula ϕ, there is a bound on the amount of changes contained in its minimal models, then second-order quantification in ${\downarrow}\phi$ can be replaced by a finite enumeration.

Definition. If $\beta : \mathbf{Dyn} \to \mathbb{N} = A \mapsto \beta_A$, ϕ is formula, ψ is a sentence, ϕ is β-*bounded* if for every dynamic predicate $A \in \mathbf{Dyn}$, for any $\mathfrak{A}, \mu \vDash \psi \wedge {\downarrow}\phi$, $|[\![\Delta A]\!]_\mathfrak{A}| \leqslant \beta_A$. β_A is then a bound for A in ϕ modulo ψ.

■

In other words, among the $\trianglelefteq$-minimal models of a bounded formula ϕ, there is a bound on the number of elements that can change.

We say that a formula is bounded modulo ψ if there exists a β such that the formula is β-bounded modulo ψ. We skip the "modulo" part to mean that ψ is valid.

Definition. If ϕ is β-bounded modulo ψ, define $\mathbf{X}_A^\beta := \mathbf{x}_1, \ldots, \mathbf{x}_{\beta_A}$, a tuple *of tuples* of fresh variables, with $|\mathbf{x}_i| = \mathrm{arity}(A)$ for $1 \leqslant i \leqslant \beta_A$.

■

Definition. If ϕ is β-bounded modulo ψ and $A \in \mathbf{Dyn}$, define an an associated *bounding formula*, $\mathit{Bound}(A)$:

$$\mathit{Bound}(A) := (\forall x.\ A^\star(x) \leftrightarrow A(x)) \vee \mathit{Diff}(A^\star, A, \mathbf{X}_A^\beta)$$

where for all relational symbols S, S' of same arity a, for all $\mathbf{X_k} := \mathbf{x_1}, \ldots, \mathbf{x_k}$ with $k \geqslant 0$ and such that $|\mathbf{x_i}| = a$ for $1 \leqslant i \leqslant k$:

$$Diff(S', S, \mathbf{X_k}) := \forall \mathbf{x}. \ S'(\mathbf{x}) \leftrightarrow Split(S, \mathbf{X_k}, \mathbf{x})$$
$$Split(S, \mathbf{X_k}, \mathbf{x}) := (S(\mathbf{x}) \wedge \mathbf{x} \notin \mathbf{X_k}) \vee (\neg S(\mathbf{x}) \wedge \mathbf{x} \in \mathbf{X_k})$$

The quantified $\mathbf{x}$ also has size a. $\mathbf{x} \in \mathbf{X_k}$ is shorthand for $\mathbf{x} = \mathbf{x_1} \vee \ldots \vee \mathbf{x} = \mathbf{x_k}$. Again, note that each $\mathbf{x_i}$ is a tuple of variables (of size a) and that $\mathbf{X_k}$ is a k-sized tuple *of tuples* of variables. ∎

A bounding formula for A declares that there are at most β_A tuples such that every tuple involved in a change (i.e. such that it is true in A and false in $A^\star$, or the reverse) is also one of the β_A tuples, and that those are the only such tuples.

More precisely:

- *Diff* takes predicates S' and S as well as tuples $\mathbf{X_k} = \mathbf{x_1}, \ldots, \mathbf{x_k}$. It states that a tuple a of elements (in some structure $\mathfrak{A}$ satisfying the formula) is in $[\![S']\!]_{\mathfrak{A}} \Delta [\![S]\!]_{\mathfrak{A}}$ iff it is equal to the interpretation of some $\mathbf{x_i}$.

- *Bound* assumes the existence of some bound β_A for its argument A and specifies that either $A^\star$ and A are equal, or $Diff(A^\star, A, \mathbf{X_A^\beta})$ holds. The first part ("$A^\star$ iff A") is necessary because while $\mathbf{X_A^\beta}$ must be interpreted by some actual tuples, β_A is an overapproximation of the size of the difference between A and $A^\star$; so it's necessary to explicitly allow for no difference at all.

Lemma 7.3.2. If ϕ is β-bounded modulo ψ, then for any $A \in \mathbf{Dyn}$, if $\mathfrak{A}, \mu \vDash \psi \wedge \downarrow \phi$ then there is M_A such that $\mathfrak{A}, \mathbf{X_A^\beta} \mapsto M_A \vDash Bound(A)$.

Proof. By minimality of $\mathfrak{A}$ and boundedness of ϕ modulo ψ, there are $k \leqslant \beta_A$ tuples in $[\![\Delta A]\!]_{\mathfrak{A}}$. If $k = 0$ then $\mathfrak{A} \vDash \forall \mathbf{x}. \ A(\mathbf{x}) \leftrightarrow A^\star(\mathbf{x})$. Otherwise let $\{\mathbf{p_1}, \ldots, \mathbf{p_k}\} = [\![\Delta A]\!]_{\mathfrak{A}}$. Define $M_A := \mathbf{u_1}, \ldots, \mathbf{u_{\beta_A}}$ where $\mathbf{u_i} := \mathbf{p_{i \bmod k}}$. We must show:

$$\mathfrak{A} \vDash \forall \mathbf{x}. \ A^\star(\mathbf{x}) \leftrightarrow ((A(\mathbf{x}) \wedge (\bigwedge_i \mathbf{x} \neq \mathbf{u_i})) \vee (\neg A(\mathbf{x}) \wedge (\bigvee_i \mathbf{x} = \mathbf{u_i})))$$

(This is just an inlining of the definition of *Bound*). If $e \in [\![A^\star]\!]_{\mathfrak{A}}$, then either $e \in [\![\Delta A]\!]_{\mathfrak{A}}$, so $e \notin [\![A]\!]_{\mathfrak{A}}$ and there is $1 \leqslant j \leqslant k$ such that $e = \mathbf{p_j} = \mathbf{u_j}$, and so $\mathfrak{A}, \mathbf{x} \mapsto e \vDash \neg A(\mathbf{x}) \wedge (\bigvee_i \mathbf{x} = \mathbf{u_i})$, or $\mathbf{x} \notin [\![\Delta A]\!]_{\mathfrak{A}}$, so $e \in [\![A]\!]_{\mathfrak{A}}$ and for all $1 \leqslant i \leqslant k$, $e \neq \mathbf{p_i}$, so for all $1 \leqslant j \leqslant \beta_A$, $e \neq \mathbf{u_i}$. So $\mathfrak{A}, \mathbf{x} \mapsto e \vDash A(\mathbf{x}) \wedge (\bigwedge_i \mathbf{x} \neq \mathbf{u_i})$.

The case for $e \notin [\![A^\star]\!]_{\mathfrak{A}}$ is similar. □

Lemma 7.3.3. If ϕ is β-bounded and supported modulo ψ, then $\downarrow \phi$ is first-order definable modulo ψ.

Proof. In the following we assume variable freshness wherever necessary.

For any $\mathfrak{A}, \mu \vDash \psi \wedge \downarrow \phi$, any $A \in \mathbf{Dyn}$, by lemma 7.3.2, we get $\mathfrak{A}, \mathbf{X_A^\beta} \mapsto M_A \vDash Bound(A)$ for some $M_A = \mathbf{m_1}, \ldots, \mathbf{m_{\beta_A}}$. Consider any set U of arity(A)-sized tuples from $dom(\mathfrak{A})$ such that $[\![A]\!]_{\mathfrak{A}} \Delta U \subseteq [\![\Delta A]\!]_{\mathfrak{A}}$. For some $0 \leqslant q \leqslant \beta_A$ there are $\mathbf{x_{k_1}}, \ldots, \mathbf{x_{k_q}} \in \mathbf{X_A^\beta}$ such that

$$\mathfrak{A}, S' \mapsto U, \mathbf{x_{k_1}} \mapsto \mathbf{m_{k_1}}, \ldots, \mathbf{x_{k_q}} \mapsto \mathbf{m_{k_q}} \vDash Diff(S', A, (\mathbf{x_{k_1}}, \ldots, \mathbf{x_{k_q}}))$$

Note how the case $q = 0$ makes $Diff(S', A, ())$ degenerate to $\forall \mathbf{x}. \ S'(\mathbf{x}) \leftrightarrow A(\mathbf{x})$.

Moving from the meta language to the logic, we restate the above as:

$$\mathfrak{A}, S' \mapsto U \models \exists X_A^\beta.\ Bound(A) \wedge \bigvee_{X \subseteq X_A^\beta} Diff(S', A, X)$$

Since $\psi_1(Supp', \mathbf{Dyn}') \models \Delta A' \subseteq \Delta A$, for any $A' \in \mathbf{Dyn}'$, modulo *Bound* and ψ we have:

$$\exists X_A.\ Bound(A) \wedge \forall A'.\ \psi_1(Supp', \mathbf{Dyn}') \to \psi_2(Supp', \mathbf{Dyn}')$$

$$\equiv \exists X_A.\ Bound(A) \wedge \forall A'.\ \psi_1(Supp', \mathbf{Dyn}') \to \left(\psi_2(Supp', \mathbf{Dyn}') \wedge \bigvee_{X \subseteq X_A^\beta} Diff(A', A, X) \right)$$

$$\equiv \exists X_A.\ Bound(A) \wedge \forall A'.\ \bigwedge_{X \subseteq X_A^\beta} \left(Diff(A', A, X) \wedge \psi_1(Supp', \mathbf{Dyn}') \right) \to \psi_2(Supp', \mathbf{Dyn}')$$

$$\equiv \exists X_A.\ Bound(A) \wedge$$
$$\forall A'.\ \bigwedge_{X \subseteq X_A^\beta} \left(\psi_1(Supp', \mathbf{Dyn}') \to \psi_2(Supp', \mathbf{Dyn}') \right) \{A'(x) \mapsto Split(A, X, x)\}$$

Note that the substitution $\{A'(x) \mapsto Split(A, X, x)\}$ captures the arguments of A' and uses them as arguments for *Split*; the substitution is simply applying the equivalence $Diff(A', A, X) = \forall x.\ A'(x) \leftrightarrow Split(A, X, x)$ at the syntax level.

Since all occurrences of A' have been removed, the $\forall A'$. can be dropped. Starting from the formula provided by lemma 7.3.1, repeated application of the above for all $A \in \mathbf{Dyn}$ removes all 2nd-order quantification from $\downarrow\phi$: As syntactic sugar, define

$$\psi := \psi_1(Supp', \mathbf{Dyn}') \to \psi_2(Supp', \mathbf{Dyn}')$$

With a bounded support formula ϕ, and $\mathbf{Dyn} = (A_1, \dots, A_p)$, for any $1 \leqslant i \leqslant p$, any tuple Y of arity(A_i)-sized tuples of variables, define the substitution

$$\sigma_i(Y) := A'_i(x) \mapsto Split(A_i, Y, x)$$

Modulo ψ, we have:

$$\downarrow\phi \equiv \downarrow\phi \wedge \bigwedge_{A_i \in \mathbf{Dyn}} \exists X_{A_i}^\beta.\ Bound(A_i)$$

$$\equiv \exists X_{A_1}^\beta, \dots, X_{A_m}^\beta.\ \downarrow\phi \wedge Bounds \qquad\qquad (\text{where } Bounds := \textstyle\bigwedge_{A_i \in \mathbf{Dyn}} Bound(A_i))$$

$$\equiv \exists X_{A_1}^\beta, \dots, X_{A_m}^\beta.\ x\phi \wedge (\forall A'_0, \dots, A'_m.\ \psi) \wedge Bounds$$

$$\equiv \exists X_{A_1}^\beta, \dots, X_{A_m}^\beta.\ \phi \wedge \left(\forall A'_0, \dots, A'_{n-1}.\ \bigwedge_{Y_{A_m} \subseteq X_{A_m}^\beta} \psi\{\sigma_m(Y_{A_m})\} \right) \wedge Bounds$$

$$\vdots$$

$$\equiv \exists X_{A_1}^\beta, \dots, X_{A_m}^\beta.\ \phi \wedge \left(\bigwedge_{Y_{A_1} \subseteq X_{A_1}^\beta} \cdots \bigwedge_{Y_{A_m} \subseteq X_{A_m}^\beta} \psi\{\sigma_m(Y_{A_m})\} \dots \{\sigma_1(Y_{A_1})\} \right) \wedge Bounds$$

$\square$

7.3.3. Simplifying FO quantifiers

We can show the level of quantifier nesting is not increased much by minimisation.

Lemma 7.3.4. If ϕ is supported, ϕ is bounded modulo ψ, η is a first-order quantifier alternation prefix, and $\phi \in \eta$ and $\neg\phi \in \eta$, then $\downarrow\phi$ is equivalent modulo ψ to an FO formula with a quantifier prefix in $\exists^*\eta\forall^*$.

Proof. By inspection of the first-order formula equivalent to $\downarrow\phi$ (modulo ψ) given by lemma 7.3.3, and by trivial syntactic manipulation: $\bigwedge_{Y_{A_i} \subseteq X_{A_i}} (\psi_1(Supp', \mathbf{Dyn'}) \to \psi_2(Supp', \mathbf{Dyn'}))\{A_i'(x) \mapsto Split(A_i, Y_{A_i}, x)\}$ is a conjunction of formulas with quantifier prefix in η (by assumption on ϕ), and $Bound(A_i)$ contains only universal quantifiers in prenex form. $\square$

Definition. Let $d \geqslant 0$. In an n-BFL, the formula $SameTree(x,y)$ specifies that x and y are in the same tree. The formula $\delta_{\leqslant d}(x,y)$ specifies that x and y are at distance $\leqslant d$. We skip the proofs for each.

$$SameTree(x,y) := parent^n(x) = parent^n(y)$$

$$\delta_{\leqslant d}(x_0, y_d) := \exists\, \{y_{d'-1}, x_{d'} \mid 0 < d' \leqslant d\}. \left(\bigwedge_{0 \leqslant i \leqslant d} SameTree(x_i, y_i)\right)$$

$$\wedge \left(\bigwedge_{0 \leqslant i < d} y_i = x_{i+1} \vee Link(y_i, x_{i+1}) \vee Link^\star(y_i, x_{i+1})\right)$$

The goal of the following lemma is to remove an existential quantifier in front of a formula. Recall that $\ell_{\Theta_n}(d)$ bounds the size of any d-sized ball in any $\mathfrak{A} : \Theta_n$ (cf. lemma 4.3.6 p.67).

Lemma 7.3.5. If $d \geqslant 0$, and ψ is a formula then $\exists x.\, \delta_{\leqslant d}(x,y) \wedge \psi$ is equivalent modulo $\mathcal{T}_n$ to:

$$\bigvee_{1 \leqslant k \leqslant \ell_{\Theta_n}(d)} \left(\left(\exists x_1, ..., x_k.\, \bigwedge_{1 \leqslant i \leqslant k} \delta_{\leqslant d}(x_i, y) \wedge \bigwedge_{1 \leqslant j \neq i \leqslant k} x_i \neq x_j\right)\right.$$

$$\left.\wedge \forall x_1, ..., x_k.\, \left(\bigwedge_{1 \leqslant i \leqslant k} \delta_{\leqslant d}(x_i, y) \wedge \bigwedge_{1 \leqslant j \neq i \leqslant k} x_i \neq x_j\right) \to \bigvee_{1 \leqslant i \leqslant k} \psi\{x \mapsto x_i\}\right)$$

Proof. By lemma 4.3.6, there are at most $\ell_{\Theta_n}(d)$ nodes at distance at most d from y.

Suppose $\mathfrak{A}, \mu \vDash \psi$ and $\delta_{\mathfrak{A}}(\mu(x), \mu(y)) \leqslant d$. Let $1 \leqslant k \leqslant \ell_{\Theta_n}(d)$ be the number of elements at distance at most d from $\mu(y)$. The kth disjunct of the above sentence is satisfied by taking those k elements as existential witnesses of for $x_1, \ldots, x_k$ (left conjunct): they are at distance $\leqslant d$ from y, all different. The only universal witnesses for $x_1, \ldots, x_k$ (right conjunct) that satisfy the premise are those k elements, and one of them is $\mu(x)$.

The other way around, let θ_k be the satisfied disjunct; by the existential part of θ_k, its universal part is not vacuously true. $\square$

Each disjunct in the formula above supposes a different number of distinct elements in $\mathcal{B}_{\mathfrak{A}}(\mu(y), d)$. For each possible answer, the universal part affirms that at least of those elements can play the role of x in ψ. We could reduce the number of disjunctions by counting the number of trees around an element instead of the number of elements; then the $x_1, \ldots, x_k$ would be tree roots and the final disjunct ($\psi\{x \mapsto x_i\}$) would be over all descendants of the correct root.

Theorem 7.3.6. If ϕ is $(\mathcal{V}, d)$-preserved then $\phi \wedge \textit{Support}_D$ is bounded modulo $\mathfrak{T}_n$.

Proof. Let $\mathfrak{A} : \Theta_n$ such that $\mathfrak{A}, \mu \models {\downarrow}(\phi \wedge \textit{Support}_D)$. Let $A \in \mathbf{Dyn}$, $e \in [\![\Delta A]\!]_{\mathfrak{A}}$ and $a \in e$. By lemma 4.3.5, if $a \notin \mathcal{B}_{\mathfrak{A}}(\mu(\mathcal{V}), d+1)$ then there is a $(\mu(\mathcal{V}), d)$-sub $\mathfrak{B}$ of $\mathfrak{A}$ such that $\mathfrak{B} \neq \mathfrak{A}$ and $\mathfrak{B} \vDash \textit{Support}_D$. So $\mathfrak{B} <_\alpha \mathfrak{A}$, by definition of subs. By lemma 4.3.3, $\mathfrak{B} : \Theta_n$. And since ϕ is $(\mathcal{V}, d)$-preserved, $\mathfrak{B}, \mu \vDash \phi$. Absurd by minimality of $\mathfrak{A}$. So $[\![\Delta A]\!]_{\mathfrak{A}} \subseteq \mathcal{B}_{\mathfrak{A}}(\mu(\mathcal{V}), d+1)^{\mathrm{arity}(A)}$.

By lemma 4.3.6, $|\mathcal{B}_{\mathfrak{A}}(\mu(\mathcal{V}), d+1)| \leqslant |\mathcal{V}| \ell_{\Theta_n}(d+1)$. $\qquad\square$

Lemma 7.3.7. If ϕ' is $(\mathcal{V}, d)$-preserved, in $\exists^* \forall^*$ and in $\forall^* \exists^*$ then ${\downarrow}(\phi' \wedge \textit{Support}_D)$ is in $\exists^* \forall^*$ and in $\forall^* \exists^*$ modulo $\mathfrak{T}_n$.

Proof. Let $\phi := \phi' \wedge \textit{Support}_D$. For ${\downarrow}\phi \in \exists^* \forall^*$ modulo $\mathfrak{T}_n$, we get β-boundedness modulo $\mathfrak{T}_n$ (for some $\beta : \mathbf{Dyn} \to \mathbb{N}$) with lemma 7.3.6, ϕ is supported by definition, and so we get the syntactic fragment with lemma 7.3.4.

Inspecting the formula equivalent to ${\updownarrow}\phi$ modulo $\mathfrak{T}_n$ as shown in 7.3.3, we have:

$$\downarrow\phi \equiv \exists X_{A_1}, ..., X_{A_m}. \ F$$

where

$$F := \phi \wedge \left(\bigwedge_{Y_{A_1} \subseteq X_{A_1}} \cdots \bigwedge_{Y_{A_m} \subseteq X_{A_m}} \psi\{\sigma_m(Y_{A_m})\} \ldots \{\sigma_1(Y_{A_1})\} \right) \wedge \textit{Bounds}$$

It is easy to syntactically check that $F \in \forall^* \exists^*$ (remember that $\phi' \in \exists^* \forall^*$, $\phi' \in \forall^* \exists^*$, and that $\textit{Support}_D$ is a universal sentence). Now, we show that each witness in $\exists X_{A_i}$ can be always be found at distance $\leqslant d+1$ from $\mu(\mathcal{V})$ for any $\mathfrak{A} : \Theta_n$ such that $\mathfrak{A}, \mu \models {\downarrow}\phi$.

As a reminder, $\textit{Bound}(A_i) := (\forall x. \ A_i^\star(x) \leftrightarrow A_i(x)) \vee \textit{Diff}(A^\star, A_i, X_{A_i})$ and is satisfied for every $A_i \in \mathrm{Dyn}$. Let $\mathfrak{A}, (X_{A_i} \mapsto M_{A_i})_{1 \leqslant i \leqslant m}, \mu \vDash F$. If $\textit{Bound}(A_i)$ is satisfied by its 1st disjunct, then $[\![\Delta A_i]\!]_{\mathfrak{A}} = \emptyset$ and the choice of witness is irrelevant.

Otherwise, by *Diff*, the interpretations for X_{A_i} are exactly the members of $[\![\Delta A_i]\!]_{\mathfrak{A}}$. By lemma 4.3.5, for every $e \in [\![\Delta A]\!]_{\mathfrak{A}}$ and every $a \in e$, there is some $c \in \mu(\mathcal{V})$ such that $\delta_{\mathfrak{A}}(a, c) \leqslant d+1$.

By the above, we have modulo $\mathfrak{T}_n$:

$$\downarrow\phi \equiv \exists X_{A_1}, ..., X_{A_m}. \left(\bigwedge_{1 \leqslant i \leqslant m} \bigwedge_{x \in X_{A_i}} \bigwedge_{x \in x} \bigvee_{y \in \mathcal{V}} \delta_{\leqslant d+1}(x, y) \right) \wedge F$$

Repeated application of lemma 7.3.5 yields that modulo $\mathfrak{T}_n$, ${\downarrow}\phi$ is equivalent to a disjunction of formulas with the following shape:

$$(\exists X_{A_1}, ..., X_{A_m}. \ G_1) \wedge \forall X_{A_1}, ..., X_{A_m}. \ G_1 \to F'$$

where G_1 is existential and $F' \in \forall^* \exists^*$. G_1 is a conjunction of existential formulas of the form $\delta_{\leqslant d+1}(x, y)$ or $x \neq y$. F' is a disjunction of formulas in $\forall^* \exists^*$ (see the details of the construction in lemma 7.3.5), so ${\downarrow}\phi \in \forall^* \exists^*$ modulo $\mathfrak{T}_n$. $\qquad\square$

We can move to the main result of this section:

7.3.4. Main theorem for elimination of $\downarrow$

Theorem 7.3.8. If $\mathcal{V}; d \vdash_{\lhd} \phi$, then $\phi \wedge \mathcal{T}_n$ is in both $\exists^*\forall^*$ and $\forall^*\exists^*$.

Proof. Remember that $\mathcal{T}_n$ is a universal sentence, so it suffices to prove that ϕ is in both $\exists^*\forall^*$ and $\forall^*\exists^*$ modulo $\mathcal{T}_n$. We proceed by induction on the derivation, proving for all $\mathcal{V}, d$.

Dynamic$_{\lhd}$, Static$_{\lhd}$, Invariant$_{\lhd}$, Weak$_{\lhd}$, Bool$_{\lhd}$, Guard$_{\lhd}$ Trivial.

$\exists$-**Guard$_{\lhd}$** $\phi = \exists x.\ \alpha(x,y) \wedge \phi' \in \exists^*\forall^*$ by induction hypothesis on ϕ'. Moreover, $x \neq y$, so by functionality of Link, Link* (given by $\mathcal{T}_n$) and equality, we have $\phi \equiv (\exists x.\ \alpha(x,y)) \wedge (\forall x.\ \alpha(x,y) \to \phi)$; since $\phi' \in \forall^*\exists^*$ by induction hypothesis, $\phi \in \forall^*\exists^*$ as well.

$\forall_{\lhd}$ $\forall x.\ \phi \in \forall^*\exists^*$ using the induction. For $\exists^*\forall^*$: no subformula of ϕ is ψ such that $x \in \langle \psi \rangle$ and of the form $\psi = \downarrow \psi'$ since Circum$_{\lhd}$ adds all variables to $\mathcal{V}$ and adds 1 to d, and Invariant$_{\lhd}$ (which removes variables from $\mathcal{V}$) requires a formula with $d = 0$. So $\phi \in$ FO. We now proceed by induction over the number of $\exists$ quantifiers under $\forall x$. .

We put ϕ in conjunctive normal form, treating quantified $\exists$ subformulas as atoms, and distributing $\forall x$. over $\wedge$, pushing it inside other toplevel $\forall$, and pushing universal quantifiers up. Note that by distributing $\forall x$. over $\wedge$, the number of $\exists$ under $\forall x$. does not increase.

Every conjunct of the resulting formula is of the form

$$\forall x.\ \bigvee \theta_i$$

where each θ_i is a literal or a formula of the form

$$\exists y.\ \alpha(y,z) \wedge \theta_i'$$

where $\mathcal{V}', d' \vdash_{\lhd} \theta_i'$ for some $\mathcal{V}', d'$, since the existential subformulas have been treated as atoms, pre formulas are all provable (for Invariant$_{\lhd}$) and guards are all provable (for Guard$_{\lhd}$). Moreover $x \neq z$: by assumption $x \notin \mathcal{V}$ (otherwise $\forall_{\lhd}$ could not have been applied), and $\exists$-Guard$_{\lhd}$ adds z to the context. The only rule that could have removed z is Invariant$_{\lhd}$, but that requires $d = 0$, and $\exists$-Guard$_{\lhd}$ increases d by 1 (rules are monotone w.r.t. d). Now, by induction on the number of θ_i' where x occurs. If there is none, every $\exists$ and $\forall x$. can all be swapped and we are done.

Otherwise the clause is of the form

$$\forall x.\ (\exists y.\ \alpha(y,z) \wedge \theta_i') \vee \theta$$

where x occurs in θ_i'. By the same argument as in the previous case ($\exists$-Guard$_{\lhd}$), this is equivalent modulo $\mathcal{T}_n$ to:

$$(\forall x.\ \theta \vee \exists y.\ \alpha(y,z)) \wedge (\forall x.\ \theta \vee \forall y.\ \alpha(y,z) \to \theta_2)$$
$$\equiv (\forall x.\ \theta \vee \exists y.\ \alpha(y,z)) \wedge \forall y.\ (\forall x.\ \theta \vee \neg\alpha(y,z) \vee \theta_2) \qquad \text{(up to variable renaming)}$$

In the left conjunct, there is one less θ_i' where x occurs, so done by the innermost induction hypothesis. In the right conjunct, there is one less $\exists$ under $\forall x$. and $\theta \vee \alpha(y,z) \vee \theta_2$ is $\vdash_{\lhd}$-provable, so done by the outermost induction hypothesis.

Circum$_\lhd$ Lemma 4.3.3 implies that *Symlink* is $(\emptyset, 0)$-preserved, so by the premise of Circum$_\lhd$ and theorem 4.4.5, $\phi \wedge Symlink$ is $(\mathcal{V}, \mathrm{d})$-preserved. Moreover, *Symlink* is a universal sentence so $\phi \wedge Symlink$ is in $\exists^* \forall^*$ and $\forall^* \exists^*$. So by lemma 7.3.7 (applicable thanks to the induction hypothesis), and Weak$_\lhd$, we are done.

$\square$

8. Conclusion

8.1. Related work

Circumscription   The use of circumscription for nonmonotonic reasoning dates back to McCarthy [McCarthy, 1980], and attempts at inferring non-changes logically can be found by following mentions of the frame problem [Reiter, 1991]. We use an instanciation of unified circumscription, a new flavor of circumscription which generalises [Doherty et al., 1998]. Previous works on taming circumscription require global syntactic properties of the formulas [Doherty et al., 1997, Nonnengart et al., 1999, Conradie, 2006] and only consider satisfiability or FO-definability of circumscribed formulas; we introduce a modular method for building circumscribed formulas that live in a tight fragment of FO. The resulting classes of formulas are incomparable with those presented here. Moreover, the goal of circumscription is usually to approximate commonsense reasoning. In contrast, we want to use nonmonotonicity for the purpose of graph rewriting rule synthesis.

There is a wide variety of forms of circumscription, and some make the connection to contemporary logic programming by representing semantics such as stable model and FLP [Ferraris et al., 2007, Ferraris et al., 2011, Wan et al., 2014, Zhang and Ying, 2010, Yang et al., 2011, Bartholomew et al., 2011]; we hope other connections can be made between the form of circumscription we developed here and these fields.

Biological modelling   Logical approaches to biological and chemical modelling exist in many flavors, using linear logic or other resource-aware logics [Despeyroux, 2016, Boniolo et al., 2010], or logic rules for specification and modality for queries [Eker et al., 2002], or explicitly representing reactions as first-class citizens and using circumscription for describing weakest preconditions to reactions [Doherty et al., 2004]. The focus, so far, has been on verification of models or pure knowledge representation. Here, we offer a classically model-based approach and maintain intuitive semantics. We follow the tradition of taking graph rewriting as a domain-specific language for biology [Boutillier et al., 2018, Chylek et al., 2014, Faeder et al., 2009] and we try to retain enough expressivity for large-scale, natural knowledge representation while ensuring reasonable tractability. More generally, logical approaches such a boolean networks are widely used to model gene interactions [Novère, 2015].

Another approach is [Basso-Blandin et al., 2015], which extends Kappa using slice categories over a particular kind of graph (reminiscent of bigraphs) and their homorphisms. The slice is over a master graph which represents all possible entities and their interactions. "Nuggets" of knowledge are given as graphs together with a homomorphism towards the master. Nuggets of knowledge are updated using categorical operations, and in some cases a set of knowledge nuggets can be compiled to Kappa. This approach seems promising for the connection between ontologies in biology [Stevens et al., 2000, Bard, 2004] and rule-based execution.

Transition structures   We compare with two well-known approaches. Hoare logic is often about verifying properties of programs, by obtaining a proof of {P}p{Q} for statements P, Q and a

program p. Corresponding assertions in our setting would be of the form $P^\circ \to Q^\star$, with a question mark as a placeholder for p: we are concerned with synthesising p from a specification. The p we wish to synthesize, however, is not any program satisfying $P^\circ \to Q^\star$; we want it to be among the "best" programs, in the sense that it also edits the memory as little as possible. Other differences are that our programs run in constant time, and that we want to mix pre- and post-conditions more freely. There have been approaches to Hoare logic-based program synthesis ([Mamouras, 2016]); in the future work section we describe an idea for "rule merging" which would yield something similar to a Hoare sequence rule $\{P\}p\{Q\}, \{Q\}p'\{R\} \Rightarrow \{P\}p; p'\{R\}$.

A logic on pre- and post-conditions such as the one introduced here naturally makes one think of an extension to more frames, either by adding new copies of **Dyn**, such as $\mathbf{Dyn}^{\star\star\cdots}$, or in the style of [Reiter, 2001] by adding natural numbers and an additional dimension to every dynamic statement: one can then write for instance $A(7, x)$ instead of $A(x)$ to mean that $A(x)$ is true at frame 7, or even $\forall i : \text{int. } A(i, x) \to A(i + 1, x)$. A modal version of this approach would be dynamic logic, but 1) we see this approach as a logic on *traces* rather on rules, about which see [Laurent et al., 2018] for a trace query language on Kappa traces, and 2) the crucial notion of minimality, especially on an unbounded number of frames, could resurrect variants of the frame problem [Reiter, 1991].

Guarded fragment   Many constructions in this thesis are based on *guards*. After the introduction of $\vdash_{\lhd}$, they become quite constrained (they are morally of bounded degree) anyway, but before that, one could try to extend the correctness results on $\vdash_\rho$ to looser guards, such as the loosely guarded fragment [Hodkinson, 2002], or even the guarded negation fragment [Bárány et al., 2011].

8.2. Conclusion

This thesis is an attempt to lay the foundations for a more capable molecular biology modelling framework. It takes rule-based modelling as a starting point and extends it through logic. Models are transitions between states. Transition signatures are separated into precondition $(\text{Link}, \text{Has}, \text{A}, \ldots)$, postcondition $(\text{Link}^\star, \text{Has}^\star, \text{A}^\star, \ldots)$ and static information $(\text{N}, \text{f}, \text{parent})$. For instance, $\text{Enzyme}(x) \wedge \neg\text{Active}(x) \wedge \text{Active}^\star(x)$ means that x is an enzyme and becomes active. Any context is allowed. Anything else can happen. State signature drop the postcondition part.

A reasonable specialisation is the class of n-BFLs, where the static part describes height-bounded trees and where $\text{Link}, \text{Link}^\star$ are functional. The static part represents the inner protein structure, while the dynamic part represents transient protein state and protein-protein interactions.

Even modulo the theory of n-BFLs, first-order satisfiability is not decidable. But decidability in the $\exists^*\forall^*$ fragment is decidable.

There is a change order $\unlhd$ between transitions. If $\mathfrak{A} \unlhd \mathfrak{B}$, $\mathfrak{A}$ and $\mathfrak{B}$ start from the same state but $\mathfrak{A}$ "does less" than $\mathfrak{B}$. This order is relevant to biological modelling because execution of protein-protein interactions requires a closed-world policy regarding which changes can occur. An operator $\downarrow$ is added to first-order logic. For any formula ϕ, $\downarrow\phi$ is the set of $\unlhd$-minimal models of ϕ. $\downarrow$ represents the end of the data gathering phase and the beginning of the execution phase.

Circumscription is a class of 2nd-order macros created by [McCarthy, 1980], used to capture minimal models of formulas along various orders. Unified circumscription generalises various common forms of circumscription, and can be instantiated to define $\downarrow$. So it is possible to define $\downarrow$ syntactically using 2nd-order logic.

Transitions can be summarised by transition patterns, of the form $(\mathfrak{A}, K)$ with $K \subseteq \text{dom}(\mathfrak{A})$. While an $\mathfrak{n}$-BFL $\mathfrak{A}$ can be "found" in another $\mathfrak{n}$-BFL $\mathfrak{B}$ if there is an embedding from $\mathfrak{A}$ to $\mathfrak{B}$, a pattern adds the constraint that the embedding must add no new Link (or Link*) to elements of K. Without patterns, negative constraints such as $\forall y.\ \neg\text{Link}(c, y)$ could not be encoded. Embeddings that agree with that constraint are matches.

There are two ways to generate (finite or infinite) sequences of sates. One is operational, based on sets of rules (a rule is just a pattern with more information): given a partial match from a rule to a state, a next state is constructed. The other is denotational, based on sets of transitions: any sequence from the set is correct if the post- and preconditions of every successive transition are equal. All sequences are defined relative to some starting state.

The main contribution of this thesis is to provide conditions where those two notions coincide and where the operational part can be finitely described:

For every set of transitions, there is an associated set of canonical patterns (and corresponding rules). If a formula ϕ is provable in the deductive system $\vdash_\rho$ then, for some starting states, the denotational runs of $[\![\phi]\!]$ are exactly the executions induced by the canonical rules associated to $[\![\phi]\!]$.

A rule-based rewriting engine starts from a file, which is a finite object. In general, the canonical rules associated to $[\![\phi]\!]$ may be an infinite set. If a formula ϕ is provable in another deductive system, $\vdash_\triangleleft$, then that set of rules is finite. The execution of those rules may not match the denotational runs; for that, one needs $\vdash_\rho$.

This is the desired correspondence between a denotational notion of runs and an operation notion of execution, given by a finite object. That finite object is essentially a subset of the models of a formula ϕ within a bounded size, so it can effectively be constructed.

An additional property of $\vdash_\triangleleft$ is that, modulo the theory of $\mathfrak{n}$-BFLs, the formulas it proves are all in the $\exists^*\forall^*$ fragment and also in the $\forall^*\exists^*$ fragment. But they are not negation-closed. Given the earlier decidability result, this opens the door to querying. The elimination of $\downarrow$ happens thanks to its 2nd-order, syntactic definition.

A special-purpose syntax for programs is here to make biological modelling easier. Every program P corresponds to a formula ϕ_P, and ϕ_P is always provable both in $\vdash_\triangleleft$ and $\vdash_\rho$, up to equivalence.

8.3. Future work

Application to a biological use case   In the short term, we plan on applying the *program* syntax to a relevant biological use case. However, we believe that most of the practical usefulness of the logic developed here will after the development of a higher-level language that is more comfortable to use and compiles to the iota logic.

Complexity of the fragments   There are lower bounds to establish on satisfiability in $\exists^*\forall^*$ modulo $\mathcal{T}_\mathfrak{n}$ as well as on $\vdash_\rho$ and $\vdash_\triangleleft$-provable formulas.

Stochastic rates   To accurately represent biological observations, attaching reaction rates to each rule is mandatory. Given a state and a set of rules, every "match" (in the sense of some graph morphism) from a rule to the state is counted. The probability of rewriting with one match is weighted by the rate attached to the rule which is source of the match.

Rewriting in our setting is nondeterministic. We intend to introduce rates at the level of programs: a program would be of the form

$$\theta_1 \texttt{ at rate}_1 \texttt{ ; } \ldots \texttt{ ; } \theta_k \texttt{ at rate}_k \texttt{ (\ldots)}$$

where each `rate`$_i$ is e.g. a floating-point number. The main difficulty is finding a representation of the canonical patterns such that the number of transitions at the denotational level is equal to the number of match at the pattern level.

Better $\vdash_\lhd$ Currently, a formula such as

$$Symlink = \forall xy. \ (\mathsf{Link}(x,y) \to \mathsf{Link}(y,x)) \land (\mathsf{Link}^\star(x,y) \to \mathsf{Link}^\star(y,x))$$

is not provable. In general, if there is a guard $\alpha(x,y) \to \ldots$, the "..." may only contain dynamic information about x or y, but not both. We think the requirement can be loosened: binary dynamic atoms that mention *both* x and y can be allowed (but dynamic atoms that only mention x cannot be allowed together with dynamic atoms that only mention y). This should not require changing the definition of preservation; we think it should be possible to define a class of *nice* clauses (as we did with good and great clauses in chapter 3), and add a rule of the form

$$\mathcal{C} \text{ nice} \quad \frac{}{\emptyset, 0 \vdash_\lhd \mathcal{C}^\circ \land \mathcal{C}^\star}$$

The rule would replace both INVARIANT$_\lhd$ and GUARD$_\lhd$. INVARIANT$_\lhd$ would be admissible, using a technique similar to the proof of lemma 5.2.2, and exactly the same formulas would be provable, albeit as tightly (context-wise) as with GUARD$_\lhd$. It seems that the following definition of a *nice clause* $\mathcal{C}$ should work: there is some x mentioned in every dynamic literal, and every $y \neq x$ mentioned in a dynamic literal is mentioned in a negative dynamic literal.

Rule merging Consider F_0 and F_1 two formulas on Θ. Let $\mathsf{Shift}(F_1)$ be F_1 where every starred symbol in F_1 is replaced with a *doubly starred symbol*, and every unstarred symbol is replaced with a starred symbol. For instance, $A^\star$ becomes $A^{\star\star}$ and A becomes $A^\star$. We would be using a signature $\Theta' = (\mathbf{Stat}, \mathbf{Dyn}, \mathbf{Dyn}^\star, \mathbf{Dyn}^{\star\star})$ with static symbols, normal dynamic symbols, starred symbols and doubly starred symbols. The formula $F_0 \land \mathsf{Shift}(F_1)$ would be a Θ'-formula.

We could of course extend this to k-state transitions, which opens questions about minimisation (but introduces issues related to the frame problem [Reiter, 1991] that we have mostly avoided in this thesis). But the 3-place case is particular since we don't need more to ask this question:

Is there a Θ-formula F such that $\llbracket F \rrbracket$ is the projection of $\mathbf{Stat}, \mathbf{Dyn}, \mathbf{Dyn}^{\star\star}$ from the models of $\llbracket F_0 \land \mathsf{Shift}(F_1) \rrbracket$?

If yes, we could iteratively apply this process and have formulas that characterize the n-step accessible states from any starting state. More prosaically, it means we could construct new rules by combining existing rules in a new dimension (by gluing them end-to-end with the shift-conjunction method given above for F_0 and F_1).

Suppose that, for every interpretation of $\mathbf{Stat}, \mathbf{Dyn}, \mathbf{Dyn}^{\star\star}$, $\llbracket F_0 \land \mathsf{Shift}(F_1) \rrbracket$ only allows a single middle state (i.e. the semantics of $\mathbf{Dyn}^\star$ are uniquely defined by the semantics of $\mathbf{Stat}, \mathbf{Dyn}, \mathbf{Dyn}^{\star\star}$). In that case Beth's definability theorem answers in the positive: there is a Θ-formula as we wish. We wonder if there are extensions of Beth's theorem that could hold when the middle state is not uniquely defined but may only be one among a finite number of states, with that bound a function of F_0 and F_1. Under $\downarrow$ this constraint is realistic.

Unguarded $\exists$ **in** $\vdash_\lhd$ As mentioned page 69, it should be possible to allow unguarded existential quantification in $\vdash_\lhd$ by changing the definition of preservation. Instead of drawing boundaries around the image of an interpretation, the definition would require the existence of a bounded-size protected set. We avoided this approach because it does not immediately yield *where* these protected elements are, and the proofs from chapter 4 on eliminating $\downarrow$ from $\vdash_\lhd$-provable formulas relies on knowing the location of these protected elements.

Removing Support   Even though support is used heavily throughout, we think it may be possible to not rely on it so much and simplify proofs in the process. First, state could be on the signature $\Theta^\circ \setminus \{\mathsf{Has}\}$, and the state joining operation $\langle \cdot, \cdot \rangle$ could add populate the interpretations of Has and Has* depending on whether an element is in the domain of the left state and/or the right state. This way, all transitions would be immediately decomposable, and would also be immediately strongly supported as well as dynamically supported.

The harder part is that support seems to be necessary in the proof of $\downarrow$-elimination when the domain quantification is removed, but modulo boundedness, it may be possible to eliminate domain quantification by quantifying on a finite number of "possibly removed" elements (recall that in circumscription domain quantification only needs to consider subdomains of the "current" one, so if there are a finite number of candidates for removal, there are a finite number of subdomains to consider).

Effective rule synthesis   We have so far avoided the topic of complexity for rule synthesis. For a formula ϕ, can we write down the rules induced by $\mathcal{P}(\phi, \|\phi\|)$ in reasonable time? Those patterns are of size $\mathcal{O}(|\mathbf{Stat}_{\mathsf{fun}}|^{2\|\phi\|n^2})$ (see lemma 4.3.6 page 67), i.e. exponential in the quantifier depth of ϕ and in n^2. For a given $s \geqslant 0$, the number patterns of size s (up to iso) is bounded by $2^{su} \times (s+1)^{2s}$, where u is the number of unary predicates in Θ. The first term is for the names and dynamic properties of elements. The second term is for each functional link relations (Link and Link*). So we are a priori in double exponential time.

First, an observation: formulas induced by programs are of the form $\bigvee \downarrow\psi_i$, and each pattern set can be synthesised in parallel. It seems reasonable to hope for a large number of small ψ_is.

Second, we have good hope of finding conditions where there exists a pattern set P such that $\mathsf{Insts}(P) = \mathsf{Insts}(\mathcal{P}(\phi, \|\phi\|))$ (same instances implies same runs), and such that each $p \in P$ is of size *polynomial* in the number of quantifiers of ϕ: quantifier depth bounds the maximum *distance* between "important" pattern elements, but the number of quantifiers bounds the *number* of "important" pattern trees. Trees have size exponential in n, however.

The size of P itself, however, would be exponential in the number of quantifiers of ϕ and in n.

Full-blown specification language   First, an extension of programs P: we would like to extend the syntax with

$$\mathtt{react} \ \ +J_1 \to J'_1, \ldots, +J_m \to J'_m$$

where each J_i has the same form as the I_j's in the definition of P, J'_i is quantifier-free, $\langle J'_i \rangle \subseteq \langle J_i \rangle$ and $+J_i \to J'_i$ is compiled to

$$\forall \neg J_i \to J_i^\star \to J'_i$$

There may be additional constraint on the polarity of literals and whether J'_i is purely post. In plain english, this would add universal invariants of the form "if J_i becomes true, ensure J'_i". The idea is that with formula of the form $+A \to B$, the variables mentioned in B are already

known to be active. So if their tree is cleared, $+A$ becomes false, validating the implication. We are confident that this type of formula is provable in $\vdash_{\rho}$. $\vdash_{\triangleleft}$ may need to be extended.

Second, a longer-term goal is to develop a full-fledged programming language for molecular biologists. Any valid program would compile to a formula provable both in $\vdash_{\rho}$ and $\vdash_{\triangleleft}$.

A. Additional proofs

To make reading easier, we reproduce the definition of $\mathcal{T}_n$ given page 61:

$$\mathcal{T}_n := Symlink \wedge Funlink \wedge \mathcal{T}'_n$$

where

$$\mathcal{T}'_n := \forall x. PSpec(x) \wedge H^n(x)$$

$$PSpec(x) := \left(\bigwedge_{f \in \mathbf{Stat}_{fun}} f(x) \neq x \rightarrow parent(f(x)) = x \right)$$

$$\wedge \left(parent(x) \neq x \rightarrow \bigvee_{f \in \mathbf{Stat}_{fun}} f(parent(x)) = x \right)$$

$$H^i(x) := \bigwedge_{f \in \mathbf{Stat}_{fun}} f(x) \neq x \rightarrow H^{i-1}(f(x)) \qquad \text{(for } 0 \leqslant i < n)$$

$$H^{-1}(x) := \bot$$

$$FunLink' := \forall x, y, z. \, Link(x,y) \rightarrow Link(x,z) \rightarrow y = z \qquad FunLink := FunLink \wedge FunLink^\star$$

$$Symlink' := \forall x, y. \, Link(x,y) \rightarrow Link(y,x) \qquad\qquad Symlink := Symlink \wedge Symlink^\star$$

Lemma 4.1.1. If $\mathfrak{A} : \Theta$, $\mathfrak{A} : \Theta_n$ iff $\mathfrak{A} \vDash \mathcal{T}_n$.

Proof. Trivially, $\mathfrak{A} \vDash Symlink \wedge FunLink$ iff $[\![Link]\!]_{\mathfrak{A}}$ and $[\![Link^\star]\!]_{\mathfrak{A}}$ are symmetric and functional. Let $p := [\![parent]\!]_{\mathfrak{A}}$. Now:

$\mathfrak{A} : \Theta_n$ **implies** $\mathfrak{A} \vDash \mathcal{T}_n$ For any $a \in dom(\mathfrak{A})$, we first verify the left conjunct of $PSpec(a)$: if $f(a) = b$ for some $b \neq a$ and $f \in [\![\mathbf{Stat}_{fun} \setminus \{parent\}]\!]_{\mathfrak{A}}$, by definition of an n-BFL we have $p(f(a)) = a$.

For the right conjunct, the definition of an n-BFL also gives $f(p(a)) = a$ for some $f \in [\![\mathbf{Stat}_{fun} \setminus \{parent\}]\!]_{\mathfrak{A}}$.

Now we show that for all $0 \leqslant k \leqslant n$, if the longest path away from the root starting from a in $G_{\mathfrak{A}}$ is of length at most k, then a satisfies H^n. By induction on n, the base case is trivial: we have $k = 0$ and a satisfies $H^0(x) = \bigwedge_{f \in f} f(x) = x$. Otherwise take any child $f(a)$ of a (for some $f \in f$), the longest path from $f(a)$ is of length at most $k-1$, so by induction we have $f(a)$ satisfies H^{n-1}. Since this is for any f, a satisfies H^n.

$\mathfrak{A} \vDash \mathcal{T}_n$ **implies** $\mathfrak{A} : \Theta_n$ Since any element satisfies H^n, there is no forward path of length $> n$. The last thing to check is that we have a tree, not a DAG. Take $a, b \in dom(\mathfrak{A})$, $f, g \in [\![\mathbf{Stat}_{fun} \setminus \{parent\}]\!]_{\mathfrak{A}}$ such that $a \neq f(a) = g(b) \neq b$. Since $PSpec(a)$ and $PSpec(b)$ hold in $\mathfrak{A}$, the first conjuncts gives us respectively $p(f(a)) = a$ and $p(g(b)) = b$. We assumed $f(a) = g(b)$, so

$a = b$. This also shows that the graph of p contains the parent relation of the loop-free union of the graphs of $[\![\mathbf{Stat}_{\mathrm{fun}} \setminus \{\mathrm{parent}\}]\!]_{\mathfrak{A}}$. To see that the graph of p is contained in it, consider the second conjunct of *PSpec*: whenever p is not the identity on some a, there must be some f such that $f(p(a)) = a$. $\square$

Lemma A.0.1. $\cdot \vDash_{\bar{\rho}} \mathsf{Has}_0 \wedge \mathsf{Has}_0^\star$

Proof. Immediate by lemma 5.2.2. $\square$

A.1. Building context-free and retractable formulas

Proof of lemma 3.3.1:

Lemma 3.3.1. If ϕ is (I, d)-context-free then ϕ is (I', d')-context-free for any $d' \geqslant d$, $I' \vDash I$.

Proof. Take any $d' \geqslant d$, $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \to \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d')) \to \mathfrak{B}$, $\mathfrak{A}, \mu \vDash \phi$ and $\mathfrak{B} \vDash \forall (I' \vee m)$. Since $\mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d') \subseteq \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)$, by lemma 2.4.4, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}$; moreover $\mathfrak{B} \vDash \forall (I' \vee m)$ trivially implies $\mathfrak{B} \vDash \forall (I \vee m)$ so by assumption on ϕ we get $\mathfrak{B}, m \circ \mu \vDash \phi$. $\square$

Proof of lemma 3.3.2:

Lemma 3.3.2. If ϕ is (I, d)-retractable then ϕ is (I', d')-retractable for any $d' \geqslant d$, $I' \vDash I$.

Proof. Take any $d' \geqslant d$, $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \to \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d')) \to \mathfrak{B}$, $\mathfrak{B}, m \circ \mu \vDash \phi$ and $\mathfrak{B} \vDash \forall (I \vee m)$. Since $\mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d') \subseteq \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)$, by lemma 2.4.4, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}$; moreover $\mathfrak{B} \vDash \forall (I' \vee m)$ trivially implies $\mathfrak{B} \vDash \forall (I \vee m)$ so by assumption on ϕ we get $\mathfrak{A}, \mu \vDash \phi$. $\square$

Proof of lemma 3.3.3

Lemma 3.3.3. A quantifier-free formula ϕ is $(\top, 0)$-context-free.

Proof. Let $\mathfrak{A}, \mu \vDash \phi$, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(0)) \to \mathfrak{B}$. Since ϕ is quantifier-free and m is an embedding, $\mathfrak{B}, m \circ \mu \vDash \phi$. $\square$

Proof of lemma 3.3.4

Lemma 3.3.4. A quantifier-free formula ϕ is $(\top, 0)$-retractable.

Proof. Take $d = 0$ as retraction radius. For $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \to \mathrm{dom}(\mathfrak{A})$, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d)) \to \mathfrak{B}$ with $\mathfrak{B}, m \circ \mu \vDash \phi$, m is an embedding so $\mathfrak{A}, \mu \vDash \phi$. $\square$

Proof of lemma 3.3.5:

Lemma 3.3.5. If ϕ_1 is (I, d_1)-context-free (resp. retractable) and ϕ_2 is (I, d_2)-context-free (resp. retractable) then $\phi_1 \wedge \phi_2$ is $(I, \max(d_1, d_2))$-context-free (resp. retractable).

Proof. Take $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi_1 \wedge \phi_2 \rangle \to \mathrm{dom}(\mathfrak{A})$, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(\max(d_1, d_2)) \to \mathfrak{B}$ such that $\mathfrak{A}, \mu \vDash \phi_1 \wedge \phi_2$ (resp. $\mathfrak{B}, m \circ \mu \vDash \phi_1 \wedge \phi_2$) and $\mathfrak{B} \vDash \forall (I \vee m)$. We want $\mathfrak{B}, m \circ \mu \vDash \phi_1 \wedge \phi_2$ (resp. $\mathfrak{A}, \mu \vDash \phi_1 \wedge \phi_2$).
By lemma 2.4.4, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}(\mu(\langle \phi_1 \rangle) \cup [\![\Delta]\!]_{\mathfrak{A}}, d_1) \to \mathfrak{B}$. So by definition of context-free (resp. retractable), $\mathfrak{B}, m \circ \mu \vDash \phi_1$ (resp. $\mathfrak{A}, \mu \vDash \phi_1$). Similarly, $\mathfrak{B}, m \circ \mu \vDash \phi_2$ (resp. $\mathfrak{A}, \mu \vDash \phi_2$). $\square$

Proof of lemma 3.3.6:

Lemma 3.3.6. If ϕ_1 is (I, d_1)-context-free (resp. retractable) and ϕ_2 is (I, d_2)-context-free (resp. retractable) then $\phi_1 \vee \phi_2$ is $(I, \max(d_1, d_2))$-context-free (resp. retractable).

Proof. Take $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi_1 \vee \phi_2 \rangle \rightarrow \mathrm{dom}(\mathfrak{A}), m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(\max(d_1, d_2))) \rightarrow \mathfrak{B}$ such that $\mathfrak{A}, \mu \vDash \phi_1 \vee \phi_2$ (resp. $\mathfrak{B}, m \circ \mu \vDash \phi_1 \vee \phi_2$) and $\mathfrak{B} \vDash \forall(I \vee m)$. We want $\mathfrak{B}, m \circ \mu \vDash \phi_1 \vee \phi_2$ (resp. $\mathfrak{A}, \mu \vDash \phi_1 \vee \phi_2$).

For $i \in \{1, 2\}$, suppose $\mathfrak{A}, \mu \vDash \phi_i$ (resp. $\mathfrak{B}, m \circ \mu \vDash \phi_i$). By lemma 2.4.4, $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}(\mu(\langle \phi_i \rangle) \cup \llbracket \Delta \rrbracket_{\mathfrak{A}}, d_i) \rightarrow \mathfrak{B}$. So by definition of context-free (resp. retractable), $\mathfrak{B}, m \circ \mu \vDash \phi_i$ (resp. $\mathfrak{A}, \mu \vDash \phi_i$). $\square$

Proof of lemma 3.3.7:

Lemma 3.3.7. If ϕ is (I, d)-context-free and $\alpha(x, y)$ is a guard, then $\forall x.\ \alpha(x, y) \rightarrow \phi$ is $(I, d+1)$-context-free.

Proof. Let $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \setminus \{x\} \rightarrow \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d+1)) \rightarrow \mathfrak{B}$ with $\mathfrak{A}, \mu \vDash \forall x.\ \alpha(x, y) \rightarrow \phi$ and $\mathfrak{B} \vDash \forall(I \vee m)$. We want $\mathfrak{B}, m \circ \mu \vDash \forall x.\ \alpha(x, y) \rightarrow \phi$.

Take any $a \in \mathrm{dom}(\mathfrak{B})$ such that $\mathfrak{B}, m \circ \mu, x \mapsto a \vDash \alpha(x, y)$. Let $b := \mu(y)$. By lemma 3.2.1, $\delta_{\mathfrak{B}}(a, m(b)) \leqslant 1$. Since $b \in \mathrm{Im}(\mu)$, by definition of a match, $a \in m(\mathfrak{A})$. Since m is an embedding, $\mathfrak{A}, x \mapsto m^{-1}(a), y \mapsto b \vDash \alpha(x, y)$. Since $\mathfrak{A}, \mu \vDash \forall x.\ \alpha(x, y) \rightarrow \phi$, we get $\mathfrak{A}, \mu, x \mapsto m^{-1}(a) \vDash \phi$. Moreover, by lemma 3.2.1, $\delta_{\mathfrak{A}}(m(a), b) \leqslant 1$, so $\mathcal{B}_{\mathfrak{A}}^{\Delta, \nu}(d) \subseteq \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d+1)$ where $\nu := \mu, x \mapsto m^{-1}(a)$. We can now apply lemma 2.4.4 to m and get $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \nu}(d)) \rightarrow \mathfrak{B}$. Since ϕ is (I, d)-context-free, $\mathfrak{B}, \nu \vDash \phi$. Since we took any $a \in \mathrm{dom}(\mathfrak{B})$, we have $\mathfrak{B}, \mu \vDash \forall x.\ \alpha(x, y) \rightarrow \phi$. $\square$

Proof of lemma 3.3.8:

Lemma 3.3.8. If ϕ is (I, d)-retractable then for any x, $\forall x.\ \phi$ is (I, d)-retractable.

Proof. Let d be a retraction radius for ϕ, I. Let $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \setminus \{x\} \rightarrow \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d)) \rightarrow \mathfrak{B}$ such that $\mathfrak{B}, m \circ \mu \vDash \forall x.\ \phi$ and $\mathfrak{B} \vDash \forall(I \vee m)$. We want $\mathfrak{A}, \mu \vDash \forall x.\ \phi$.

Take any $a \in \mathrm{dom}(\mathfrak{A})$. Since $\mathfrak{B} \vDash \forall x.\ \phi$, $\mathfrak{B}, m \circ \mu, x \mapsto m(a) \vDash \phi$. Since ϕ is (I, d)-retractable and $m^{-1}(m \circ \mu, x \mapsto m(a)) = \mu, x \mapsto a$, we get $\mathfrak{A}, \mu, x \mapsto a \vDash \phi$. $\square$

Proof of lemma 3.3.9:

Lemma 3.3.9. If ϕ is (I, d)-context-free then for any x, $\exists x.\ \phi$ is (I, d)-retractable.

Proof. Let $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \setminus \{x\} \rightarrow \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d)) \rightarrow \mathfrak{B}$ such that $\mathfrak{A}, \mu \vDash \exists x.\ \phi$ and $\mathfrak{B} \vDash \forall(I \vee m)$. We want $\mathfrak{B}, m \circ \mu \vDash \exists x.\ \phi$.

There is $a \in \mathrm{dom}(\mathfrak{A})$ such that $\mathfrak{A}, \mu, x \mapsto a \vDash \phi$. Since ϕ is (I, d)-context-free, $\mathfrak{B}, m \circ \mu, x \mapsto m(a) \vDash \phi$. $\square$

Proof of lemma 3.3.10:

Lemma 3.3.10. If ϕ is (I, d)-retractable and $\alpha(x, y)$ is a guard then $\exists x.\ \alpha(x, y) \wedge \phi$ is $(I, d+1)$-retractable.

Proof. Let d be a retraction radius for ϕ, I, let $\mathfrak{A}, \mathfrak{B}, \mu : \langle \phi \rangle \setminus \{x\} \rightarrow \mathrm{dom}(\mathfrak{A})$ and $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta, \mu}(d+1)) \rightarrow \mathfrak{B}$ with $\mathfrak{B}, m \circ \mu \vDash \exists x.\ \alpha(x, y) \wedge \phi$ and $\mathfrak{B} \vDash \forall(I \vee m)$. We want $\mathfrak{A}, \mu \vDash \exists x.\ \alpha(x, y) \wedge \phi$.

There is some $a \in \mathrm{dom}(\mathfrak{B})$ such that $\mathfrak{B}, m \circ \mu, x \mapsto a \vDash \alpha(x, y) \wedge \phi$. Let $b := \mu(y)$. By lemma 3.2.1, $\delta_{\mathfrak{B}}(a, m(b)) \leqslant 1$. Since $b \in \mathrm{Im}(\mu)$, by definition of a match $a \in m(\mathfrak{A})$. Since m is an

embedding, $\mathfrak{A}, x \mapsto m^{-1}(a), y \mapsto b \vDash \alpha(x,y)$, so again by lemma 3.2.1, $\delta_{\mathfrak{A}}(m^{-1}(a),b) \leqslant 1$. So $\mathcal{B}_{\mathfrak{A}}^{\Delta,\nu}(d) \subseteq \mathcal{B}_{\mathfrak{A}}^{\Delta,\mu}(d+1)$ where $\nu := \mu, x \mapsto m^{-1}(a)$. We can now apply lemma 2.4.4 to m and get $m : (\mathfrak{A}, \mathcal{B}_{\mathfrak{A}}^{\Delta,\nu}(d)) \to \mathfrak{B}$. Since φ is (I,d)-retractable, $\mathfrak{A}, \nu \vDash \varphi$. So $\mathfrak{A}, \mu \vDash \exists x.\ \alpha(x,y) \wedge \varphi$. $\qquad\square$

Index

Index

Bibliography

[Bárány et al., 2011] Bárány, V., Ten Cate, B., and Segoufin, L. (2011). Guarded negation. In *International Colloquium on Automata, Languages, and Programming*, pages 356–367. Springer.

[Bard, 2004] Bard, J. B. L. (2004). Ontologies in biology: design, application and future challenges. *Nature Reviews Genetics*, 5:213–222.

[Bartholomew et al., 2011] Bartholomew, M., Lee, J., and Meng, Y. (2011). First-order extension of the FLP stable model semantics via modified circumscription. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22, page 724.

[Basso-Blandin et al., 2015] Basso-Blandin, A., Fontana, W., and Harmer, R. (2015). A knowledge representation meta-model for rule-based modelling of signalling networks. In *Proceedings of the Eleventh International Workshop on Developments in Computational Models, DCM 2015, Cali, Colombia, October 28, 2015.*, pages 47–59.

[Boniolo et al., 2010] Boniolo, G., D'Agostino, M., and Di Fiore, P. P. (2010). Zsyntax: A formal language for molecular biology with projected applications in text mining and biological prediction. *PLOS ONE*, 5(3):1–12.

[Börger et al., 1997] Börger, E., Grädel, E., and Gurevich, Y. (1997). *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer.

[Boutillier et al., 2017] Boutillier, P., Ehrhard, T., and Krivine, J. (2017). Incremental update for graph rewriting. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, pages 201–228.

[Boutillier et al., 2018] Boutillier, P., Maasha, M., Li, X., Medina-Abarca, H. F., Krivine, J., Feret, J., Cristescu, I., Forbes, A. G., and Fontana, W. (2018). The Kappa platform for rule-based modeling. *Bioinformatics*, 34(13):583–592.

[Chaudhri et al., 2014] Chaudhri, V. K., Elenius, D., Goldenkranz, A., Gong, A., Martone, M. E., Webb, W., and Yorke-Smith, N. (2014). Comparative analysis of knowledge representation and reasoning requirements across a range of life sciences textbooks. *Journal of Biomedical Semantics*, 5:51.

[Chylek et al., 2014] Chylek, L., Harris, L., Tung, C.-S., Faeder, J., Lopez, C., and Hlavacek, W. (2014). Rule-based modeling (...). *Wiley interdisciplinary reviews Systems biology and medicine*, 6(1):13–36.

[Cohen, 2015] Cohen, P. R. (2015). DARPA's big mechanism program. *Physical Biology*, 12(4):045008.

[Conradie, 2006] Conradie, W. (2006). On the strength and scope of DLS. *Journal of Applied Non-Classical Logics*, 16(3-4):279–296.

Bibliography

[Danos et al., 2012] Danos, V., Feret, J., Fontana, W., Harmer, R., Hayman, J., Krivine, J., Thompson-Walsh, C., and Winskel, G. (2012). Graphs, rewriting and pathway reconstruction for rule-based models. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2012)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[Danos et al., 2007] Danos, V., Feret, J., Fontanta, W., Harmer, R., and Krivine, J. (2007). Rule based modeling of biological signaling. In Caires, L. and Vasconcelos, V. T., editors, *Proceedings of CONCUR 2007*, volume 4703 of *LNCS*, pages 17–41. Springer.

[De Kleer and Konolige, 1989] De Kleer, J. and Konolige, K. (1989). Eliminating the fixed predicates from a circumscription. *Artificial Intelligence*, 39(3):391–398.

[Demir et al., 2010] Demir, E., P Cary, M., Paley, S., Fukuda, K., Lemer, C., and Vastrik, I. (2010). The BioPAX community standard for pathway data sharing. *Nature Biotechnology*, 28.

[Despeyroux, 2016] Despeyroux, J. (2016). (Mathematical) logic for systems biology. In *CMSB*, pages 3–12. Springer.

[Doherty et al., 2004] Doherty, P., Kertes, S., Magnusson, M., and Szalas, A. (2004). Towards a logical analysis of biochemical pathways. In *European Workshop on Logics in Artificial Intelligence*, pages 667–679. Springer.

[Doherty et al., 1997] Doherty, P., Łukaszewicz, W., and SzaŁas, A. (1997). Computing circumscription revisited: A reduction algorithm. *Journal of Automated Reasoning*, 18(3):297–336.

[Doherty et al., 1998] Doherty, P., Łukaszewicz, W., and Szałas, A. (1998). General domain circumscription and its effective reductions. *Fundamenta Informaticae*, 36(1):23–55.

[Eiter et al., 1993] Eiter, T., Gottlob, G., and Gurevich, Y. (1993). Curb your theory! a circumscriptive approach for inclusive interpretation of disjunctive information. In *Proceedings of IJCAI'93*, volume 13, pages 634–634.

[Eker et al., 2002] Eker, S., Knapp, M., Laderoute, K., Lincoln, P., Meseguer, J., and Sonmez, K. (2002). Pathway logic: symbolic analysis of biological signaling. *Biocomputing*, pages 400–412.

[Etherington, 1986] Etherington, D. W. (1986). *Reasoning with incomplete information : investigations of non-monotonic reasoning*. PhD thesis, University of British Columbia.

[Faeder et al., 2009] Faeder, J. R., Blinov, M. L., and Hlavacek, W. S. (2009). *Rule-Based Modeling of Biochemical Systems with BioNetGen*, pages 113–167. Humana Press, Totowa, NJ.

[Ferraris et al., 2007] Ferraris, P., Lee, J., and Lifschitz, V. (2007). A new perspective on stable models. In *IJCAI*, volume 7, pages 372–379.

[Ferraris et al., 2011] Ferraris, P., Lee, J., and Lifschitz, V. (2011). Stable models and circumscription. *Artificial Intelligence*, 175(1):236.

[Fisher and Henzinger, 2007] Fisher, J. and Henzinger, T. A. (2007). Executable cell biology. *Nature Biotechnology*, 25(11):1239–1249.

[Fontana and Buss, 1996] Fontana, W. and Buss, L. (1996). The barrier of objects: From dynamical systems to bounded organizations. Working papers, International Institute for Applied Systems Analysis.

[Grädel, 1999] Grädel, E. (1999). On the restraining power of guards. *J. Symb. Log.*, 64(4):1719–1742.

[Gulwani et al., 2017] Gulwani, S., Polozov, O., Singh, R., et al. (2017). Program synthesis. *Foundations and Trends® in Programming Languages*, 4(1-2):1–119.

[Harel et al., 2001] Harel, D., Kozen, D., and Tiuryn, J. (2001). Dynamic logic. In *Handbook of philosophical logic*, pages 99–217. Springer.

[Hodkinson, 2002] Hodkinson, I. (2002). Loosely guarded fragment of first-order logic has the finite model property. *Studia Logica*, 70(2):205–240.

[Hunter, 1990] Hunter, L. (1990). Artificial intelligence and molecular biology. *AI magazine*, 2:27–36.

[Husson and Krivine, 2019] Husson, A. and Krivine, J. (2019). A tractable logic for molecular biology. In Bogaerts, B., Erdem, E., Fodor, P., Formisano, A., Ianni, G., Inclezan, D., Vidal, G., Villanueva, A., Vos, M. D., and Yang, F., editors, Proceedings 35th International Conference on *Logic Programming* (Technical Communications) , Las Cruces, NM, USA, September 20-25, 2019, volume 306 of *Electronic Proceedings in Theoretical Computer Science*, pages 101–113. Open Publishing Association.

[Krivine, 2017] Krivine, J. (2017). Systems biology. *SIGLOG News*, 4(3):43–61.

[Laurent et al., 2018] Laurent, J., Medina-Abarca, H. F., Boutillier, P., Yang, J., and Fontana, W. (2018). A trace query language for rule-based models. In *International Conference on Computational Methods in Systems Biology*, pages 220–237. Springer.

[Lazebnik, 2002] Lazebnik, Y. (2002). Can a biologist fix a radio? – or, what I learned while studying apoptosis. *Cancer Cell*, 2(3):179–182.

[Libkin, 2013] Libkin, L. (2013). *Elements of finite model theory*. Springer Science & Business Media.

[Mamouras, 2016] Mamouras, K. (2016). Synthesis of strategies using the hoare logic of angelic and demonic nondeterminism. *arXiv preprint arXiv:1606.09110*.

[Mason et al., 1999] Mason, C. S., Springer, C. J., Cooper, R. G., Superti-Furga, G., Marshall, C. J., and Marais, R. (1999). Serine and tyrosine phosphorylations cooperate in raf-1, but not b-raf activation. *The EMBO journal*, 18(8):2137–2148.

[McCarthy, 1980] McCarthy, J. (1980). Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1):27 – 39. Special Issue on Non-Monotonic Logic.

[McCarthy, 1986] McCarthy, J. (1986). Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28(1):89 – 116.

[Miller and Sirán, 2005] Miller, M. and Sirán, J. (2005). Moore graphs and beyond: A survey of the degree/diameter problem. *The electronic journal of combinatorics*, 1000:DS14–Dec.

Bibliography

[Nonnengart et al., 1999] Nonnengart, A., Ohlbach, H. J., and Szałas, A. (1999). Elimination of predicate quantifiers. In *Logic, Language and Reasoning*, pages 149–171. Springer.

[Novère, 2015] Novère, N. L. (2015). Quantitative and logic modelling of molecular and gene networks. *Nature Reviews Genetics*, 16:146–158.

[Reiter, 1991] Reiter, R. (1991). The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. *Artificial intelligence and mathematical theory of computation: papers in honor of John McCarthy*, 27:359–380.

[Reiter, 2001] Reiter, R. (2001). *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT press.

[Stevens et al., 2000] Stevens, R., Goble, C. A., and Bechhofer, S. (2000). Ontology-based knowledge representation for bioinformatics. *Briefings in Bioinformatics*, 1(4):398–414.

[UniProt Consortium, 2019] UniProt Consortium, T. (2019). Uniprotkb - p09560 (raf1_xenla) [uniprot.org/uniprot/p09560].

[UniProt Consortium, 2018] UniProt Consortium, T. (2018). UniProt: the universal protein knowledgebase. *Nucleic Acids Research*, 46(5):2699–2699.

[Wan et al., 2014] Wan, H., Xiao, Z., Yuan, Z., Zhang, H., and Zhang, Y. (2014). Computing general first-order parallel and prioritized circumscription. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI'14, pages 1105–1111. AAAI Press.

[Yaman et al., 2018] Yaman, F., Adler, A., and Beal, J. (2018). AI challenges in synthetic biology engineering. In *The Thirtieth AAAI Conference on Innovative Applications of Artificial Intelligence (IAAI-18)*.

[Yang et al., 2011] Yang, Q., You, J.-H., and Feng, Z. (2011). Integrating rules and description logics by circumscription. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, AAAI'11. AAAI PRess.

[Zhang and Ying, 2010] Zhang, H. and Ying, M. (2010). Decidable fragments of first-order language under stable model semantics and circumscription. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI'10, pages 375–380. AAAI Press.