

Mathematical Structures in Computer Science

<http://journals.cambridge.org/MSC>

Additional services for *Mathematical Structures in Computer Science*:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)



Ticket Entailment is decidable

VINCENT PADOVANI

Mathematical Structures in Computer Science / Volume 23 / Issue 03 / June 2013, pp 568 - 607
DOI: 10.1017/S0960129512000412, Published online: 09 July 2012

Link to this article: http://journals.cambridge.org/abstract_S0960129512000412

How to cite this article:

VINCENT PADOVANI (2013). Ticket Entailment is decidable. *Mathematical Structures in Computer Science*, 23, pp 568-607 doi:10.1017/S0960129512000412

Request Permissions : [Click here](#)

Ticket Entailment is decidable

VINCENT PADOVANI

*Equipe Preuves, Programmes et Systèmes,
Université Paris VII – Denis Diderot,
Case 7014,
75205 PARIS Cedex 13, France
Email: padovani@pps.jussieu.fr*

Received 19 June 2010; revised 6 March 2012

We prove the decidability of the logic T_{\rightarrow} of Ticket Entailment. This issue was first raised by Anderson and Belnap within the framework of relevance logic, and is equivalent to the question of the decidability of type inhabitation in simply typed combinatory logic with the partial basis BB'IW. We solve the equivalent problem of type inhabitation for the restriction of simply typed lambda calculus to hereditarily right-maximal terms.

1. Introduction

The partial bases built using the atomic combinators B, B', C, I, K, W of combinatory logic are well known for being closely connected with propositional logic. The types of their combinators form the axioms of implicational logic systems that have been studied now for well over 70 years (Trigg *et al.* 1994). The partial basis BB'IW corresponds, through the types of its combinators, to the system T_{\rightarrow} of *Ticket Entailment*, which was introduced and motivated in Anderson and Belnap (1975) and Anderson *et al.* (1990). The system T_{\rightarrow} consists of *modus ponens* and four axiom schemes that range over the following types for each atomic combinator:

- B : $(\chi \rightarrow \psi) \rightarrow ((\phi \rightarrow \chi) \rightarrow (\phi \rightarrow \psi))$
- B' : $(\phi \rightarrow \chi) \rightarrow ((\chi \rightarrow \psi) \rightarrow (\phi \rightarrow \psi))$
- I : $\phi \rightarrow \phi$
- W : $(\phi \rightarrow (\phi \rightarrow \chi)) \rightarrow (\phi \rightarrow \chi)$

The type inhabitation problem for BB'IW is the problem of deciding for a given type whether there exists within this basis a combinator of this type. This problem is equivalent to the problem of deciding whether a given formula can be derived in T_{\rightarrow} .

Surprisingly, the question of the decidability of T_{\rightarrow} has remained unsolved since it was raised in Anderson and Belnap (1975), though the problem has been thoroughly explored within the framework of relevance logic with proofs of decidability and undecidability for several related systems. For instance, the system R_{\rightarrow} of *Relevant Implication* (which corresponds to the basis BCIW) and the system E_{\rightarrow} of *Entailment* (Anderson and Belnap 1975) are both decidable (Kripke 1959), whereas the extensions R , E , T of R_{\rightarrow} , E_{\rightarrow} , T_{\rightarrow} to a larger set of connectives (\rightarrow , \wedge , \vee) are undecidable (Urquhart 1984).

In 2004, a partial decidability result for the type inhabitation problem was proposed in Broda *et al.* (2004) for a restricted class of formulas – the class of 1-unary formulas in




Fig. 1. The principle of our proof of the decidability of type inhabitation for HRM-terms.

which every maximal negative subformula has arity at most 1. Broda, Dams, Finger and Silva e Silva’s approach is based on a translation of the problem into a type inhabitation problem for the *hereditary right-maximal* (HRM) terms of lambda calculus (Trigg *et al.* 1994; Bunder 1996; Broda *et al.* 2004). The closed HRM-terms form the closure under β -reduction of all translations of $BB'IW$ -terms, so the type inhabitation problem within the basis $BB'IW$ is equivalent to the type inhabitation problem for HRM-terms.

In this paper we use the same approach as used by Broda, Dams, Finger and Silva e Silva, and prove that the type inhabitation problem for HRM-terms is decidable, and thus conclude that the logic T_{\rightarrow} is decidable[†].

1.1. Organisation of the paper

In Section 2, we recall the definition of hereditarily right-maximal terms and the equivalence between the decidability of type inhabitation for $BB'IW$ and the decidability of type inhabitation for HRM-terms. The principle of our proof is shown in Figure 1.

In Sections 3 and 4, we provide for each formula ϕ a partial characterisation of the inhabitants of ϕ in normal form and of minimal size. We show that all of these inhabitants belong to two larger sets of terms, *viz.* the set of *compact* and *locally compact* inhabitants of ϕ .

In Section 5, we show how to associate with each locally compact inhabitant M of a formula ϕ , a labelled tree with the same tree structure as M – we call this tree the *shadow* of M . We then define for shadows the analogue of compactness for terms, and prove that the shadow of a compact term is itself compact.

[†] In the course of the publication of this article we heard of work in progress by Katalin Bimbò and Michael Dunn towards a solution that seems to be based on a different approach.

Finally, in Section 6, we prove that for each formula ϕ the set of all compact shadows of inhabitants of ϕ is a finite set (hence the set of compact inhabitants of ϕ is also a finite set), and then prove that this set is effectively computable from ϕ . The proof appeals to the Higman and Kruskal Theorems – more precisely, to Mellies’ Axiomatic Kruskal Theorem.

The decidability of the type inhabitation problem for HRM-terms and the decidability of T_{\rightarrow} , follow from a final key result: given an arbitrary formula ϕ , this formula is inhabited if and only if there exists a compact shadow with the same tree structure as an inhabitant of ϕ , and our key lemma proves that the existence of such a shadow is decidable.

1.2. Preliminaries

Section 2 assumes some familiarity with pure and simply typed lambda calculus and the usual notions of α -conversion, β -reduction and β -normal form (Barendregt 1984; Krivine 1993). These last three notions are not essential to our discussion, as we later focus exclusively on a particular set of simply typed terms in β -normal form. We shall now briefly recall the definitions and results used in Section 2.

The set of *terms of pure lambda calculus* (λ -terms) is defined inductively as follows:

- Every variable x is a λ -term.
- If M is a λ -term and x is a variable, then $(\lambda x M)$ is a λ -term.
- If M, N are λ -terms, then (MN) is a λ -term.

The terms given by the second and third rules are called *abstractions* and *applications*, respectively. The parentheses surrounding applications and abstractions are often omitted if no ambiguity arises. We write

$$\lambda x_1 \dots x_n. MN_1 \dots N_p$$

to abbreviate

$$(\lambda x_1 (\dots (\lambda x_n (((MN_1) \dots) N_p)) \dots)).$$

For instance, $\lambda xy.x(xy)z$ stands for $(\lambda x(\lambda y((x(xy))z)))$.

The *bound variables* of M are all x such that λx occurs in M . A variable x is *free* in M if and only if any of the following holds:

- $M = x$.
- $M = \lambda y.N$, $y \neq x$ and x is free in N .
- $M = NP$ and x is free in N or free in P .

A *closed* term is a term with no free variables. The *raw substitution of N for x in M* , written $M\langle x \leftarrow N \rangle$, is the term obtained by substituting N for every free occurrence of x in M (that is, every occurrence of x that is not in the scope of a λx). We require this substitution to avoid variable capture (for all y free in N , no free occurrence of x in M is in the scope of a λy):

- If $y = x$, then $y\langle x \leftarrow N \rangle$ is equal to N , otherwise it is equal to y .
- $(\lambda x.M)\langle x \leftarrow N \rangle = \lambda x.M$.

- If $y \neq x$ and y is free in N , then $(\lambda y.M)\langle x \leftarrow N \rangle$ is undefined.
- If $y \neq x$, y is not free in N and $M\langle x \leftarrow N \rangle = M'$, then $(\lambda y.M)\langle x \leftarrow N \rangle = \lambda y.M'$.
- If $M_1\langle x \leftarrow N \rangle = M'_1$ and $M_2\langle x \leftarrow N \rangle = M'_2$, then $(M_1M_2)\langle x \leftarrow N \rangle = (M'_1M'_2)$.

α -conversion is defined as the least binary relation \equiv_α such that:

- $x \equiv_\alpha x$.
- If $M \equiv_\alpha M'$, y is not free in M' and $M'\langle x \leftarrow y \rangle = M''$, then $(\lambda x.M) \equiv_\alpha (\lambda y.M'')$.
- If $M_1 \equiv_\alpha M'_1$ and $M_2 \equiv_\alpha M'_2$, then $(M_1M_2) \equiv_\alpha (M'_1M'_2)$.

For instance, $\lambda x.y \equiv_\alpha \lambda z.y \not\equiv_\alpha \lambda y.y$. It is a common practice to consider λ -terms up to α -conversion, but we will not do this in our exposition.

β -reduction is the least binary relation β satisfying:

- If $M \equiv_\alpha (\lambda x.N)P$ and $N\langle x \leftarrow P \rangle = N'$, then $M\beta N'$.
- If $M\beta M'$, then $(\lambda x.M)\beta(\lambda x.M')$, $(MN)\beta(M'N)$ and $(NM)\beta(NM')$.

In the first rule, x is not necessarily free in N , so we may have $N = N'$ – in particular, free variables may disappear in the process of reduction.

We write β^* for the reflexive and transitive closure of β . A term M is in β -normal form, or is β -normal, if there is no M' such that $M\beta M'$. A term M is normalising if there is a normal N , which is called the normal form of M , such that $M\beta^*N$. It is strongly normalising if there is no infinite sequence $M = M_0\beta M_1\beta M_2\dots$

It is well known that β -conversion enjoys the Church–Rosser property: if $M\beta^*N$ and $M\beta^*N'$, then there exist two α -convertible P, P' such that $N\beta^*P$ and $N'\beta^*P'$. As a consequence, if a term is normalising, its normal form is unique up to α -conversion.

The judgment ‘assuming x_1, \dots, x_n are of types ψ_1, \dots, ψ_n , the term M has type ϕ ’, written $\{x_1 : \psi_1, \dots, x_n : \psi_n\} \vdash M : \phi$, where $\psi_1, \dots, \psi_n, \phi$ are formulas of propositional calculus and x_1, \dots, x_n are distinct variables, is defined by:

- $\Gamma \vdash x : \psi$ for each $x : \psi \in \Gamma$.
- If $\Gamma \cup \{x : \phi\} \vdash M : \psi$, then $\Gamma \vdash \lambda x.M : \phi \rightarrow \psi$.
- If $\Gamma \vdash M : \phi \rightarrow \psi$ and $\Gamma \vdash N : \phi$, then $\Gamma \vdash (MN) : \psi$

The simply typable terms are all M for which there exist Γ, ϕ such that $\Gamma \vdash M : \phi$. Note that Γ contains all variables free in M . The following properties are well known:

- **Strong normalisation:** If $\Gamma \vdash M : \phi$, then M is strongly normalising.
- **Subject reduction:** If $\Gamma \vdash M : \phi$ and $M\beta N$, then $\Gamma \vdash N : \phi$.

2. From BB'IW to simply typed lambda calculus

The aim of this section is to provide a precise characterisation of simply typable terms that are typable with inhabited types in BB'IW so as to transform the problem of type inhabitation in BB'IW into a type inhabitation problem in lambda calculus. The types of atomic combinators in BB'IW are also types for their respective counterparts $\lambda f g x.f(gx)$, $\lambda f g x.g(fx)$, $\lambda x.x$, $\lambda h x.hxx$ in lambda calculus, so to each inhabited type ϕ in BB'IW, there corresponds at least one closed λ -term of type ϕ . Moreover, subject reduction and strong normalisation (see above) also ensure the existence of a closed normal λ -term of type ϕ .

What we lack is a criterion to distinguish amongst all typed normal forms those that are reducts of translations of combinators within $\mathbf{BB}'\mathbf{IW}$.

The material and results of this section are not new (Bunder 1996; Broda *et al.* 2004), and the contents of Sections 2.3 and 2.4, apart from Lemma 2.10, may be skipped entirely to go immediately to the study of stable parts and blueprints in Section 3.

The definition of hereditarily right-maximal terms is an adaptation of the definition given in Bunder (1996). The proof of Lemma 2.6 (subject reduction for HRM-terms) is similar to the proof of Property 2.4 on page 375 of Broda *et al.* (2004). The right-to-left implication of Lemma 2.10 can be deduced from Property 2.20 on page 390 of Broda *et al.* (2004), though our proof method seems to be simpler.

2.1. Lambda calculus

Let \mathcal{X} be a countably infinite set of variables $x, y, z \dots$ together with a one-to-one function \mathcal{O} from \mathcal{X} to \mathbb{N} . For all x, y in \mathcal{X} , we write $x < y$ when $\mathcal{O}(x) < \mathcal{O}(y)$. Throughout the rest of the paper, when we write *term*, we will always mean a term of lambda calculus built over these variables. For each term M , we write $\text{Free}(M)$ for the strictly increasing sequence of all free variables of M .

Terms are *not* identified modulo α -conversion – apart from in Section 2, all terms considered will be in normal form, and we will even use the Greek letters α and β with a new meaning at the beginning of Section 3. However, we do adopt the usual convention according to which two distinct λ 's may not bind the same variable in a term, and no variable can be simultaneously free and bound in the same term.

2.2. Hereditarily right-maximal terms

Definition 2.1. The set of *hereditarily right-maximal* (HRM) terms is inductively defined as follows:

- (1) Each variable x is HRM.
- (2) If M is HRM and x is the greatest free variable of M , then $\lambda x.M$ is HRM.
- (3) If M, N are HRM, and for each free variable x of M there exists a free variable y of N such that $x \leq y$, then (MN) is HRM.

The second rule ensures that all HRM-terms are λ_I -terms, that is, terms in which every subterm $\lambda x.M$ is such that x is free in M . As a consequence, the set of free variables of an HRM-term is preserved under β -reduction. As we shall see in Lemma 2.6, right-maximality can also be preserved at the cost of appropriate bound variable renamings.

In the third rule, if N is closed, then so is M . When M and N are non-closed terms, the greatest free variable of M is less than or equal to the greatest free variable of N . For instance, if $f < g < x$ and $h < x$, then $\lambda f g x.f(gx)$, $\lambda f g x.g(fx)$, $\lambda x.x$, $\lambda h x.hxx$ are HRM, whereas $\lambda y z.zy$ is not, no matter whether $y < z$ or $y > z$.

Definition 2.2. Let Ω be a function mapping each variable to a formula in such a way that $\Omega^{-1}(\phi)$ is an infinite set for each ϕ . We extend this function to the set of all strictly increasing finite sequences of variables by letting $\Omega(x_1, \dots, x_n) = (\Omega(x_1), \dots, \Omega(x_n))$.

Definition 2.3. The judgment $M : \phi$ (in words ‘ M has type ϕ with respect to Ω ’) is defined by:

- If $\Omega(x) = \phi$, then $x : \phi$.
- If $x : \chi$, $M : \psi$ and $\lambda x.M$ is HRM, then $\lambda x.M : \chi \rightarrow \psi$.
- If $M : \chi \rightarrow \psi$, $N : \chi$ and (MN) is HRM, then $(MN) : \psi$.

The function Ω will remain fixed throughout the exposition. Accordingly, the type of a term M with respect to Ω will be called *the* type of M without any further reference to the choice of Ω . Note that every typed term is HRM.

Definition 2.4. We write Λ_{NF} for the set of all typed terms in β -normal form. We say any closed term $M \in \Lambda_{\text{NF}}$ of type ϕ is an Λ_{NF} -inhabitant of ϕ .

The next lemma is the well-known subformula property of simply typed lambda calculus:

Lemma 2.5 (subformula property). Let M be a Λ_{NF} -inhabitant of ϕ . The types of the subterms of M are subformulas of ϕ .

2.3. Subject reduction of hereditarily right-maximal terms

Lemma 2.6. If there exists a closed $M : \phi$, then ϕ is Λ_{NF} -inhabited.

Proof.

- (1) The proof of the fact that for every variable y and for every $N : \phi$, there exists $N' \equiv_x N$ such that $N' : \phi$ and every bound variable of N' is strictly greater than y is left as an exercise.
- (2) We prove the following proposition by induction on P . Let P, Q be typed HRM-terms. Suppose:
 - x and Q have the same type,
 - If Q is closed and $x \in \text{Free}(P)$, then $x = \min(\text{Free}(P))$
 - If Q is not closed, then for all $z \in \text{Free}(P)$:
 - If $z < x$, then $z \leq \max(\text{Free}(Q))$.
 - If $x < z$, then $\max(\text{Free}(Q)) < z$.
 - if Q is not closed, then $\max(\text{Free}(Q)) < z$ for all bound variables z of P .

Then $R = P \langle x \leftarrow Q \rangle$ is defined, HRM and have the same type as P . We consider cases:

- P is a variable:
The proposition is clear in this case
- $P = \lambda z.P'$:
Then $\text{Free}(P') = \text{Free}(P) \cdot (z)$. By the induction hypothesis, $R' = P' \langle x \leftarrow Q \rangle$ is defined, HRM and have the same type as P' . The variable z is still the greatest free variable of R' and z is not free in Q , hence $R = \lambda z.R'$.

— $P = (P_1P_2)$:

By the induction hypothesis, $R_i = P_i\langle x \leftarrow Q \rangle$ is defined, HRM and have the same type as P_i for each $i \in \{1, 2\}$. We still need to check that $R = (R_1R_2)$ is HRM. Assume x is free in P and P_1 is not closed. There are three sub-cases:

– $\max(\text{Free}(P_1)) > x$:

Then

$$\max(\text{Free}(P_1)) = \max(\text{Free}(R_1)) \leq \max(\text{Free}(P_2)) = \max(\text{Free}(R_2)).$$

– $\max(\text{Free}(P_1)) < x$:

The term Q cannot be closed, and

$$\max(\text{Free}(P_1)) = \max(\text{Free}(R_1)) \leq \max(\text{Free}(Q)).$$

We have either

$$\max(\text{Free}(P_2)) = x$$

$$\max(\text{Free}(R_2)) = \max(\text{Free}(Q))$$

or

$$\max(\text{Free}(P_2)) > x$$

$$\max(\text{Free}(P_2)) = \max(\text{Free}(R_2)).$$

– $\max(\text{Free}(P_1)) = x$:

There are two subcases:

• $\max(\text{Free}(P_2)) > x$:

Then $\max(\text{Free}(P_2)) = \max(\text{Free}(R_2))$. If Q is closed, then $\text{Free}(P_1) = (x)$ and R_1 is closed. Otherwise,

$$\max(\text{Free}(R_1)) = \max(\text{Free}(Q)) \leq \max(\text{Free}(P_2)).$$

• $\max(\text{Free}(P_2)) = x$:

If Q is closed, then $\text{Free}(P_1) = \text{Free}(P_2) = (x)$ and R_1, R_2 are closed. Otherwise,

$$\max(\text{Free}(R_1)) = \max(\text{Free}(R_2)) = \max(\text{Free}(Q)).$$

(3) Assume $N : \phi$ and N is not in normal form. We will prove the existence of $N' : \phi$ such that $N\beta N'$ by induction on N .

If $N = \lambda x.P$ or $N = (N_1N_2)$ with N_1 or N_2 not in normal form, then the existence of N' follows from the induction hypothesis and the fact that β -reduction preserves the set of free variables of an HRM-term. Otherwise, $N = (\lambda x.P)Q$, where for each free variable z of $\lambda x.P$, we have $z < x$ and there exists a free variable y of Q such that $z < y$. By (1), there exists $P' \equiv_x P$ such that $P' : \phi$ and no bound variable of P' is less than or equal to a free variable of Q . The variable x is the greatest free variable of P' . By (2), the term $N' = P'\langle x \leftarrow Q \rangle$ is well defined, HRM and of the type ϕ . Moreover, $N\beta N'$.

- (4) We can now prove the lemma. The term M is a simply typable HRM-term. The strong normalisation property implies the existence of a normal form N of M . The term N is still a closed term. Finally, by (1), there exists $N' \equiv_{\alpha} N$ such that $N' : \phi$, that is, ϕ is Λ_{NF} -inhabited. \square

2.4. Equivalence between inhabitation in $\text{BB}'\text{IW}$ and Λ_{NF} -inhabitation

In the next three lemmas we write $\phi_1 \dots \phi_n \rightarrow \psi$ to mean the formula

$$(\phi_1 \rightarrow (\dots (\phi_n \rightarrow \psi) \dots))$$

if $n > 0$, and the formula ψ otherwise. We write $\vdash_{\text{BB}'\text{IW}} \phi$ for the judgment ‘there exists within the basis $\text{BB}'\text{IW}$ a combinator of type ϕ ’.

Lemma 2.7. If $\vdash_{\text{BB}'\text{IW}} \phi$, then ϕ is Λ_{NF} -inhabited.

Proof. If $f < g < x$ and $h < x$, then $\lambda x.x$, $\lambda f g x.f(gx)$, $\lambda f g x.g(fx)$ and $\lambda h x.hxx$ are HRM. For each type ϕ of an atomic combinator, the variables f, g, h, x can be chosen so that one of these terms has type ϕ . The set of all formulas ϕ for which there exists a closed M of type ϕ is closed under *modus ponens*. By Lemma 2.6, every such formula is Λ_{NF} -inhabited. \square

Lemma 2.8. If $\vdash_{\text{BB}'\text{IW}} \chi \rightarrow \psi$, then

$$\vdash_{\text{BB}'\text{IW}} (\phi_1 \dots \phi_n \rightarrow \chi) \rightarrow (\phi_1 \dots \phi_n \rightarrow \psi)$$

for all ϕ_1, \dots, ϕ_n .

Proof. The proof is by induction on n using left-applications of B . \square

Lemma 2.9. Suppose (i_1, \dots, i_n) , (j_1, \dots, j_m) , (k_1, \dots, k_p) are strictly increasing sequences of integers, $\{k_1, \dots, k_p\} = \{i_1, \dots, i_n, j_1, \dots, j_m\}$, $n = 0$ or $(n > 0, m > 0, i_n \leq j_m)$. If

$$(1) \vdash_{\text{BB}'\text{IW}} \omega_{i_1} \dots \omega_{i_n} \rightarrow (\chi \rightarrow \psi),$$

$$(2) \vdash_{\text{BB}'\text{IW}} \omega_{j_1} \dots \omega_{j_m} \rightarrow \chi,$$

then $\vdash_{\text{BB}'\text{IW}} \omega_{k_1} \dots \omega_{k_p} \rightarrow \psi$.

Proof. The proof is by induction on $n+m$. The proposition is obviously true if $n = m = 0$, so we assume $n+m > 0$. So $m > 0$. We consider cases:

— $n = 0$:

Then $(j_1, \dots, j_m) = (k_1, \dots, k_p)$. We have:

$$\vdash_{\text{BB}'\text{IW}} (\chi \rightarrow \psi) \rightarrow ((\omega_{j_m} \rightarrow \chi) \rightarrow (\omega_{j_m} \rightarrow \psi)) \quad (\text{i})$$

$$\vdash_{\text{BB}'\text{IW}} (\omega_{j_m} \rightarrow \chi) \rightarrow (\omega_{j_m} \rightarrow \psi), \quad (\text{ii})$$

where (i) is a type for B and (ii) follows from (i), (1) and *modus ponens*.

If $m = 1$, then $\vdash_{\text{BB}'\text{IW}} \omega_{j_1} \rightarrow \psi$ follows from (ii), (2) and *modus ponens*. Otherwise,

$$\vdash_{\text{BB}'\text{IW}} \omega_{j_1} \dots \omega_{j_m} \rightarrow \psi$$

follows from (ii), (2) and the induction hypothesis.

— $n > 0$:

We consider two sub-cases:

– $m > 1$ and $i_n \leq j_{m-1}$:

Hence:

$$\vdash_{\mathbf{BB}'\text{IW}} (\chi \rightarrow \psi) \rightarrow ((\omega_{j_m} \rightarrow \chi) \rightarrow (\omega_{j_m} \rightarrow \psi)) \quad (\text{iii})$$

$$\begin{aligned} \vdash_{\mathbf{BB}'\text{IW}} (\omega_{i_1} \dots \omega_{i_n} \rightarrow (\chi \rightarrow \psi)) \\ \rightarrow (\omega_{i_1} \dots \omega_{i_n} \rightarrow ((\omega_{j_m} \rightarrow \chi) \rightarrow (\omega_{j_m} \rightarrow \psi))) \end{aligned} \quad (\text{iv})$$

$$\vdash_{\mathbf{BB}'\text{IW}} \omega_{i_1} \dots \omega_{i_n} \rightarrow ((\omega_{j_m} \rightarrow \chi) \rightarrow (\omega_{j_m} \rightarrow \psi)), \quad (\text{v})$$

where (iii) is a type for \mathbf{B} and (iv) follows from (iii) and Lemma 2.8. (v) then follows from (iv), (1) and *modus ponens*.

We now have $k_p = j_m$ and

$$\{k_1, \dots, k_{p-1}\} = \{i_1, \dots, i_n, j_1, \dots, j_{m-1}\}.$$

Since $i_n \leq j_{m-1}$, we have

$$\vdash_{\mathbf{BB}'\text{IW}} \omega_{k_1} \dots \omega_{k_{p-1}} \rightarrow (\omega_{j_m} \rightarrow \psi)$$

by (v), (2) and the induction hypothesis.

– $m = 1$ or ($m > 1$ and $i_n > j_{m-1}$):

Hence:

$$\vdash_{\mathbf{BB}'\text{IW}} (\omega_{j_m} \rightarrow \chi) \rightarrow ((\chi \rightarrow \psi) \rightarrow (\omega_{j_m} \rightarrow \psi)) \quad (\text{vi})$$

$$\vdash_{\mathbf{BB}'\text{IW}} (\omega_{j_1} \dots \omega_{j_m} \rightarrow \chi) \rightarrow (\omega_{j_1} \dots \omega_{j_{m-1}} \rightarrow ((\chi \rightarrow \psi) \rightarrow (\omega_{j_m} \rightarrow \psi))) \quad (\text{vii})$$

$$\vdash_{\mathbf{BB}'\text{IW}} \omega_{j_1} \dots \omega_{j_{m-1}} \rightarrow ((\chi \rightarrow \psi) \rightarrow (\omega_{j_m} \rightarrow \psi)) \quad (\text{viii})$$

$$\vdash_{\mathbf{BB}'\text{IW}} \omega_{n_1} \dots \omega_{n_q} \rightarrow (\omega_{j_m} \rightarrow \psi), \quad (\text{ix})$$

where (vi) is a type for \mathbf{B}' and (vii) follows from (vi) and Lemma 2.8. (viii) then follows from (vii), (2) and *modus ponens*. Then, writing

$$\{n_1, \dots, n_q\} = \{j_1, \dots, j_{m-1}, i_1, \dots, i_n\},$$

(ix) follows from (viii), (1) and the induction hypothesis.

If $j_m > i_n$,

$$(n_1, \dots, n_q, j_m) = (k_1, \dots, k_p).$$

Otherwise,

$$j_m = i_n$$

$$n_q = i_n$$

$$(n_1, \dots, n_q) = (k_1, \dots, k_p)$$

and

$$\vdash_{\text{BB}'\text{IW}} \omega_{k_1} \dots \omega_{k_{p-1}} \rightarrow (\omega_{i_n} \rightarrow (\omega_{i_n} \rightarrow \psi)) \quad (\text{x})$$

$$\vdash_{\text{BB}'\text{IW}} (\omega_{i_n} \rightarrow (\omega_{i_n} \rightarrow \psi)) \rightarrow (\omega_{i_n} \rightarrow \psi) \quad (\text{xi})$$

$$\vdash_{\text{BB}'\text{IW}} (\omega_{k_1} \dots \omega_{k_{p-1}} \rightarrow (\omega_{i_n} \rightarrow (\omega_{i_n} \rightarrow \psi))) \rightarrow (\omega_{k_1} \dots \omega_{k_{p-1}} \rightarrow (\omega_{i_n} \rightarrow \psi)) \quad (\text{xii})$$

$$\vdash_{\text{BB}'\text{IW}} \omega_{k_1} \dots \omega_{k_{p-1}} \rightarrow (\omega_{i_n} \rightarrow \psi), \quad (\text{xiii})$$

where (x) is just (ix), and (xi) is a type for W. Then (xii) follows from (xi) and Lemma 2.8, and (xiii) follows from (x), (xii) and *modus ponens*.

Finally, (xiii) is just $\vdash_{\text{BB}'\text{IW}} \omega_{k_1} \dots \omega_{k_p} \rightarrow \psi$. \square

Lemma 2.10. For every formula ϕ , we have $\vdash_{\text{BB}'\text{IW}} \phi$ if and only if ϕ is Λ_{NF} -inhabited.

Proof. The left to right implication is Lemma 2.7. Using Lemma 2.9 when M is an application, an immediate induction on M shows that if $M : \psi$, $\text{Free}(M) = (x_1, \dots, x_n)$ and $x_1 : \chi_1, \dots, x_n : \chi_n$, then $\vdash_{\text{BB}'\text{IW}} \chi_1 \dots \chi_n \rightarrow \psi$. \square

3. Stable parts and blueprints

3.1. Introduction

Lemma 2.10 showed that the decidability of type inhabitation for $\text{BB}'\text{IW}$ is equivalent to the decidability of Λ_{NF} -inhabitation. The rest of the paper is devoted to the elaboration of a decision algorithm for the latter problem.

The problem we shall examine throughout Sections 3 and 4 is that if an inhabitant is not of minimal size, is there any way to transform it (with the help of grafts and/or another compression of some sort) into a smaller inhabitant with the same type? This question is not easy because we are dealing with a lambda calculus restricted with strong structural constraints (right-maximality). There are, however, simple situations in which an inhabitant is obviously not of minimal size.

Consider a Λ_{NF} -inhabitant M and two subterms N, P of M such that P is a strict subterm of N . Suppose:

- N, P are applications with the same type or abstractions with the same type.
- $\text{Free}(N) = X = (x_1, \dots, x_n)$.
- $\text{Free}(P) = Y = (y_0^1, \dots, y_{p_1}^1, \dots, y_0^n, \dots, y_{p_n}^n)$.
- $\Omega(X) = (\chi_1, \dots, \chi_n)$.
- $\Omega(Y) = (\chi_0^1, \dots, \chi_{p_1}^1, \dots, \chi_0^n, \dots, \chi_{p_n}^n)$.
- $\chi_j^i = \chi_i$ for each i, j .

In this case M is not of minimal size. Indeed, we can rename the free variables of P (letting $\rho(y_j^i) = x_i$) so as to obtain a term P' with the same size as P and the same type, and with the same free variables as N . The subterm N of M can be replaced with P' in M . The resulting term is a Λ_{NF} -inhabitant with the same type but of strictly smaller size.

However, this simple property is far from sufficient to characterise the minimal inhabitants of a formula: there are indeed formulas with inhabitants of arbitrary size

in which this situation never occurs. What we need is a more flexible way to reduce the size of non-minimal inhabitants. In particular, we need a better understanding of the available freedom of action if we are to rename the free variables of a term – possibly occurrence by occurrence – and if we want to ensure that right-maximality is preserved. This section is devoted to the proof of two key lemmas that delimit this freedom:

- In Sections 3.2, 3.3 and 3.3 we show how to build from any term $M \in \Lambda_{\text{NF}}$ a partial tree labelled with formulas. This partial tree is called the *blueprint* of M . This blueprint can be seen as a synthesised version of M that contains all and only the information required to determine whether a (non-uniform) renaming of the free variables of M will preserve hereditarily right-maximality.
- In Sections 3.5 and 3.6 we introduce a rewriting relation on blueprints that allows us to ‘extract’ sequences of formulas from a blueprint.
- In section 3.7 we prove our two key lemmas:
 - Lemma 3.15 clarifies the link between the blueprints of M and $\lambda x.M$ (provided both are in Λ_{NF}). This lemma proves, in particular, that the sequence of the types of the free variables of M (that is, $\Omega(\text{Free}(M))$) can always be extracted from its blueprint.
 - Lemma 3.16 shows that for *every* sequence of formulas $\bar{\phi}$ that can be extracted from the blueprint of M , there exists a (non-uniform) renaming of the free variables of M that will produce a term N with the same type and the same blueprint as M , and such that $\Omega(\text{Free}(N)) = \bar{\phi}$.

As a continuation of our first example, we will examine the consequences of this last result. Consider again a Λ_{NF} -inhabitant M and two subterms N, P of M such that P is a strict subterm of N and N, P are applications with the same type or abstractions with the same type. Suppose:

- The sequence $\Omega(\text{Free}(N))$ can be extracted from the blueprint of P .

This situation is a generalisation of the preceding one (in our first example $\Omega(X)$ could also be extracted from the blueprint of P – see Definition 3.10). The term M is still not of minimal size. Indeed, we may use the second key lemma to prove the existence of a (non-uniform) renaming of the free variables of P that will produce a term P' with the same type as P such that $\text{Free}(P') = \text{Free}(N)$. The term N can be replaced with P' in M .

3.2. Partial trees and trees

Definition 3.1. Let (\mathbf{A}, \leq) be the set of all finite sequences over the set \mathbb{N}_+ of natural numbers, ordered by prefix ordering. Elements of \mathbf{A} are called *addresses*. We define a *partial tree* to be any function π whose domain is a set of addresses. For each partial tree π and for each address a , we let $\pi_{|a}$ denote the partial tree $c \mapsto \pi(a \cdot c)$ of domain $\{c \mid a \cdot c \in \text{dom}(\pi)\}$.

Definition 3.2. For all partial trees π, π' and for every address a , we use $\pi[a \leftarrow \pi']$ to denote the partial tree π'' such that $\pi''_{|a} = \pi'$ and $\pi''(b) = \pi(b)$ for all $b \in \text{dom}(\pi)$ such that $a \not\leq b$.

Definition 3.3. A *tree domain* is a set $A \subseteq \mathbb{A}$ such that for all $a \in A$:

- Every prefix of a is in A .
- For every integer $i > 0$, if $a \cdot (i) \in A$, then $a \cdot (j) \in A$ for each $j \in \{1, \dots, i - 1\}$.

A tree domain A is *finitely branching* if and only if for each $a \in A$, there exists an $i > 0$ such that $a \cdot (i)$ is undefined. We define a *tree* to be any function whose domain is a tree domain.

In the rest of this paper, terms will be freely identified with trees. We identify:

- x with the tree mapping ε to x ;
- $\lambda x.M$ with the tree τ mapping ε to λx and such that $\tau_{|(1)}$ is the tree of M ;
- (M_1M_2) with the tree τ mapping ε to $@$ and such that $\tau_{|(i)}$ is the tree of M_i for each $i \in \{1, 2\}$.

3.3. Blueprints

Definition 3.4. Let \mathfrak{S} be the signature consisting of all formulas and all symbols of the form $@_\phi$ where ϕ is a formula. Each formula is considered as a symbol of null arity. Each $@_\phi$ has arity 2.

We define a *blueprint* to be any finite partial tree $\alpha : A \rightarrow \mathfrak{S}$ satisfying the condition that for each $a \in A$, if $\alpha(a) = @_\phi$, then $\alpha_{|a(1)}$ and $\alpha_{|a(2)}$ are of non-empty domains. We define a *rooted blueprint* to be a blueprint α such that $\varepsilon \in \text{dom}(\alpha)$.

For each $\mathcal{S} \subseteq \mathfrak{S}$, we define a \mathcal{S} -*blueprint* to be any blueprint whose image is a subset of \mathcal{S} . We write $\mathbb{B}(\mathcal{S})$ for the set of all \mathcal{S} -blueprints, and $\mathbb{B}_\varepsilon(\mathcal{S})$ for the set of all rooted \mathcal{S} -blueprints.

Definition 3.5. For every blueprint α and every address a , the *relative depth of a in α* is the number of $b \in \text{dom}(\alpha)$ such that $b < a$. The *relative depth of α* is defined as 0 if α has an empty domain, and the maximal relative depth of an address in α otherwise.

In the rest of this paper, we will use the following notation for blueprints (see Figure 2):

- $\emptyset_{\mathbb{B}}$ denotes the blueprint of the empty domain.
- ϕ abbreviates $\varepsilon \mapsto \phi$.
- $@_\phi(\alpha_1, \alpha_2)$ denotes the (rooted) blueprint α such that

$$\begin{aligned} \alpha(\varepsilon) &= \phi \\ \alpha_{|(1)} &= \alpha_1 \\ \alpha_{|(2)} &= \alpha_1. \end{aligned}$$

- For any sequence $\bar{a} = (a_1, \dots, a_k)$ of pairwise incomparable addresses, $*_{\bar{a}}(\alpha_1, \dots, \alpha_k)$ denotes the blueprint α of minimal domain such that $\alpha_{|a_i} = \alpha_i$ for each $i \in [1, \dots, k]$.
- $*(\alpha_1, \dots, \alpha_k)$ denotes the blueprint $*_{\bar{a}}(\alpha_1, \dots, \alpha_k)$ such that $\bar{a} = ((1), \dots, (k))$.

For each blueprint α , the choice of $\bar{a}, \alpha_1, \dots, \alpha_k$ such that $\alpha = *_{\bar{a}}(\alpha_1, \dots, \alpha_k)$ is obviously not unique. The sequence $(\alpha_1, \dots, \alpha_k)$ may contain an arbitrary number of empty blueprints, hence the sequence \bar{a} may be of arbitrary length. Also, α can be rooted (if $k = 1, a_1 = \varepsilon$ and α_1 is rooted) or empty (if $k = 0$ or $\alpha_1 = \dots = \alpha_k = \emptyset_{\mathbb{B}}$). These ambiguities will not




Fig. 2. The construction of blueprints using the notation of Section 3.3. In the upper diagram, the blueprints α and β must be non-empty. Although $\alpha_1, \dots, \alpha_k$ are displayed from left to right, the sequence (a_1, \dots, a_k) need not be lexicographically ordered.

be difficult to deal with, but they will require us to take a few precautions in our proofs and definitions by induction on blueprints.

3.4. Blueprint of a term

Definition 3.6. For all $M \in \Lambda_{NF}$, the *stable part* of M is the set of all $a \in \text{dom}(M)$ such that $\text{Free}(M|_a) \subseteq \text{Free}(M)$ and $M|_a$ is a variable or an application.

It is easy to check that our conventions (no variable is simultaneously free and bound in a term) ensure that the stable part of a term does not depend on the choice of variable names. Since M is in normal form, M has an empty stable part if and only if it is closed.

Definition 3.7. For all $M \in \Lambda_{NF}$, we define the *blueprint of M* to be the function α mapping each a in the stable part of M to:

- ψ if $M|_a$ is a variable of type ψ ,
- $@_{\psi}$ if $M|_a$ is an application of type ψ .

We will use $M \Vdash \alpha$ to denote the judgment ‘ M has blueprint α ’ (Figure 3).

If $M = (M_1M_2) \in \Lambda_{NF}$, $M : \phi$, $M_1 \Vdash \alpha_1$, $M_2 \Vdash \alpha_2$, then each α_i has a non-empty domain and $(M_1M_2) \Vdash @_{\phi}(\alpha_1, \alpha_2)$ – in other words, the so-called blueprint of M is indeed a blueprint, provided the blueprints of M_1 and M_2 are blueprints. When $M = \lambda x.M_1$, the blueprint of M has the form $*(\alpha)$ – the relation between α and the blueprint of M_1 in this case will be clarified by Lemma 3.15.

Lemma 3.8. For all $M \in \Lambda_{NF}$ and forall $a \cdot b \in \text{dom}(M)$:

- (1) If $\text{Free}(M|_{a \cdot b}) \subseteq \text{Free}(M)$, then $\text{Free}(M|_{a \cdot b}) \subseteq \text{Free}(M|_a)$.
- (2) If $M|_a \Vdash \alpha$ and $M|_{a \cdot b} \Vdash \beta$, then $\alpha|_b = \beta$.




Fig. 3. An element of Λ_{NF} with its blueprint ($x_0 < x_1 < y_1, x_2 < x_3 < y_0 < y_2, x_1 < y_0 < y_2$).




Fig. 4. Full reductions of $@_{\psi}(\chi \rightarrow \psi, @_{\chi}(\phi \rightarrow \chi, \phi))$ to $\emptyset_{\mathbb{B}}$.

Proof. The first part is a consequence of our bound variable convention (see Section 2.1), since if

$$\begin{aligned} \text{Free}(M) &= X \\ \text{Free}(M|_a) &= X' \cup Y \end{aligned}$$

where $X' \subseteq X$ and X, Y are disjoint, then every element of $\text{Free}(M|_{a \cdot b})$ in X is also an element of X' . Thus if $a \cdot b$ is in the stable part of M , then b is also in the stable part of $M|_a$.

The second part is equivalent to the first. □




Fig. 5. Principle of blueprint reduction.

3.5. Extraction of the formulas of a blueprint

Definition 3.9. The judgment ‘ β is the blueprint obtained by extracting the formula ϕ at the address a in the blueprint α ’ (written $\alpha \triangleright_{\phi}^a \beta$) is defined inductively by:

- (1) $\phi \triangleright_{\phi}^{\varepsilon} \emptyset_{\mathbb{B}}$.
- (2) If $\alpha \triangleright_{\phi}^a \beta$, then $@_{\psi}(\gamma, \alpha) \triangleright_{\phi}^{(2)^a} *(\gamma, \beta)$.
- (3) If $\alpha \triangleright_{\phi}^a \beta$, then $*_{(b,c_1,\dots,c_k)}(\alpha, \gamma_1, \dots, \gamma_k) \triangleright_{\phi}^{b^a} *_{(b,c_1,\dots,c_k)}(\beta, \gamma_1, \dots, \gamma_n)$.

In (2) we assume, of course, that α and γ are non-empty. In (3) we assume $b \neq \varepsilon$ in order to avoid circularity.

For instance (see Figure 4):

$$\begin{aligned}
 - \quad & @_{\psi}(\chi \rightarrow \psi, @_{\chi}(\phi \rightarrow \chi, \phi)) \triangleright_{\phi}^{(2,2)} *(\chi \rightarrow \psi, *(\phi \rightarrow \chi, \emptyset_{\mathbb{B}})) \\
 & \quad \triangleright_{\phi \rightarrow \chi}^{(2,1)} *(\chi \rightarrow \psi, *(\emptyset_{\mathbb{B}}, \emptyset_{\mathbb{B}})) \\
 & \quad \triangleright_{\chi \rightarrow \psi}^{(1)} *(\emptyset_{\mathbb{B}}, *(\emptyset_{\mathbb{B}}, \emptyset_{\mathbb{B}})) = \emptyset_{\mathbb{B}}. \\
 - \quad & @_{\psi}(\chi \rightarrow \psi, @_{\chi}(\phi \rightarrow \chi, \phi)) \triangleright_{\phi}^{(2,2)} *(\chi \rightarrow \psi, *(\phi \rightarrow \chi, \emptyset_{\mathbb{B}})) \\
 & \quad \triangleright_{\chi \rightarrow \psi}^{(1)} *(\emptyset_{\mathbb{B}}, *(\phi \rightarrow \chi, \emptyset_{\mathbb{B}})) \\
 & \quad \triangleright_{\phi \rightarrow \chi}^{(2,1)} *(\emptyset_{\mathbb{B}}, *(\emptyset_{\mathbb{B}}, \emptyset_{\mathbb{B}})) = \emptyset_{\mathbb{B}}.
 \end{aligned}$$

When $\alpha \triangleright_{\phi}^a \beta$, the blueprint β can be seen as α in which the formula ϕ at a is erased together with all $@$'s in the path to a . At each $@$, this path must follow the right-hand branch of $@$ (see Figure 5). The constraints on the construction of blueprints imply the existence of at least one such path in every non-empty blueprint, even if it is not the blueprint of a term.

3.6. Sets of extractible sequences

Definition 3.10. For each formula ϕ , let \triangleright_{ϕ} be the relation defined by $\alpha \triangleright_{\phi} \beta$ if and only if there exists a such that $\alpha \triangleright_{\phi}^a \beta$. We write \triangleright_{ϕ}^+ for the transitive closure of \triangleright_{ϕ} . For each α , we write $\mathbf{IF}(\alpha)$ for the set of all sequences (ϕ_1, \dots, ϕ_n) such that $\alpha \triangleright_{\phi_n}^+ \dots \triangleright_{\phi_1}^+ \emptyset_{\mathbb{B}}$.

The set $\mathbf{IF}(\alpha)$ is what we called ‘set of extractible sequences of α ’ in Section 3.1. Note that $\mathbf{IF}(\emptyset_{\mathbb{B}}) = \{\varepsilon\}$. If $\alpha \neq \emptyset_{\mathbb{B}}$, then all elements of $\mathbf{IF}(\alpha)$ are non-empty sequences. Note also that each \triangleright -reduction strictly decreases the cardinality of the domain of a blueprint,




Fig. 6. Shuffling of two sequences. The chunks of F and G need not have the same size – some of them can be empty. Every contraction of the resulting sequence belongs to $\otimes(F, G)$. Each contraction also belongs to $\odot(F, G)$ when F, G are non-empty and the last chunk G^p of G is non-empty.

therefore $\mathbb{IF}(\alpha)$ is a finite set for all α . We now introduce the notion of a *shuffle*, which will allow us to characterise $\mathbb{IF}(\alpha)$ according to the structure of α .

Definition 3.11. A *contraction* of a sequence F is either the sequence F or a sequence $G \cdot (f) \cdot H$ where $G \cdot (f) \cdot H$ is a contraction of F .

Definition 3.12. For all finite sequences F_1, \dots, F_n , we define a *shuffle* of (F_1, \dots, F_n) to be any sequence $F_1^1 \cdots F_n^1 \cdots F_1^p \cdots F_n^p$ such that $F_i^1 \cdots F_i^p = F_i$ for each i . For each tuple of sets of finite sequences $(\mathcal{F}_1, \dots, \mathcal{F}_n)$, we write $\otimes(\mathcal{F}_1, \dots, \mathcal{F}_n)$ for the closure under contraction of the set of shuffles of elements of $\mathcal{F}_1 \times \cdots \times \mathcal{F}_n$.

Definition 3.13. Given two non-empty finite sequences F_1, F_2 , we define a *right-shuffle* of (F_1, F_2) to be any sequence $F_1^1 \cdot F_2^1 \cdots F_1^p \cdot F_2^p$ such that $F_i^1 \cdots F_i^p = F_i$ for each i and $F_2^p \neq \varepsilon$. For each pair of sets of non-empty finite sequences $(\mathcal{F}_1, \mathcal{F}_2)$, we write $\odot(\mathcal{F}_1, \mathcal{F}_2)$ for the closure under contraction of the set of right-shuffles of elements $\mathcal{F}_1 \times \mathcal{F}_2$.

The principle of (right-)shuffling is shown in Figure 6. The following properties follow from our definitions, and will be used without further reference:

- (1) If $\alpha = \emptyset_{\mathbb{B}}$, then $\mathbb{IF}(\alpha) = \{\varepsilon\}$.
- (2) If $\alpha = \phi$, then $\mathbb{IF}(\alpha) = \{(\phi)\}$.
- (3) If $\alpha = *_a(\alpha_1, \dots, \alpha_k)$, then $\mathbb{IF}(\alpha) = \otimes(\mathbb{IF}(\alpha_1), \dots, \mathbb{IF}(\alpha_k))$.
- (4) If $\alpha = @_{\phi}(\alpha_1, \alpha_2)$, then $\mathbb{IF}(\alpha) = \odot(\mathbb{IF}(\alpha_1), \mathbb{IF}(\alpha_2))$.

3.7. Abstraction versus extraction

Lemma 3.14. Suppose $\{a_1, \dots, a_p\} = \{b_1, \dots, b_p\}$, and:

- $\alpha \triangleright_{\chi}^{a_1} \cdots \triangleright_{\chi}^{a_p} \beta$.
- $\alpha \triangleright_{\chi}^{b_1} \cdots \triangleright_{\chi}^{b_p} \beta'$.

Then $\beta = \beta'$.

Proof. The proof is by an easy induction on α . □

Recall that for every strictly increasing sequence of variables $X = (x_1, \dots, x_n)$, we write $\Omega(X)$ for the sequence of the types of x_1, \dots, x_n . We will now clarify the link between the blueprint α of a term M and the blueprint of $\lambda x.M$.




Fig. 7. How the blueprint of M evolves into the blueprint of $\lambda x.M$

In particular, the next lemma shows that if $M, \lambda x.M \in \Lambda_{NF}$, then M and $\lambda x.M$ have blueprints α and β if and only if there exist a_0, \dots, a_p such that

$$\begin{aligned} \{a_0, \dots, a_p\} &= \{a \mid M|_a = x\} \\ \alpha &\triangleright_{\chi}^{a_0} \dots \triangleright_{\chi}^{a_p} \alpha' \\ \beta &= *(\alpha') \end{aligned}$$

(see Figure 7).

Lemma 3.15. Suppose $M \in \Lambda_{NF}$ has blueprint α , with $\text{Free}(M) = (x_1, \dots, x_n)$ and $\Omega(x_1, \dots, x_n) = (\chi_1, \dots, \chi_n)$. For each $i \in [0, \dots, n]$:

- Let α_i be the restriction of α to $\text{dom}(\alpha) \cap \{a \mid \text{Free}(M|_a) \subseteq \{x_1, \dots, x_i\}\}$.
- Let β_i be the blueprint of $\lambda x_{i+1} \dots x_n.M$.

Then:

- (1) For each $i \in [0, \dots, n]$, we have $\text{dom}(\beta_i) = \{1^{n-1} \cdot a \mid a \in \text{dom}(\alpha_i)\}$ and $\beta_{i|1^{n-i}} = \alpha_i$.
- (2) For each $i \in]0, \dots, n]$:
 - (a) There exist $a_0^i, \dots, a_{p_i}^i$ such that

$$\{a_0^i, \dots, a_{p_i}^i\} = \{a \mid M|_a = x_i\}$$

and

$$\alpha_i \triangleright_{\chi_i}^{a_0^i} \dots \triangleright_{\chi_i}^{a_{p_i}^i} \alpha_{i-1}.$$

(b) If

$$\{b_0, \dots, b_{p_i}\} = \{a \mid M|_a = x_i\}$$

and

$$\alpha_i \triangleright_{\chi_i}^{b_0} \dots \triangleright_{\chi_i}^{b_{p_i}} \alpha',$$

then $\alpha' = \alpha_{i-1}$.(3) $(\chi_1, \dots, \chi_n) \in \mathbb{F}(\alpha)$.

Proof. Property (1) follows immediately from the definition of a blueprint. Since $\alpha_n = \alpha$ and $\alpha_0 = \emptyset_{\mathbb{B}}$, Property (3) follows from Property (2a). Property (2b) follows from Property (2a) and Lemma 3.14. To prove (2a), we first introduce the following notation.

For each $N \in \Lambda_{\text{NF}}$, we let ρ_N be the least partial function satisfying the following conditions:

- For every blueprint γ , we have $\rho_N(\varepsilon, \gamma) = \gamma$.
- For every finite sequence of variables Y and for every blueprint γ , if

$$\begin{aligned} \rho_N(Y, \gamma) &= \delta \\ \{b \mid N|_b = y\} &= \{b_0, \dots, b_m\} \end{aligned}$$

and

$$\delta \triangleright_{\chi}^{b_0} \dots \triangleright_{\chi}^{b_m} \delta',$$

then $\rho_M((y) \cdot Y, \gamma) = \delta'$.

By Lemma 3.14, if

$$\{b \mid N|_b = y\} = \{b_0, \dots, b_m\} = \{c_0, \dots, c_m\}$$

and

$$\begin{aligned} \delta \triangleright_{\chi}^{b_0} \dots \triangleright_{\chi}^{b_m} \delta' \\ \delta \triangleright_{\chi}^{c_0} \dots \triangleright_{\chi}^{c_m} \delta'' \end{aligned}$$

then $\delta' = \delta''$, so ρ_N is indeed a function. For each finite sequence of variables Y' and for each blueprint γ , we let $\mu_N(Y', \gamma)$ be the restriction of γ to $\text{dom}(\gamma) \cap \{b \mid \text{Free}(N|_b) \subseteq Y'\}$.

We shall prove by induction on M that for all pairs (X, X') such that $\text{Free}(M) = X \cdot X'$, we have $\mu_M(X, \alpha) = \rho_N(X', \alpha)$. In particular, for all $i > 0$, we have

$$\begin{aligned} \alpha_{i-1} &= \mu_M((x_1, \dots, x_{i-1}), \alpha) \\ &= \rho_M((x_i, \dots, x_n), \alpha) \\ &= \rho_M((x_i), \rho_M((x_{i+1}, \dots, x_n), \alpha)) \\ &= \rho_M((x_i), \mu_M((x_1, \dots, x_i), \alpha)) \\ &= \rho_M((x_i), \alpha_i), \end{aligned}$$

so (2a) holds.

The case $X' = \varepsilon$ is immediate, so we will assume that X' is a non-empty suffix of $\text{Free}(M)$ and consider cases:

— M is equal to a variable:

This case follows immediately from the definitions.

— $M = (M_1 M_2)$, $M_1 \Vdash \gamma_1$ and $M_2 \Vdash \gamma_2$:

There exist X_1, X_2, X'_1, X'_2 such that:

$$\begin{aligned} X_1 \cup X_2 &= X \\ X'_1 \cup X'_2 &= X' \\ \text{Free}(M_j) &= X_j \cdot X'_j \quad \text{for each } j \in \{1, 2\}. \end{aligned}$$

We have $\alpha = @_{\psi}(\gamma_1, \gamma_2)$ where ψ is the type of M , and

$$\mu_M(X, \alpha) = *(\mu_{M_1}(X_1, \gamma_1), \mu_{M_2}(X_2, \gamma_2)).$$

By the induction hypothesis,

$$\mu_{M_i}(X_i, \gamma_i) = \rho_{M_i}(X'_i, \gamma_i)$$

for each i . The sequence X' is non-empty, so the final elements of X' and X'_2 are equal.

Assume $X' = X'' \cdot (x)$ and $X'_2 = X''_2 \cdot (x)$. If x is not the final element of X'_1 ,

$$\begin{aligned} \rho_M(X', \alpha) &= \rho_M(X'' \cdot (x), @_{\psi}(\gamma_1, \gamma_2)) \\ &= \rho_M(X'_1 \cup X''_2, *(\gamma_1, \rho_{M_2}((x), \gamma_2))) \\ &= *(\rho_{M_1}(X'_1, \gamma_1), \rho_{M_2}(X''_2, \rho_{M_2}((x), \gamma_2))) \\ &= *(\rho_{M_1}(X'_1, \gamma_1), \rho_{M_2}(X''_2 \cdot (x), \gamma_2)) \\ &= *(\rho_{M_1}(X'_1, \gamma_1), \rho_{M_2}(X'_2, \gamma_2)). \end{aligned}$$

Otherwise, $X'_1 = X''_1 \cdot (x)$ and we have

$$\begin{aligned} \rho_M(X', \alpha) &= \rho_M(X', @_{\psi}(\gamma_1, \gamma_2)) \\ &= \rho_M(X''_1 \cup X''_2, *(\rho_{M_1}((x), \gamma_1), \rho_{M_2}((x), \gamma_2))) \\ &= *(\rho_{M_1}(X''_1, \rho_{M_1}((x), \gamma_1)), \rho_{M_2}(X''_2, \rho_{M_2}((x), \gamma_2))) \\ &= *(\rho_{M_1}(X''_1 \cdot (x), \gamma_1), \rho_{M_2}(X''_2 \cdot (x), \gamma_2)) \\ &= *(\rho_{M_1}(X'_1, \gamma_1), \rho_{M_2}(X'_2, \gamma_2)). \end{aligned}$$

In either case

$$\begin{aligned} \rho_M(X', \alpha) &= *(\rho_{M_1}(X'_1, \gamma_1), \rho_{M_2}(X'_2, \gamma_2)) \\ &= *(\mu_{M_1}(X_1, \gamma_1), \mu_{M_2}(X_2, \gamma_2)) \\ &= \mu_M(X, \alpha). \end{aligned}$$

— $M = \lambda x.M_1$, $M_1 \Vdash \gamma_1$:

By the induction hypothesis,

$$\begin{aligned} \mu_{M_1}(X, \gamma_1) &= \rho_{M_1}(X' \cdot (x), \gamma_1) \\ &= \rho_{M_1}(X', \rho_{M_1}((x), \gamma_1)) \\ &= \rho_{M_1}(X', \mu(X \cdot X', \gamma_1)) \\ &= \rho_{M_1}(X', \alpha_{(1)}). \end{aligned}$$




Fig. 8. A non-uniform renaming of the variables of M , based on an alternate extraction of the formulas of its blueprint.

Moreover,

$$\mu_{M_1}(X, \gamma_1) = \mu_{M_1}(X, \mu_{M_1}(X \cdot X', \gamma_1)) = \mu_{M_1}(X, \alpha_{|(1)}).$$

Therefore,

$$\mu_{M_1}(X, \alpha_{|(1)}) = \rho_{M_1}(X', \alpha_{|(1)}),$$

so

$$\mu_{M_1}(X, \alpha) = \rho_{M_1}(X', \alpha). \quad \square$$

Thus the full sequence of the types of the free variables of M can be extracted from its blueprint. The next lemma shows that, conversely, for each sequence $\bar{\chi}$ in $\mathbb{F}(\alpha)$, there exists a term N with the same domain and blueprint, that has the same type as M , and such that the sequence of types of the free variables of N is equal to $\bar{\chi}$ – see Figure 8.

Lemma 3.16. Let $M \in \Lambda_{\text{NF}}$ be a term with blueprint α . Suppose

$$\alpha \triangleright_{\omega_m}^{b_0^m} \cdots \triangleright_{\omega_m}^{b_{p_m}^m} \cdots \triangleright_{\omega_1}^{b_0^1} \cdots \triangleright_{\omega_1}^{b_{p_1}^1} \not\triangleright_{\mathbb{B}}.$$

Then for every strictly increasing sequence of variables $Y = (y_1, \dots, y_m)$ such that $\Omega(Y) = (\omega_1, \dots, \omega_m)$, there exists N with the same domain and blueprint and with the same type as M , and such that $\text{Free}(N) = Y$ and $\{b \mid N|_b = y_i\} = \{b_1^i, \dots, b_{p_i}^i\}$ for each i .

Proof. The proof is by induction on M :

— M is a variable:

This case is obvious.

— $M = (M_1 M_2)$:

This case follows easily from the induction hypothesis.

— $M = \lambda x. M_1 : \phi \rightarrow \psi$ with $M_1 \Vdash \gamma$:

Let $Y' = (y_1, \dots, y_m, x)$. By Lemma 3.15 (2a), there exist a_1, \dots, a_p such that

$$\{a_1, \dots, a_p\} = \{a \mid M|_a = x\}$$

and

$$\gamma \triangleright_{\phi}^{a_0} \cdots \triangleright_{\phi}^{a_p} \gamma' = \alpha_{|1}.$$

Now

$$\alpha \triangleright_{\omega_m}^{b_0^m} \cdots \triangleright_{\omega_m}^{b_p^m} \dots \triangleright_{\omega_1}^{b_0^1} \cdots \triangleright_{\omega_1}^{b_{p_1}^1} \not\in \mathbb{B},$$

so each b_j^i has the form $(1) \cdot c_j^i$. Furthermore,

$$\gamma \triangleright_{\phi}^{a_0} \cdots \triangleright_{\phi}^{a_p} \triangleright_{\omega_m}^{c_0^m} \cdots \triangleright_{\omega_m}^{c_{p_m}^m} \dots \triangleright_{\omega_1}^{c_0^1} \cdots \triangleright_{\omega_1}^{c_{p_1}^1} \not\in \mathbb{B}.$$

By the induction hypothesis, there exists N_1 with the same domain and blueprint and with the same type as M_1 such that

$$\begin{aligned} \text{Free}(N_1) &= Y' \\ \{a \mid N_1|_a = x\} &= \{a_0, \dots, a_p\} \\ \{c \mid N_1|_c = y_i\} &= \{c_0^i, \dots, c_{p_i}^i\} \text{ for each } i. \end{aligned}$$

By Lemma 3.15 (2b), we have $\lambda x.N_1 \Vdash \alpha$, so we may take $N = \lambda x.N_1$. □

4. Vertical compressions and compact terms

This section provides a partial characterisation of minimal inhabitants. In Section 4.1, we make a simple observation on the relative depths of their blueprints, together with the following easy consequence of the subformula property (Lemma 2.5). If M is a minimal Λ_{NF} -inhabitant of ϕ , then for all addresses a in M the blueprint of $M|_a$ has relative depth at most $k \times p$, where:

- k is the number of λ in the path from the root to M to a .
- p is the number of subformulas of ϕ .

We say any Λ_{NF} -inhabitant satisfying this condition is *locally compact*. In Section 4.2 we introduce the notion of a *vertical compression* of a blueprint. A (strict) vertical compression of β is obtained by taking any address b in β and then grafting $\beta|_b$ at any address $a < b$ such that $\beta(a) = \beta(b)$. The vertical compressions of β are all blueprints obtained by applying this transformation to β zero or more times. The key property of these compressions is (see Figure 9):

- If M has blueprint β and α is a vertical compression of β , the compression of β into α can be mimicked by a compression of M into an HRM-term in the following sense. Assuming $\alpha = \beta[a \leftarrow \beta|_b]$ (the base case), the term $Q = M[a \leftarrow M|_b]$ is *not* in general an HRM-term. However, *there exists* an HRM-term M' with the same domain as Q and with the same type as M . Moreover, M' and M are applications with the same type or abstractions with the same type.

We again consider a Λ_{NF} -inhabitant M and two addresses a, b such that $a < b$, and $M|_a$ and $M|_b$ are applications with the same type or abstractions with the same type. Suppose:

- there exists a vertical compression α' of the blueprint of $M|_b$ such that the sequence $\Omega(\text{Free}(M|_a))$ can be extracted from α' .




Fig. 9. How the compression of terms is able to follow the compression of blueprints.

This situation is a generalisation of the last example in Section 3.1 (in which α' was equal to the blueprint of $M|_b$, which is thereby a trivial compression of this blueprint). The term M is not minimal. Indeed, the key property above implies the existence of a term N of blueprint α' whose size is not greater than the size of $M|_b$, and such that $N, M|_b, M|_a$ are applications with the same type, or abstractions with the same type. By Lemma 3.16, there exists a term P with the same type and domain as N such that $\text{Free}(P) = \text{Free}(M|_a)$. The graft of P at a yields an inhabitant of strictly smaller size.

We will say all inhabitants in which the preceding situation does not occur are *compact*. All inhabitants of minimal size are of course compact. As we shall see in Section 6, we will not need a sharper characterisation of minimal inhabitants. For every formula ϕ , the set of compact inhabitants of ϕ is actually a *finite* set, and our decision method will consist of the exhaustive computation of their domains.

4.1. Depths of the blueprints of minimal inhabitants

Definition 4.1. Two terms $M, M' \in \Lambda_{\text{NF}}$ are of the same kind if and only if they are both variables or both applications, or both abstractions, and if they have the same type.

Definition 4.2. For all formulas ϕ , we write $\text{Sub}(\phi)$ for the set of all subformulas of ϕ .

Definition 4.3. Let $M \in \Lambda_{\text{NF}}$. Let a be any address in M . Let (a_1, \dots, a_m) be the strictly increasing sequence of all prefixes of a . Let $(\lambda x_1, \dots, \lambda x_k)$ be the subsequence of $(M(a_1), \dots, M(a_m))$ consisting of all labels of the form λx . We write $\Lambda(M, a)$ for (x_1, \dots, x_k) .

Definition 4.4. Let M be a Λ_{NF} -inhabitant of ϕ . We say that M is *locally compact* if for all addresses a in M , the blueprint of $M|_a$ has relative depth at most $|\Lambda(M, a)| \times |\text{Sub}(\phi)|$.

Lemma 4.5. Let M be a Λ_{NF} -inhabitant of ϕ . If M is not locally compact, then there exist two addresses b, b' such that $b < b'$, $M|_b$ and $M|_{b'}$ are of the same kind and $\text{Free}(M|_b) = \text{Free}(M|_{b'})$. Moreover, M is not a Λ_{NF} -inhabitant of ϕ of minimal size.

Proof. For each address a in $\text{dom}(M)$, let α_a be the blueprint of $M|_a$ and let $X_a = \text{Free}(M|_a)$. We assume the existence of an α_a with relative depth

$$n > |\Lambda(M, a)| \times |\text{Sub}(\phi)|.$$




Fig. 10. Proof of Lemma 4.5.

So there exist $b_1, \dots, b_{n+1} \in \text{dom}(\alpha_a)$ such that $b_1 < \dots < b_n < b_{n+1}$. By Lemma 3.8 (1), we have

$$X_{a.b_n} \subseteq \dots \subseteq X_{a.b_1} \subseteq \Lambda(M, a).$$

By Lemma 2.5, each $\phi_{a.b_i}$ is a subformula of ϕ . Hence there exist i, j such that $i < j$ and

$$(X_{a.b_i}, \phi_{a.b_i}) = (X_{a.b_j}, \phi_{a.b_j}),$$

that is, $M|_{a.b_i}$ and $M|_{a.b_j}$ are applications with the same type and with the same free variables (see Figure 10). Now, let $M' = M[a \cdot b_i \leftarrow M|_{a.b_j}]$. The term M' is a Λ_{NF} -inhabitant of ϕ of strictly smaller size. \square

4.2. Vertical compression of a blueprint

Definition 4.6. Let \uparrow be the least reflexive and transitive binary relation on blueprints satisfying the condition that if $a, b \in \text{dom}(\beta)$, $a < b$ and $\beta(a) = \beta(b)$, then $\beta[a \leftarrow \beta|_b] \uparrow \beta$.

Lemma 4.7. Suppose $M \in \Lambda_{\text{NF}}$, $M : \phi$, $M \Vdash \beta$ and $\alpha \uparrow \beta$. There exists a term $M' \in \Lambda_{\text{NF}}$ of the same kind as M , with blueprint α and such that $|\text{dom}(M')| \leq |\text{dom}(M)|$.

Proof. It is enough to consider the case of $\alpha = \beta[a \leftarrow \beta|_b]$ with $a, b \in \text{dom}(\beta)$, $a < b$ and $\beta(a) = \beta(b)$.

We prove the existence of M' by induction on the length of a .

— $a = \varepsilon$:

In this case M is necessarily an application and $\beta(\varepsilon) = \beta(b) = @_\phi$, so $M|_b$ is an application with type ϕ , and we can take $M' = M|_b$.

— $a \neq \varepsilon$.

We consider cases:

- $M = (M_1 M_2)$, $M_1 \Vdash \beta_1$, $M_2 \Vdash \beta_2$, $a = (i) \cdot a_i$ and $b = (i) \cdot b_i$:
By the induction hypothesis, there exists M'_i with blueprint

$$\alpha_i = \beta_i[a_i \leftarrow \beta_{i|b_i}] = \beta_i[a_i \leftarrow \beta_{i|b}],$$

of the same kind as M_i and such that $\text{dom}(M'_i) \leq \text{dom}(M_i)$. Let $j = 1$ if $i = 2$, otherwise, let $j = 2$. Let $(M'_j, \alpha_j) = (M_j, \beta_j)$. Let $X = (x_1, \dots, x_n)$ be the strictly increasing sequence of all variables free or bound in M'_2 . Let $Y = (y_1, \dots, y_n)$ be a strictly increasing sequence of variables such that $\Omega(X) = \Omega(Y)$ and y_1 is greater than or equal to the greatest variable of M'_1 . Let M'_2 be the term obtained by replacing each x_i by y_i in M'_2 . We can take $M' = (M'_1 M'_2)$.

- $M = \lambda x.M_1$, $M_1 \Vdash \beta_1$, $x : \chi$, $a = (1) \cdot a_1$ and $b = (1) \cdot b_1$:
Since $a, b \in \text{dom}(\beta)$, we also have $a_1, b_1 \in \text{dom}(\beta_1)$. By the induction hypothesis, there exists M'_1 of the same kind as M_1 , with blueprint

$$\alpha_1 = \beta_1[a_1 \leftarrow \beta_{1|b_1}]$$

and such that $\text{dom}(M'_1) \leq \text{dom}(M_1)$. By Lemma 3.15 (2a), there exist $\gamma_1, c_0, \dots, c_p$ such that

$$\begin{aligned} \{c_0, \dots, c_p\} &= \{c \mid M_{|c} = x\} \\ \beta_1 &\triangleright_{\chi}^{c_0} \cdots \triangleright_{\chi}^{c_p} \gamma_1 \\ \beta &= *(\gamma_1). \end{aligned}$$

Since $a, b \in \text{dom}(\alpha)$, we know a_1 and c_i are incomparable addresses for all i . Hence

$$\alpha_1 = \beta_1[a_1 \leftarrow \beta_{1|b_1}] \triangleright_{\chi}^{c_0} \cdots \triangleright_{\chi}^{c_p} \gamma_1[a_1 \leftarrow \beta_{1|b_1}] = \beta[a \leftarrow \beta_{|b}]_{|(1)} = \alpha_{|1}.$$

By Lemma 3.16, there exists a term M''_1 with the same type and domain as M'_1 such that the greatest variable y free in M''_1 has type χ and

$$\{c \mid M''_{1|c} = y\} = \{c_0, \dots, c_p\}.$$

By Lemma 3.15 (2b), we have $\lambda y.M''_1 \Vdash \alpha$, so we may take $M' = \lambda y.M''_1$. \square

Definition 4.8. We say a term $M \in \Lambda_{\text{NF}}$ is *compact* if there are no a, b, a' such that $a < b$, $M_{|a}$ and $M_{|b}$ are of the same kind, $M_{|b} \Vdash \alpha_b$, $a' \uparrow \alpha_b$ and $\Omega(\text{Free}(M_{|a})) \in \mathbb{F}(a')$.

Lemma 4.9.

- (1) Every Λ_{NF} -inhabitant of minimal size is compact.
- (2) Every compact Λ_{NF} -inhabitant of ϕ is locally compact.

Proof. Let M be an arbitrary Λ_{NF} -inhabitant of ϕ .

- (1) Assume M is not compact. Let a, b be such that $a < b$, $M_{|a}$ and $M_{|b}$ are of the same kind, $M_{|b} \Vdash \alpha_b$, $a' \uparrow \alpha_b$, $\text{Free}(M_{|a}) = X_a$ and $\Omega(X_a) \in \mathbb{F}(a')$ – see Figure 11. By Lemma 4.7, there exists a term $N \in \Lambda_{\text{NF}}$ with blueprint a' , of the same kind as $M_{|b}$ and such that $|\text{dom}(N)| \leq |\text{dom}(M_{|b})|$. By Lemma 3.16, there exists $P \in \Lambda_{\text{NF}}$ with

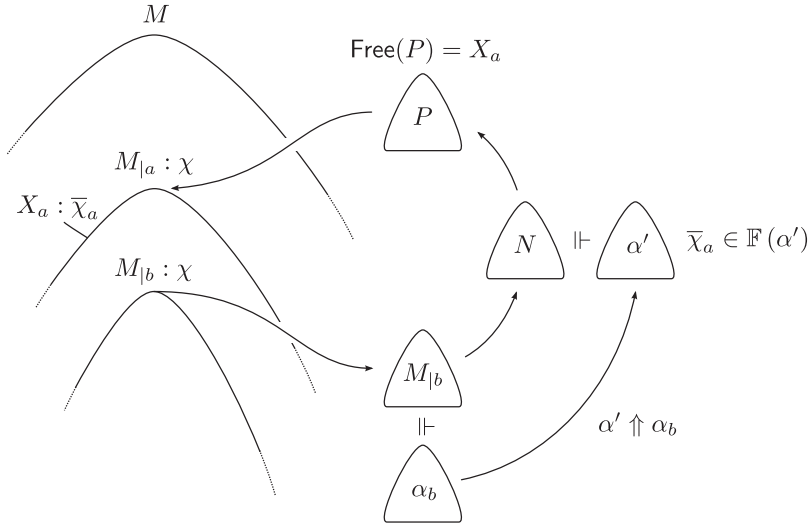


Fig. 11. Proof of Lemma 4.9, part (1).

blueprint α' , of the same kind as N , such that $\text{dom}(P) = \text{dom}(N)$ and $\text{Free}(P) = X_a$. The term $M[a \leftarrow P]$ is then a Λ_{NF} -inhabitant of ϕ of smaller size.

- (2) Suppose M meets the conditions of Lemma 4.5. Let $\alpha_{b'}$ be the blueprint of $M|_{b'}$. By Lemma 3.15 (3), we have $\Omega(\text{Free}(M|_b)) = \Omega(\text{Free}(M|_{b'})) \in \mathbb{F}(\alpha_{b'})$. Since the relation \uparrow is reflexive, M is not compact. □

5. Shadows

So far we have isolated two properties shared by all minimal inhabitants (Lemma 4.9). We shall now exploit these properties to design a decision method for the inhabitation problem.

In Sections 5.1 and 5.2 we show how to associate with each locally compact inhabitant M of a formula ϕ , a tree with the same domain as M , which we call the *shadow* of M . At each address a , this tree is labelled with a triple of the form $(\bar{\chi}_a, \gamma_a, \phi_a)$ where ϕ_a is the type of $M|_a$, the sequence $\bar{\chi}_a$ is $\Omega(\text{Free}(M|_a))$, and γ_a is a ‘transversal compression’ of the blueprint α_a of $M|_a$ (Definitions 5.1 and 5.2). Recall that $\bar{\chi}_a \in \mathbb{F}(\alpha_a)$ – see Lemma 3.15 (3). The blueprint γ_a can be viewed as a synthesised version of α_a with the same relative depth but smaller ‘width’, and such that $\bar{\chi}_a \in \mathbb{F}(\gamma_a) \subseteq \mathbb{F}(\alpha_a)$.

Each tree prefix of the shadow of M belongs to a finite set that is effectively computable from ϕ and the domain of this prefix. In particular, we can compute all possible values for its labels, even without full knowledge of M – or even without knowledge of the existence of M . The key property satisfied by this shadow at every address a is:

- For each $\gamma' \uparrow \gamma_a$, there exists $\alpha' \uparrow \alpha_a$ such that $\mathbb{F}(\gamma') \subseteq \mathbb{F}(\alpha')$.

This property is sufficient to detect the non-compactness of M for a pair of addresses (a, b) simply from the knowledge of $\bar{\chi}_a, \phi_a, \gamma_b, \phi_b$ and the arity of the nodes at a and b .

Indeed, we suppose $a < b$, $\phi_a = \phi_b$ and the nodes at a and b have the same arity (1 or 2), and assume:

— There exists $\gamma' \uparrow \gamma_b$ such that $\bar{\chi}_a \in \mathbf{IF}(\gamma')$.

Then $M_{|a}$ and $M_{|b}$ are of the same kind and there exists $\alpha' \uparrow \alpha_b$ such that

$$\bar{\chi}_a = \Omega(\text{Free}(M_{|a})) \in \mathbf{IF}(\gamma') \subseteq \mathbf{IF}(\alpha'),$$

therefore M is not compact.

In Section 5.2, what we call a *shadow* is merely a tree $a \mapsto (\bar{\chi}_a, \gamma_a, \phi_a)$ of a certain shape, regardless of whether this tree is the shadow of a term or not. This shadow is *compact* if there is no pair (a, b) as above. Of course, the shadow of a compact term is always compact in this sense.

In Section 6 we will prove that for every formula ϕ , the set of shadows of the compact inhabitants of ϕ is a finite set that is effectively computable from ϕ (hence the same property holds for the set of compact inhabitants of ϕ), and we will deduce from this key property the decidability of type inhabitation for HRM-terms.

5.1. Blueprint equivalence and transversal compression

Definition 5.1. We let \equiv be the least binary relation on blueprints such that:

- (1) $\emptyset_{\mathbb{B}} \equiv \emptyset_{\mathbb{B}}$.
- (2) $\phi \equiv \phi$.
- (3) If $\alpha_1 \equiv \beta_1$, $\alpha_2 \equiv \beta_2$, then $\textcircled{\alpha}_\phi(\alpha_1, \alpha_2) \equiv \textcircled{\alpha}_\phi(\beta_1, \beta_2)$.
- (4) If $|\bar{a}| = |\bar{b}| = n$ and $\alpha_i \equiv \beta_i$ for each $i \in [1, \dots, n]$, then

$$*_\bar{a}(\alpha_1, \dots, \alpha_n) \equiv *_\bar{b}(\beta_1, \dots, \beta_n).$$

In (3), we assume $\alpha_1, \alpha_2, \beta_1, \beta_2$ are non-empty, and in (4), we assume that the elements of each sequence \bar{a}, \bar{b} are pairwise incomparable addresses. To avoid circularity, we also assume $a \neq \varepsilon$ or $b \neq \varepsilon$, and $\alpha_i, \beta_i \neq \emptyset_{\mathbb{B}}$ for at least one i .

To some extent, this equivalence allows us to consider blueprints regardless of the exact values of addresses. For instance,

$$*_\bar{a}(\alpha_1, \dots, \alpha_n) \equiv *(\alpha_1, \dots, \alpha_n) \equiv *(\alpha_n, \dots, \alpha_1),$$

also

$$*(*(\alpha, \beta), \gamma) \equiv *(\alpha, \beta, \gamma) \equiv *(\alpha, *(\beta, \gamma)),$$

and so on. It is easy to check that $\alpha \equiv \beta$ implies $\mathbf{IF}(\alpha) = \mathbf{IF}(\beta)$ – this property will be used without further reference.

Definition 5.2. For each $m \in \mathbb{N}$, we let \curvearrowright_m be the least binary relation such that:

- (1) if $\gamma_1 \equiv \dots \equiv \gamma_m \equiv \gamma_{m+1} \neq \emptyset_{\mathbb{B}}$, then $*_{\bar{a}}(\gamma_1, \dots, \gamma_m) \curvearrowright_m *_{\bar{a}(b)}(\gamma_1, \dots, \gamma_m, \gamma_{m+1})$,
- (2) if $\alpha = *_\bar{a}(\alpha_1, \dots, \alpha_n)$, $\beta = *_\bar{b}(\beta_1, \dots, \beta_p)$ and $\alpha \curvearrowright_m \beta$, then:

$$(a) \textcircled{\alpha}_\phi(\alpha, \gamma) \curvearrowright_m \textcircled{\alpha}_\phi(\beta, \gamma),$$

$$(b) \textcircled{\alpha}_\phi(\gamma, \alpha) \curvearrowright_m \textcircled{\alpha}_\phi(\gamma, \beta),$$

$$(c) \ *_{\bar{a}(c)}(\alpha_1, \dots, \alpha_n, \gamma) \prec_m \ *_{\bar{b}(c)}(\beta_1, \dots, \beta_p, \gamma).$$

We define an *m-compression* of β to be any α such that $\alpha \prec_m \beta$. The *width* of β is defined as the least $m \in \mathbb{N}$ for which there is no α such that $\alpha \prec_m \beta$.

Again the elements of $\bar{a} \cdot (b)$, $\bar{a} \cdot (c)$ and $\bar{b} \cdot (c)$ must be pairwise incomparable addresses, and α, β, γ must be non-empty. Note that for all non-empty β , we have $\emptyset_{\mathbb{B}} \prec_0 \beta$, so the empty blueprint is the only blueprint of null width. If β has width $m > 0$, then for all addresses a , for $\beta|_a = *_{\bar{a}}(\gamma_1, \dots, \gamma_k)$ and for each $\gamma_i \neq \emptyset_{\mathbb{B}}$, the sequence $(\gamma_1, \dots, \gamma_k)$ contains no more than m blueprints \equiv -equivalent to γ_i . For instance, if ϕ, ψ, χ are distinct formulas, $*(\phi, \phi, \phi, \phi, \psi, \psi, \chi)$ has width 3, $*(\omega, @_{\omega}(*(\phi, \psi), \phi), @_{\omega}(*(\psi, \phi), \phi))$ has width 2, and so on.

Definition 5.3. For each $m \in \mathbb{N}$, we write \sqsubseteq_m for the reflexive and transitive closure of the union of \equiv and \prec_m . We use \sqsubseteq_m^{\max} to denote the subset of the relation \sqsubseteq_m of all pairs with a left-hand side of width at most m .

For instance, if ϕ, ψ, χ are distinct formulas:

$$\emptyset_{\mathbb{B}} \sqsubseteq_0^{\max} *(\psi, \chi, \phi) \sqsubseteq_1^{\max} *(\chi, \phi, \phi, \psi, \psi) \sqsubseteq_2^{\max} *(\phi, \phi, \phi, \psi, \psi, \chi).$$

Of course, $\alpha \sqsubseteq_m \beta$ implies $\alpha \sqsubseteq_j \beta$ for all $j \in [1, \dots, m]$ and, clearly, $\alpha \prec_m \beta$ implies $|\text{dom}(\alpha)| < |\text{dom}(\beta)|$, so \prec_m is well founded.

Definition 5.4. For all $\mathcal{S} \subseteq \mathfrak{S}$, all $d \in \mathbb{N}$ and all $m \in \mathbb{N}$:

- We let $\mathbb{B}(\mathcal{S}, d, \infty)$ be the set of \mathcal{S} -blueprints with relative depth at most d .
- We let $\mathbb{B}(\mathcal{S}, d, m)$ be the set of all blueprints in $\mathbb{B}(\mathcal{S}, d, \infty)$ of width at most m .

Lemma 5.5. For all finite $\mathcal{S} \subseteq \mathfrak{S}$, all $d \in \mathbb{N}$ and all $m \in \mathbb{N}$:

- (1) The set $\mathbb{B}(\mathcal{S}, d, m)/\equiv$ is a finite set.
- (2) A selector $\mathbb{R}(\mathcal{S}, d, m)$ for $\mathbb{B}(\mathcal{S}, d, m)/\equiv$ is effectively computable from (\mathcal{S}, d, m) .

Proof.

- (1) Let $\mathbb{B}_{\varepsilon}(\mathcal{S}, d, m)$ be the set of all rooted blueprints in $\mathbb{B}(\mathcal{S}, d, m)$. Assuming that $\mathbb{B}_{\varepsilon}(\mathcal{S}, d, m)/\equiv$ is a finite set and that a selector $\mathbb{R}_{\varepsilon}(\mathcal{S}, d, m)$ for $\mathbb{B}_{\varepsilon}(\mathcal{S}, d, m)/\equiv$ is effectively computable from (\mathcal{S}, d, m) , we prove that $\mathbb{B}(\mathcal{S}, d, m)/\equiv$ and $\mathbb{B}_{\varepsilon}(\mathcal{S}, d + 1, m)/\equiv$ are finite sets and show how to compute a selector for each set.

Let $(\alpha_1, \dots, \alpha_k)$ be an enumeration of $\mathbb{R}_{\varepsilon}(\mathcal{S}, d, m)$, and Σ_d be the set of all functions from $\{1, \dots, k\}$ to $\{0, \dots, m\}$. For each $\beta \in \mathbb{B}(\mathcal{S}, d, m)$, there exist $\beta_1, \dots, \beta_n \in \mathbb{B}_{\varepsilon}(\mathcal{S}, d, m)$ and \bar{b} such that $\beta = *_{\bar{b}}(\beta_1, \dots, \beta_n)$. We let σ_{β} be the function mapping each $i \in \{1, \dots, k\}$ to the number of occurrences of an element \equiv -equivalent to α_i in the sequence $(\beta_1, \dots, \beta_n)$. Clearly, $\sigma_{\beta} \in \Sigma_d$ and, furthermore, for all $\beta' \in \mathbb{B}(\mathcal{S}, d, m)$, we have $\beta \equiv \beta'$ if and only if $\sigma_{\beta} = \sigma_{\beta'}$, so $\mathbb{B}(\mathcal{S}, d, m)$ is a finite set.

For each $\tau \in \Sigma_d$, let

$$\rho_{\tau} = *(\alpha_1^1, \dots, \alpha_1^{\tau(1)}, \dots, \alpha_k^1, \dots, \alpha_k^{\tau(k)})$$

where each α_i^j is equal to α_i . We have $\rho_{\tau} \in \mathbb{B}(\mathcal{S}, d, m)$ and $\sigma(\rho_{\tau}) = \tau$, that is, if $\tau, \tau' \in \Sigma_d$ and $\tau \neq \tau'$, then $\rho_{\tau} \not\equiv \rho_{\tau'}$. Hence we may define $\mathbb{R}(\mathcal{S}, d, m)$ as $\{\rho_{\tau} \mid \tau \in \Sigma_d\}$.

Now the finiteness of $\mathbb{B}_\varepsilon(\mathcal{S}, d + 1, m) / \equiv$ follows immediately from the finiteness of $\mathbb{B}(\mathcal{S}, d, m)$ and the fact that if $\beta = @_\phi(\beta_1, \beta_2)$ and $\beta' = @_\psi(\beta'_1, \beta'_2)$ are elements of $\mathbb{B}_\varepsilon(\mathcal{S}, d + 1, m)$, then $\beta_1, \beta_2, \beta'_1, \beta'_2$ are non-empty elements of $\mathbb{B}(\mathcal{S}, d, m)$ and, furthermore, $\beta \equiv \beta'$ if and only if $\beta_1 \equiv \beta'_1$ and $\beta_2 \equiv \beta'_2$. The same property allows us to define $\mathbb{R}_\varepsilon(\mathcal{S}, d + 1, m)$ as the set of all blueprints of the form $@_\phi(\gamma_1, \gamma_2)$ where $@_\phi \in S$ and each γ_i is a non-empty element of $\mathbb{R}(\mathcal{S}, d, m)$.

- (2) The second part of the lemma follows by induction on d , using part (1) and the fact that $\mathbb{B}_\varepsilon(\mathcal{S}, 0, 0)$ is empty (hence $\mathbb{B}(\mathcal{S}, d, 0) = \{\emptyset_{\mathbb{B}}\}$ for all d) and if $m \in \mathbb{N}_+$, then $\mathbb{B}_\varepsilon(\mathcal{S}, 0, m)$ is the finite set of all formulas of \mathcal{S} . □

5.2. Shadow of a term

Definition 5.6. Let ϕ be a formula and \mathcal{S}_ϕ be the union of $\text{Sub}(\phi)$ (Definition 4.2) and the set of all $@_\psi$ such that $\psi \in \text{Sub}(\phi)$. For each integer k and each formula ϕ , we let $\mathfrak{R}(\phi, k) = \mathbb{R}(\mathcal{S}_\phi, k \times |\text{Sub}(\phi)|, k)$, where \mathbb{R} is the function introduced in Lemma 5.5 (2).

Definition 5.7. A *shadow* is a finite tree in which each node has arity at most 2 and is labelled with a triple of the form $(\bar{\chi}, \gamma, \psi)$, where $\bar{\chi}$ is a sequence of formulas, γ is a blueprint and ψ is a formula.

We define a ϕ -*shadow* to be any shadow Ξ satisfying the following conditions, where we assume $\Xi(\varepsilon) = (\varepsilon, \emptyset_{\mathbb{B}}, \phi)$, and for each $a \in \text{dom}(\Xi)$, we let k_a be the number of $b < a$ such that the node of Ξ at b is unary, and we let $(\bar{\chi}_a, \gamma_a, \psi_a) = \Xi(a)$:

- $\bar{\chi}_a$ is a sequence of subformulas of ϕ of length at most k_a .
- $\gamma_a \in \mathfrak{R}(\phi, k_a)$.
- $\bar{\chi}_a \in \mathbb{F}(\gamma_a)$.
- ψ_a is a subformula of ϕ .

Definition 5.8. Let M be a locally compact Λ_{NF} -inhabitant of ϕ . For each $a \in \text{dom}(M)$, let:

- $\bar{\chi}_a = \Omega(\text{Free}(M|_a))$;
- α_a be the blueprint of $M|_a$;
- $\gamma_a \in \mathfrak{R}(\phi, |\Lambda(M, a)|)$ be such that $\gamma_a \sqsubseteq_{|\Lambda(M, a)|}^{\max} \alpha_a$;
- ϕ_a be the type of $M|_a$.

We will say the tree Ξ mapping each $a \in \text{dom}(M)$ to $(\bar{\chi}_a, \gamma_a, \phi_a)$ is *the shadow of M* .

Recall that if M is a locally compact Λ_{NF} -inhabitant of ϕ , then for each address a in M , the blueprint α_a of $M|_a$ has relative depth at most $|\Lambda(M, a)| \times |\text{Sub}(\phi)|$. Every maximal $|\Lambda(M, a)|$ -compression of α_a produces a shadow α'_a with the same relative depth and of width at most $|\Lambda(M, a)|$, to which some element of $\mathfrak{R}(\phi, |\Lambda(M, a)|)$ is equivalent, thus the shadow of M is well defined. Note that the choice of γ_a is not required to be unique (though it is, since \mathbb{R} is a selector and we can actually prove that $\gamma \sqsubseteq_m^{\max} \alpha$ and $\gamma' \sqsubseteq_m^{\max} \alpha$ implies $\gamma \equiv \gamma'$, but this property is irrelevant to our discussion); we just assume that *some* γ_a is chosen for each address a in M (see Figure 12).

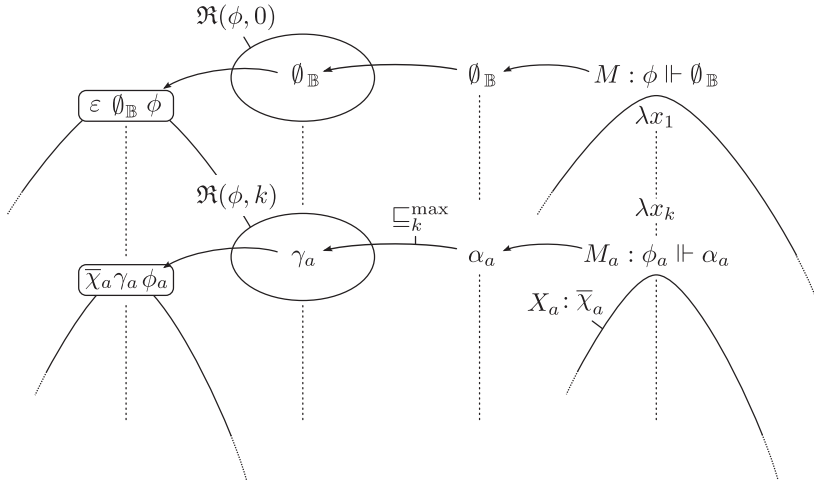


Fig. 12. A locally compact inhabitant and its shadow.

It is obvious that the shadow of M satisfies the first, second and fourth conditions in the definition of ϕ -shadows given above – in the next section, we prove that it also satisfies the third.

5.3. Compact shadows and compact inhabitants

Definition 5.9. A shadow Ξ is *compact* if and only if there are no a, b such that:

- $a < b$;
- the nodes of Ξ at a and b have the same arity;
- $\Xi(a) = (\bar{\lambda}_a, \gamma_a, \psi)$;
- $\Xi(b) = (\bar{\lambda}_b, \gamma_b, \psi)$;
- there exists $\gamma' \uparrow \gamma_b$ such that $\bar{\lambda}_a \in \mathbb{F}(\gamma')$.

This definition should be compared with the definition of compactness for a term (Definition 4.8). With the help of three auxiliary lemmas, we now prove the key lemma of Section 5: if M is a compact inhabitant – *a fortiori* locally compact by Lemma 4.9 – then the shadow of M is a compact ϕ -shadow.

Lemma 5.10. If $\alpha \uparrow \beta \sqsubseteq_1 \beta'$, then there exists α' such that $\alpha \sqsubseteq_1 \alpha' \uparrow \beta'$.

Proof.

- (1) An immediate induction on $|\text{dom}(\beta')|$ shows that if $\alpha = \beta[a \leftarrow \beta|_b]$ and $\beta \equiv \beta'$, then there exist a', b' such that $a' < b'$ and $\alpha \equiv \alpha' = \beta'[a' \leftarrow \beta'|_{b'}]$. As a consequence, an immediate induction on the length of the derivation of $\alpha \uparrow \beta$ shows that the lemma holds if $\beta \equiv \beta'$.

(2) Another induction on $|\text{dom}(\beta')|$ shows that if $\alpha \uparrow \beta \curvearrowright \beta'$, then there exists α' such that $\alpha \curvearrowright \alpha' \uparrow \beta'$. The only non-trivial case is

$$\begin{aligned} \alpha &= *_{(a_1)}(\alpha_1) \\ \beta &= *_{(a_1)}(\beta_1) && \text{with } \alpha_1 \uparrow \beta_1 \\ \beta' &= *_{(a_1, a_2)}(\beta_1, \beta_2) && \text{with } \beta_1 \equiv \beta_2. \end{aligned}$$

Since $\alpha_1 \uparrow \beta_1 \equiv \beta_2$, by (1) there exists α_2 such that $\alpha_1 \equiv \alpha_2 \uparrow \beta_2$. Hence

$$\alpha = *_{(a_1)}(\alpha_1) \curvearrowright *_{(a_1, a_2)}(\alpha_1, \alpha_2) \uparrow *_{(a_1, a_2)}(\beta_1, \beta_2) = \beta'.$$

(3) Using (1) and (2), the lemma follows by induction on the length of an arbitrary sequence $(\beta_0, \dots, \beta_n)$ such that $\beta_0 = \beta$, $\beta_n = \beta'$ and $\beta_{i-1} \equiv \beta_i$ or $\beta_{i-1} \curvearrowright \beta_i$ for each $i \in [1, \dots, n]$. □

Lemma 5.11. If $\alpha \sqsubseteq_1 \beta$, then $\mathbb{F}(\alpha) \subseteq \mathbb{F}(\beta)$.

Proof. The proof is by induction on $|\text{dom}(\beta)|$. Since $\gamma \equiv \gamma'$ implies $\mathbb{F}(\gamma) = \mathbb{F}(\gamma')$ and $|\text{dom}(\gamma)| = |\text{dom}(\gamma')|$, it is enough to consider the case where α is a 1-compression of β . The case $\alpha = *_{(a_1)}(\alpha_1)$ and $\beta = *_{(a_1, a_2)}(\alpha_1, \alpha_2)$ is clear, and the remaining cases follow easily from the induction hypothesis. □

Lemma 5.12. If $\alpha \sqsubseteq_m \beta$, the set of all elements of $\mathbb{F}(\beta)$ of length at most m is a subset of $\mathbb{F}(\alpha)$.

Proof. The proof is by induction on $|\text{dom}(\beta)|$. Again, we will only examine the case $\alpha \curvearrowright_m \beta$. The proposition is trivially true if $m = 0$. So we assume $m > 0$. The only non-trivial case is

$$\begin{aligned} \alpha &\equiv *_{\bar{a}}(\gamma_1, \dots, \gamma_m) \\ \beta &\equiv *_{\bar{a}}(\gamma_1, \dots, \gamma_m, \gamma_{m+1}) \end{aligned}$$

with $\gamma_i \equiv \gamma$ for all i . Let $\Phi = \mathbb{F}(\gamma)$. For each integer k , let $\Phi^{(k)} = \otimes(\Phi_1, \dots, \Phi_k)$ where $\Phi_i = \mathbb{F}(\gamma)$ for each i . Let $\bar{\phi} = (\phi_1, \dots, \phi_p) \in \mathbb{F}(\beta)$ be such that $p \leq m$. We have to prove that $\bar{\phi} \in \mathbb{F}(\alpha)$. For each $J \subseteq \{1, \dots, p\}$, let (j_1, \dots, j_q) be the strictly increasing enumeration of all elements of J and let $f(J) = (\phi_{j_1}, \dots, \phi_{j_q})$. We have $\bar{\phi} \in \mathbb{F}(\beta) = \Phi^{(m+1)}$, so there exist J_1, \dots, J_{m+1} such that

$$J_1 \cup \dots \cup J_{m+1} = \{1, \dots, p\},$$

and $f(J_i) \in \mathbb{F}(\gamma)$ for each $i \in \{1, \dots, m+1\}$. For each $j \in \{1, \dots, p\}$, let k_j be any element of $\{1, \dots, m+1\}$ such that $j \in J_{k_j}$. Then $J_{k_1} \cup \dots \cup J_{k_p} = \{1, \dots, p\}$, so

$$\bar{\phi} \in \otimes(\{f(J_{k_1})\}, \dots, \{f(J_{k_p})\}) \subseteq \Phi^{(p)} \subseteq \Phi^{(m)} = \mathbb{F}(\alpha).$$

□

Lemma 5.13. Let M be a locally compact Λ_{NF} -inhabitant of ϕ .

- (1) The shadow of M is a ϕ -shadow.
- (2) If M is compact, this shadow is also compact.

Proof.

- (1) For each address a in M , the sequence $\bar{\chi}_a = \Omega(\text{Free}(M|_a))$ is a subsequence of $\Omega(\Lambda(M, a))$, so the first part of the lemma follows from the definition of the shadow of M , Lemma 2.5, Lemma 3.15 (3) and Lemma 5.12.
- (2) Let Ξ be shadow of M and assume Ξ is not compact. There exist $a, b \in \text{dom}(\Xi) = \text{dom}(M)$ such that:

- $\Xi(a) = (\bar{\chi}_a, \gamma_a, \psi)$.
- $\Xi(b) = (\bar{\chi}_b, \gamma_b, \psi)$.
- The nodes at a and b in Ξ have the same arity.
- There exists $\gamma' \uparrow \gamma_b$ such that $\bar{\chi}_a \in \mathbf{F}(\gamma')$.

We have $M|_a, M|_b$ of the same kind. Let α_a, α_b be the blueprints of $M|_a, M|_b$.

Since $\gamma_b \sqsubseteq_{|\Lambda(M, a, b)|}^{\max} \alpha_b$, we have $\gamma' \uparrow \gamma_b \sqsubseteq_1 \alpha_b$. By Lemma 5.10, there exists α' such that $\gamma' \sqsubseteq_1 \alpha' \uparrow \alpha_b$. By Lemma 5.11, we have $\bar{\chi}_a \in \mathbf{F}(\gamma') \subseteq \mathbf{F}(\alpha')$, so M is not compact. \square

6. Finiteness of the set of compact ϕ -shadows

Our final step is to prove that for each formula ϕ , the set of all compact ϕ -shadows is a finite set effectively computable from ϕ .

In Definition 6.1, we will introduce a final binary relation \subseteq on blueprints. The key lemma of this section (Lemma 6.14) shows that whenever $\mathcal{S} \subset \mathfrak{S}$ is a finite set (in particular, when \mathcal{S} is the set of all subformulas of ϕ and all $@$'s tagged with a subformula of ϕ), the relation \subseteq is an almost full relation (Bezem *et al.* 2003) on the set of all \mathcal{S} -blueprints: for every infinite sequence $\gamma_1, \gamma_2, \dots$ over $\mathbf{B}(\mathcal{S})$, there exists i, j such that $i < j$ and $\gamma_i \subseteq \gamma_j$. We will prove this result with the help of Mellies' Axiomatic Kruskal Theorem (Mellies 1998). The finiteness of the set of compact ϕ -shadows follows from this key lemma with the help of König's Lemma (Lemma 6.15). The ability to compute these shadows follows directly from their definition.

By Lemma 5.13, another consequence of this result is the finiteness for each ϕ of the set of all compact Λ_{NF} -inhabitants of ϕ , although our decision method is based on the computation of *shadows* of compact terms rather than a direct computation of those terms. It is worth mentioning that the proof of Theorem 6.13 is non-constructive and that it gives no information about the complexity of our proof-search method – this question might itself become another open problem.

6.1. Almost full relations and Higman's Theorem

Definition 6.1. We let \subseteq be the relation on blueprints defined by $\alpha \subseteq \beta$ if and only if for all $\bar{\chi} \in \mathbf{F}(\alpha)$, there exists $\gamma \uparrow \beta$ such that $\bar{\chi} \in \mathbf{F}(\gamma)$.

Definition 6.2. Let \mathcal{U} be an arbitrary set. An *almost full relation (AFR)* on \mathcal{U} is a binary relation \ll such that for every infinite sequence $(u_i)_{i \in \mathbb{N}}$ over \mathcal{U} , there exist i, j such that $i < j$ and $u_i \ll u_j$.

The main aim of Section 6 is to prove the final key lemma of the paper, from which we will easily infer the decidability of Λ_{NF} -inhabitation: for each finite $\mathcal{S} \subseteq \mathfrak{S}$, the relation \subseteq is an AFR on $\mathbb{B}(\mathcal{S})$.

Proposition 6.3.

- (1) If \ll and \ll' are AFRs on \mathcal{U} , then $\ll \cap \ll'$ is an AFR on \mathcal{U} .
- (2) Suppose $\ll_{\mathcal{U}}$ is an AFR on \mathcal{U} and $\ll_{\mathcal{V}}$ is an AFR on \mathcal{V} . Let $\ll_{\mathcal{U} \times \mathcal{V}}$ be the relation defined by $(U, V) \ll_{\mathcal{U} \times \mathcal{V}} (U', V')$ if and only if $U \ll_{\mathcal{U}} U'$ and $V \ll_{\mathcal{V}} V'$. Then $\ll_{\mathcal{U} \times \mathcal{V}}$ is an AFR on $\mathcal{U} \times \mathcal{V}$.

Proof. Both results appear in Melliès (1998) in Step 4 of the proof of Theorem 1 (page 523), as a corollary of Lemma 4 (page 520). □

Definition 6.4. Let \mathcal{U} be a set and \ll be a binary relation. We write $\mathbf{S}(\mathcal{U})$ to denote the set of all finite sequences over \mathcal{U} . The relation $\ll_{\mathbf{S}}$ induced by \ll on $\mathbf{S}(\mathcal{U})$ is defined by $(U_1, \dots, U_n) \ll_{\mathbf{S}} (V_1, \dots, V_m)$ if and only if there exists a strictly monotone function $\eta : \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ such that $U_i \ll V_{\eta(i)}$ for each $i \in \{1, \dots, n\}$.

Theorem 6.5 (Higman). If \ll is an AFR on \mathcal{U} , then $\ll_{\mathbf{S}}$ is an AFR on $\mathbf{S}(\mathcal{U})$.

Proof. See Higman (1952), Kruskal (1972) and Melliès (1998). □

6.2. From rooted to unrooted blueprints

Melliès’ Axiomatic Kruskal Theorem allows us to conclude that a relation is an AFR (a ‘well binary relation’ in Melliès (1998)) if it satisfies a set of five properties or ‘axioms’ (there were six axioms in the original version of the theorem – see Melliès’ remarks at the end of his proof explaining why five axioms suffice). The details of those axioms will be given in Section 6.3.

Four of those five axioms are relatively easy to check, but the other one is more problematic. This rather technical section is entirely devoted to the proof of Lemma 6.11, which will ensure that this final axiom is satisfied. To do this, we want to prove the following proposition:

- Let \mathcal{S} be a finite subset of \mathfrak{S} and \mathcal{B}_e be a subset of $\mathbb{B}_e(\mathcal{S})$.
- Let $\mathcal{B} = \{*_a(\beta_1, \dots, \beta_n) \mid \forall i \in [1, \dots, n], \beta_i \in \mathcal{B}_e\}$.
- If \subseteq is an AFR on \mathcal{B}_e , then \subseteq is an AFR on \mathcal{B} .

Recall that $\mathbb{B}_e(\mathcal{S})$ stands for the set of all rooted \mathcal{S} -blueprints. We want to be able to extend the property that \subseteq is an AFR on a given set of rooted blueprints to the set of all blueprints that have those rooted blueprints at their minimal addresses.

Higman’s Theorem suffices to show that $\subseteq_{\mathbf{S}}$ (Definition 6.4) is an AFR on the set of finite sequences over \mathcal{B}_e . However, if we consider an infinite sequence $(\beta_i)_{i \in \mathbb{N}}$ over \mathcal{B} and transform each $\beta_i = *_a(\beta_1^i, \dots, \beta_{n_i}^i)$ where $\beta_1^i, \dots, \beta_{n_i}^i \in \mathcal{B}_e$ into $\sigma(\beta_i) = (\beta_1^i, \dots, \beta_{n_i}^i)$, the theorem will only provide two integers i, j and a strictly monotone function η such that $i < j$ and $\beta_k^i \subseteq \beta_{\eta(k)}^j$ for each $k \in \{1, \dots, n_i\}$. This is sufficient to ensure that

$$\beta_i = *_a(\beta_1^i, \dots, \beta_{n_i}^i) \subseteq *_b(\beta_{\eta(1)}^j, \dots, \beta_{\eta(n_i)}^j),$$

but not $\beta_i \subseteq \beta_j$ in general.

To bypass this difficulty, we show how for each blueprint $\beta \in \mathbf{B}(\mathcal{S})$, we can extract from the set of all vertical compressions of β a complete set of ‘followers’ of β of minimal size (Lemma 6.7). This set $\{\alpha_1, \dots, \alpha_p\}$ has the property that for each $\bar{\phi} \in \mathbf{IF}(\beta)$, there exists at least one α_i such that $\mathbf{IF}(\alpha_i)$ contains a *subsequence* of $\bar{\phi}$ – but not necessarily $\bar{\phi}$ itself. The relative depth of each α_i does not depend on the relative depth of β , but only on \mathcal{S} : it is at most $\sum_{i=1}^{1+|\mathcal{S}^\circledast|} i$, where \mathcal{S}^\circledast is the set of all binary symbols in \mathcal{S} . The lemma is proved in four steps:

- (1) We first observe that the set of all $\alpha \uparrow \beta$ with relative depth at most $\sum_{i=1}^{1+|\mathcal{S}^\circledast|} i$ is a complete set of followers. If we consider the set of all γ such that $\gamma \sqsubseteq_1^{\max} \alpha$ for at least one such α , we obtain a (possibly infinite) set closed under \equiv and finite up to \equiv . We call it the set of \mathcal{S} -residuals of β .
- (2) We prove that the set of \mathcal{S} -residuals of β is a complete set of followers of β in the same sense, that is, for each $\bar{\phi} \in \mathbf{IF}(\beta)$, there exists an \mathcal{S} -residual γ of β such that $\mathbf{IF}(\gamma)$ contains a subsequence of $\bar{\phi}$ (Lemma 6.9).
- (3) We prove that if

$$\begin{aligned} \beta &= *_a(\beta_1, \dots, \beta_n) \\ \beta' &= *_b(\beta'_1, \dots, \beta'_n, \beta'_{n+1}, \dots, \beta'_{n+k}) \end{aligned}$$

are such that $\beta_i \subseteq \beta'_i$ for each $i \in [1, \dots, n]$, and if, furthermore, β, β' have the same set of \mathcal{S} -residuals, then $\beta \subseteq \beta'$ (Lemma 6.10).

- (4) The final step is the proof of the lemma itself. The set of \mathcal{S} -residuals is finite up to \equiv (Lemma 5.5), so there are only a finite number of possible values for the set of residuals of each \mathcal{S} -blueprint. As a consequence, it is always possible to extract from an infinite sequence over \mathcal{B} an infinite sequence of blueprints with the same set of residuals. The conclusion then follows from the third step and Higman’s Theorem.

Definition 6.6. For every $\mathcal{S} \subseteq \mathfrak{S}$, we write \mathcal{S}^\circledast to denote the set of all binary symbols in \mathcal{S} .

Lemma 6.7. Let \mathcal{S} be a finite subset of \mathfrak{S} . For all $\beta \in \mathbf{B}(\mathcal{S})$ and all $\bar{\psi} \in \mathbf{IF}(\beta)$, there exists α with relative depth at most $\sum_{i=1}^{1+|\mathcal{S}^\circledast|} i$ such that $\alpha \uparrow \beta$ and such that $\mathbf{IF}(\alpha)$ contains a subsequence of $\bar{\psi}$.

Proof. We define an \mathcal{S} -linearisation to be any pair $(\gamma, \bar{\chi})$ such that $\gamma \in \mathbf{B}(\mathcal{S})$ and $\bar{\chi} \in \mathbf{IF}(\gamma)$. We say a *starting address* for $(\gamma, \bar{\chi})$ is any address b for which there exist ϕ, γ' such that $\gamma \triangleright_\phi^b \gamma'$ and $\bar{\chi} \in \mathfrak{C}(\mathbf{IF}(\gamma'), (\phi))$. We define the *path to b in γ* to be the maximal sequence $(b_1, \dots, b_n, b_{n+1})$ over $\text{dom}(\gamma)$ such that $b_1 < \dots < b_n < b_{n+1} = b$.

Given an arbitrary \mathcal{S} -linearisation $(\beta, \bar{\psi})$, we use induction on $|\text{dom}(\beta)|$ to prove the following properties simultaneously:

- (1) There exists an \mathcal{S} -linearisation $(\gamma, \bar{\chi})$ such that:
 - (a) $\gamma \uparrow \beta$ and $\bar{\chi}$ is a subsequence of $\bar{\psi}$.
 - (b) γ has relative depth at most $1 + \sum_{i=1}^{|\mathcal{S}^\circledast|} i$.
- (2) There exists an \mathcal{S} -linearisation $(\alpha, \bar{\phi})$ such that:

- (a) $\alpha \uparrow \beta$, $\bar{\phi}$ is a subsequence of $\bar{\psi}$, and if $\psi \neq \varepsilon$, then the last elements of $\bar{\phi}, \bar{\psi}$ are equal.
- (b) For each starting address b for $(\alpha, \bar{\phi})$ and for $(b_1, \dots, b_n, b_{n+1})$ equal to the path to b in α , the values $\alpha(b_1), \dots, \alpha(b_n)$ are pairwise distinct.
- (c) For all c incomparable with each starting address for $(\alpha, \bar{\phi})$, the relative depth of $(\alpha|_c)$ is at most $1 + \sum_{i=1}^{|\mathcal{S}_{@}|} i$.

Note that the conjunction of (2b) and (2c) implies that every address d in α has relative depth at most

$$|\mathcal{S}_{@}| + 1 + \sum_{i=1}^{|\mathcal{S}_{@}|} i = \sum_{i=1}^{1+|\mathcal{S}_{@}|} i.$$

Indeed, suppose d has maximal relative depth and not a starting address for $(\alpha, \bar{\phi})$. Then d must be incomparable with each starting address for $(\alpha, \bar{\phi})$. Let e be the shortest prefix of d in $\text{dom}(\alpha)$ that is incomparable with each starting address for $(\alpha, \bar{\phi})$. The address e has relative depth at most $|\mathcal{S}_{@}|$ in α , since otherwise there would exist in $\text{dom}(\alpha)$ an address $f < e$ with relative depth $|\mathcal{S}_{@}|$ and a starting address for $(\alpha, \bar{\phi})$ of the form $f \cdot f'$, with relative depth strictly greater than $|\mathcal{S}_{@}|$, which would give a contradiction. Moreover, the relative depth of d is the sum of the relative depth of e in α and the relative depth of $\alpha|_e$.

The case $\beta = \emptyset_{\mathbb{B}}$ is immediate, and if $\beta = *_a(\beta_1, \dots, \beta_n)$, $i \neq j$ and $\beta_i, \beta_j \neq \emptyset_{\mathbb{B}}$, the conclusion follows easily from the induction hypothesis. So we suppose $\beta = @_{\psi}(\beta_1, \beta_2)$.

- (1) Let d be an address of maximal length in $\beta^{-1}(@_{\psi})$, and let $\delta = @_{\psi}(\delta_1, \delta_2) = \beta|_d$. By assumption, ε is the only element of $\delta^{-1}(@_{\psi})$. As $\bar{\psi} \in \mathbf{IF}(\beta)$, there exist

$$\begin{aligned} \bar{\psi}_0 &\in \mathbf{IF}(\delta) \\ \bar{\psi}_1 &\in \mathbf{IF}(\delta_1) \\ \bar{\psi}_2 &\in \mathbf{IF}(\delta_2) \end{aligned}$$

such that $\bar{\psi}_0$ is a subsequence $\bar{\psi}$ and $\bar{\psi}_0 \in \circ(\{\bar{\psi}_1\}, \{\bar{\psi}_2\})$. By the induction hypothesis, there exists an $(\mathcal{S} - \{ @_{\psi} \})$ -linearisation $(\gamma_1, \bar{\chi}_1)$ satisfying conditions (1a) and (1b) with respect to $(\delta_1, \bar{\psi}_1)$, and an $(\mathcal{S} - \{ @_{\psi} \})$ -linearisation $(\gamma_2, \bar{\chi}_2)$ satisfying conditions (2a), (2b) and (2c) with respect to $(\delta_2, \bar{\psi}_2)$. Let $\gamma = @_{\psi}(\gamma_1, \gamma_2)$. We have $\gamma \uparrow \delta$ and $\beta(\varepsilon) = \delta(\varepsilon) = \gamma(\varepsilon)$, so $\gamma \uparrow \beta$. The blueprint γ_1 has relative depth at most

$$1 + \sum_{i=1}^{|\mathcal{S}_{@}|-1} i \leq \sum_{i=1}^{|\mathcal{S}_{@}|} i.$$

The blueprint γ_2 has relative depth at most

$$|\mathcal{S}_{@}| + \sum_{i=1}^{|\mathcal{S}_{@}|-1} i = \sum_{i=1}^{|\mathcal{S}_{@}|} i.$$

Therefore, γ has relative depth at most

$$1 + \sum_{i=1}^{|\mathcal{S}_{@}|} i.$$

Now $\bar{\chi}_2$ is a subsequence of $\bar{\psi}_2$ with the same final element, so there exists in

$$\circ(\{\bar{\chi}_1\}, \{\bar{\chi}_2\}) \subseteq \mathbf{IF}(@_{\psi}(\gamma_1, \gamma_2))$$

a subsequence $\bar{\chi}$ of $\bar{\psi}_0$. Thus $(\gamma, \bar{\chi})$ satisfies (1a) and (1b) with respect to $(\beta, \bar{\psi})$.

(2) As $\bar{\varphi} \in \mathbf{IF}(\beta)$, there exist

$$\begin{aligned} \bar{\varphi}_1 &\in \mathbf{IF}(\beta_1) \\ \bar{\varphi}_2 &\in \mathbf{IF}(\beta_2) \end{aligned}$$

such that

$$\bar{\varphi} \in \odot(\{\bar{\varphi}_1\}, \{\bar{\varphi}_2\}).$$

By the induction hypothesis, there exists an \mathcal{S} -linearisation $(\alpha_1, \bar{\phi}_1)$ satisfying conditions (1a) and (1b) with respect to $(\beta_1, \bar{\varphi}_1)$, and an \mathcal{S} -linearisation $(\alpha_2, \bar{\phi}_2)$ satisfying conditions (2a), (2b) and (2c) with respect to $(\beta_2, \bar{\varphi}_2)$.

Let $\alpha_0 = @_{\varphi}(\alpha_1, \alpha_2)$. We have $\alpha_0 \uparrow \beta$. The final elements of $\bar{\phi}_2$ and $\bar{\varphi}_2$ are equal and

$$\odot(\{\bar{\phi}_1\}, \{\bar{\phi}_2\}) \subseteq \mathbf{IF}(\alpha_0).$$

Hence, there exists in $\mathbf{IF}(\alpha_0)$ a subsequence $\bar{\phi}_0$ of $\bar{\varphi}$ with the same final element as $\bar{\varphi}$. Thus $(\alpha_0, \bar{\phi}_0)$ satisfies (2a).

For all c incomparable with each starting address for $(\alpha_0, \bar{\phi}_0)$, either $c = (1) \cdot c'$ and $c' \in \text{dom}(\alpha_1)$, or $c = (2) \cdot c''$ and $c'' \in \text{dom}(\alpha_2)$ is incomparable with each starting address in α_2 . As a consequence, the choice of α_1, α_2 ensures that $(\alpha_0, \bar{\phi}_0)$ satisfies (2c). If $(\alpha_0, \bar{\phi}_0)$ satisfies (2b), we may take $(\alpha, \bar{\phi}) = (\alpha_0, \bar{\phi}_0)$. Otherwise, some starting address b for $(\alpha_0, \bar{\phi}_0)$ does not satisfy condition (2b). Let $(b_1, \dots, b_n, b_{n+1})$ be the path to b in α . We have $b_1 = \varepsilon$, and for each $i > 0$, there exists d_i such that $b_i = (2) \cdot d_i$. The sequence (d_2, \dots, d_{n+1}) is then a path to $d = d_{n+1}$ in α_2 , and d is a starting address for $(\alpha_2, \bar{\phi}_2)$. The values $\alpha_2(d_2), \dots, \alpha_2(d_n)$ are pairwise distinct, so there must exist $i > 1$ such that $\alpha(b_i) = @_{\varphi}$. Since b_i is in the path to b , there exists in $\mathbf{IF}(\alpha_{2|d_i})$ a subsequence $\bar{\phi}'_0$ of $\bar{\phi}_0$ with the same last element as $\bar{\phi}_0$. For $\alpha'_0 = \alpha_0[\varepsilon \leftarrow \alpha_{2|d_i}]$, we have $\alpha'_0 \uparrow \beta$ and $\bar{\phi}'_0 \in \mathbf{IF}(\alpha'_0)$ and the final elements of $\bar{\phi}'_0, \bar{\phi}_0$ and $\bar{\varphi}$ are equal. By the induction hypothesis, there exists an \mathcal{S} -linearisation $(\alpha, \bar{\phi})$ satisfying (2a), (2b) and (2c) with respect to $(\alpha'_0, \bar{\phi}'_0)$. The pair $(\alpha, \bar{\phi})$ also satisfies these conditions with respect to $(\beta, \bar{\varphi})$. \square

Definition 6.8. Let \mathcal{S} be a finite subset of \mathfrak{S} . For all $\beta \in \mathbf{IB}(\mathcal{S})$ and all $\alpha \uparrow \beta$ with relative depth at most $\sum_{i=1}^{1+|\mathcal{S}@\!|} i$, we define an \mathcal{S} -residual of β to be any α_0 such that $\alpha_0 \sqsubseteq_1^{\max} \alpha$.

Note that the set of \mathcal{S} -residuals of β is $\{\emptyset_{\mathbf{B}}\}$ if $\beta = \emptyset_{\mathbf{B}}$. Otherwise, it is an infinite set – even if $\beta = \phi$, the set of residuals of β is the \equiv -equivalence class of β and contains all blueprints of the form $*_a(\phi)$ (recall that \equiv is a subset of \sqsubseteq_1 – see Definition 5.3).

Lemma 6.9. Let \mathcal{S} be a finite subset of \mathfrak{S} . For all $\beta \in \mathbf{IB}(\mathcal{S})$ and all $\bar{\varphi} \in \mathbf{IF}(\beta)$, there exists an \mathcal{S} -residual α_0 of β such that $\mathbf{IF}(\alpha_0)$ contains a subsequence of $\bar{\varphi}$.

Proof.

(1) Let γ, δ be arbitrary blueprints. Suppose $\gamma \curvearrowright_1 \delta$. We prove by induction on δ that for all $\bar{\phi} \in \mathbf{IF}(\delta)$, there exists in $\mathbf{IF}(\gamma)$ a subsequence of $\bar{\phi}$. In order to deal with the case $\delta = @_{\phi}(\delta_1, \delta_2)$, we need to prove a slightly more precise property: for all $\bar{\phi} \in \mathbf{IF}(\delta)$, there exists in $\mathbf{IF}(\gamma)$ a subsequence $\bar{\varphi}$ of $\bar{\phi}$ such that the final elements of $\bar{\phi}$ and $\bar{\varphi}$

are equal. The base case is $\delta = *_{(a_1, a_2)}(\gamma_1, \gamma_2)$, $\gamma_1 \equiv \gamma_2$ and $\gamma = *_{a_1}(\gamma_1)$, and this case is clear. Other cases follow easily from the induction hypothesis.

- (2) We can now prove the lemma. By Lemma 6.7 and the definition of an \mathcal{S} -residual, there exist α_0 and α such that $\alpha_0 \sqsubseteq_1 \alpha \uparrow \beta$, $\mathbf{IF}(\alpha)$ contains a subsequence of $\bar{\varphi}$ and α_0 is an \mathcal{S} -residual. It then follows from (1) that $\mathbf{IF}(\alpha_0)$ contains a subsequence of $\bar{\varphi}$. \square

Lemma 6.10. Let \mathcal{S} be a finite subset of \mathfrak{S} . Suppose:

- $\beta = *_{\bar{a}}(\beta_1, \dots, \beta_n) \in \mathbf{B}(\mathcal{S})$.
- $\beta' = *_{\bar{b}}(\beta'_1, \dots, \beta'_n, \beta'_{n+1}, \dots, \beta'_{n+k}) \in \mathbf{B}(\mathcal{S})$.
- $\beta_i \subseteq \beta'_i$ for each $i \in \{1, \dots, n\}$.
- The sets of \mathcal{S} -residuals of β and β' are equal.

Then $\beta \subseteq \beta'$.

Proof. Let $\bar{\varphi} \in \mathbf{IF}(\beta)$. So there exists for each $i \in [1, \dots, n]$ a sequence $\bar{\varphi}_i \in \mathbf{IF}(\beta_i)$ such that $\bar{\varphi} \in \otimes(\{\bar{\varphi}_1\}, \dots, \{\bar{\varphi}_n\})$. By assumption, there exists for each $i \in [1, \dots, n]$ an $\alpha_i \uparrow \beta'_i$ such that $\bar{\varphi}_i \in \mathbf{IF}(\alpha_i)$. As a consequence, $\bar{\varphi} \in \mathbf{IF}(*(\alpha_1, \dots, \alpha_n))$.

By Lemma 6.9, there exists an \mathcal{S} -residual α_0 of β such that $\mathbf{IF}(\alpha_0)$ contains a subsequence $\bar{\phi}$ of $\bar{\varphi}$. By assumption, α_0 is also an \mathcal{S} -residual of β' , so there exist $\alpha'_1, \dots, \alpha'_{n+k}$ and \bar{b} such that

$$\alpha_0 \sqsubseteq_1 *_{\bar{b}}(\alpha'_1, \dots, \alpha'_{n+k}) \uparrow \beta'.$$

By Lemma 5.11, we have

$$\bar{\phi} \in \mathbf{IF}(*_{\bar{b}}(\alpha'_1, \dots, \alpha'_{n+k})).$$

Hence, for each $i \in [1, \dots, n+k]$, there exists in $\mathbf{IF}(\alpha'_i)$ a subsequence of $\bar{\phi}$, which is also a subsequence of $\bar{\varphi}$. Let

$$\alpha = *_{\bar{a}}(\alpha_1, \dots, \alpha_n, \alpha'_{n+1}, \dots, \alpha'_{n+k}).$$

Then $\alpha \uparrow \beta'$, $\bar{\varphi} \in \mathbf{IF}(*(\alpha_1, \dots, \alpha_n))$, and for each $j \in [1, \dots, k]$, there exists in $\mathbf{IF}(\alpha'_{n+j})$ a subsequence of $\bar{\varphi}$. Hence, $\bar{\varphi} \in \mathbf{IF}(\alpha)$. \square

Lemma 6.11. Let \mathcal{S} be a finite subset of \mathfrak{S} and \mathcal{B}_ε be a subset of $\mathbf{B}_\varepsilon(\mathcal{S})$. Let

$$\mathcal{B} = \{ *_{\bar{a}}(\beta_1, \dots, \beta_n) \mid \forall i \in [1, \dots, n], \beta_i \in \mathcal{B}_\varepsilon \}.$$

If \subseteq is an AFR on \mathcal{B}_ε , then \subseteq is an AFR on \mathcal{B} .

Proof. Let $\mathcal{R} = \mathbf{B}(\mathcal{S}, \sum_{i=1}^{|\mathcal{S}|} i, 1)$ (see Definition 5.4). For each $\beta \in \mathcal{B}$, let $\rho(\beta)$ be the set of \mathcal{S} -residuals of β . We have $\rho(\beta) \subseteq \mathcal{R}$. Moreover, $\rho(\beta)$ is closed under \equiv (since \equiv is a subset of \sqsubseteq_1 – see Definition 5.3), that is, $\rho(\beta)$ is a union of the elements of a subset of \mathcal{R}/\equiv . By Lemma 5.5 (1), the latter is a finite set, so $\{\rho(\beta) \mid \beta \in \mathcal{B}\}$ is a finite set.

For each $\beta = *_{\bar{a}}(\beta_1, \dots, \beta_n) \in \mathcal{B}$, where \bar{a} is increasing with respect to the lexicographic ordering of addresses and $\beta_1, \dots, \beta_n \in \mathcal{B}_\varepsilon$, let $\sigma(\beta) = (\beta_1, \dots, \beta_n)$ – recall that we can take $\bar{a} = \varepsilon$, $n = 0$ if $\beta = \emptyset_{\mathbf{B}}$, and $\bar{a} = (\varepsilon)$, $n = 1$ if β is a rooted blueprint. Since $\{\rho(\beta) \mid \beta \in \mathcal{B}\}$ is a finite set, every infinite sequence over \mathcal{B} contains an infinite subsequence of blueprints with the same set of \mathcal{S} -residuals. By assumption, \subseteq is an AFR on \mathcal{B}_ε , so, by Theorem 6.5, $\subseteq_{\mathfrak{S}}$ is an AFR on $\{\sigma(\beta) \mid \beta \in \mathcal{B}\}$.

Thus, for every infinite sequence $(\beta_i)_{i \in \mathbb{N}}$ over \mathcal{B} , there exist i, j such that $i < j$, $\sigma(\beta_i) \subseteq_{\mathbb{S}} \sigma(\beta_j)$, and β_i and β_j have the same set of residuals. For

$$\begin{aligned}\sigma(\beta_i) &= (\beta_1^i, \dots, \beta_n^i) \\ \sigma(\beta_j) &= (\beta_1^j, \dots, \beta_{n+k}^j),\end{aligned}$$

there exists a subsequence $(\beta_{l_1}^i, \dots, \beta_{l_n}^i)$ of $\sigma(\beta_j)$ such that

$$\beta_1^i \subseteq \beta_{l_1}^i, \dots, \beta_n^i \subseteq \beta_{l_n}^i.$$

There also exist l_{n+1}, \dots, l_{n+k} and two sequences \bar{a} and \bar{b} such that

$$\begin{aligned}\beta_i &= *_{\bar{a}}(\beta_1^i, \dots, \beta_n^i) \\ \beta_j &= *_{\bar{b}}(\beta_1^j, \dots, \beta_{l_n}^j, \beta_{l_{n+1}}^j, \dots, \beta_{l_{n+k}}^j).\end{aligned}$$

By Lemma 6.10, we have $\beta_i \subseteq \beta_j$. □

6.3. Axiomatic Kruskal Theorem and main key lemma

The following definition is borrowed from Melliès (1998).

Definition 6.12. An *abstract decomposition system* is an 8-tuple

$$(\mathcal{T}, \mathcal{L}, \mathcal{V}, \leq_{\mathcal{T}}, \leq_{\mathcal{L}}, \leq_{\mathcal{V}}, \xrightarrow{\cdot}, \vdash)$$

where:

- \mathcal{T} is a set of *terms* noted t, u, \dots equipped with a binary relation $\leq_{\mathcal{T}}$.
- \mathcal{L} is a set of *labels* noted f, g, \dots equipped with a binary relation $\leq_{\mathcal{L}}$.
- \mathcal{V} is a set of *vectors* noted T, U, \dots equipped with a binary relation $\leq_{\mathcal{V}}$.
- $\xrightarrow{\cdot}$ is a relation on $\mathcal{T} \times \mathcal{L} \times \mathcal{V}$, for example, $t \xrightarrow{f} T$.
- \vdash is a relation on $\mathcal{V} \times \mathcal{T}$, for example, $T \vdash t$.

For each such system, we let $\triangleright_{\mathcal{T}}$ be the binary relation on \mathcal{T} defined by

$$t \triangleright_{\mathcal{T}} u \iff \exists (f, T) \in \mathcal{L} \times \mathcal{V}, \quad t \xrightarrow{f} T \vdash u.$$

An *elementary term* t is a term minimal with respect to $\triangleright_{\mathcal{T}}$, that is, a term for which there exists no u such that $t \triangleright_{\mathcal{T}} u$.

Theorem 6.13 (Melliès). Suppose $(\mathcal{T}, \mathcal{L}, \mathcal{V}, \leq_{\mathcal{T}}, \leq_{\mathcal{L}}, \leq_{\mathcal{V}}, \xrightarrow{\cdot}, \vdash)$ satisfies the following properties:

Axiom I: There is no infinite chain $t_1 \triangleright_{\mathcal{T}} t_2 \triangleright_{\mathcal{T}} \dots$.

Axiom II: The relation $\leq_{\mathcal{T}}$ is an AFR on the set of elementary terms.

Axiom III: For all t, u, u' , if $t \leq_{\mathcal{T}} u'$ and $u \triangleright_{\mathcal{T}} u'$, then $t \leq_{\mathcal{T}} u$.

Axiom IV-bis: For all t, u, f, g, T, U , if $t \xrightarrow{f} T$ and $u \xrightarrow{g} U$ and $f \leq_{\mathcal{L}} g$ and $T \leq_{\mathcal{V}} U$, then $t \leq_{\mathcal{T}} u$.

Axiom V: For all $\mathcal{W} \subseteq \mathcal{V}$, for $\mathcal{W}_{\vdash} = \{t \in \mathcal{T} \mid \exists T \in \mathcal{W}, T \vdash t\}$, if $\leq_{\mathcal{T}}$ is an AFR on \mathcal{W}_{\vdash} , then $\leq_{\mathcal{V}}$ is an AFR on \mathcal{W} .

If, furthermore, $\leq_{\mathcal{L}}$ is an AFR on \mathcal{L} , then $\leq_{\mathcal{T}}$ is an AFR on \mathcal{T} .

Proof. See Mellies (1998)[†]. □

Lemma 6.14. For each finite $\mathcal{S} \subseteq \mathfrak{S}$, the relation \subseteq is an AFR on $\mathbf{B}(\mathcal{S})$.

Proof. According to Lemma 6.11, it is sufficient to prove that \subseteq is an AFR on $\mathbf{B}_e(\mathcal{S})$. Let $(\mathcal{T}, \mathcal{L}, \mathcal{V}, \leq_{\mathcal{T}}, \leq_{\mathcal{L}}, \leq_{\mathcal{V}}, \xrightarrow{\quad}, \vdash)$ be the abstract decomposition system defined as follows:

- The set \mathcal{T} is $\mathbf{B}_e(\mathcal{S})$ – we let $\alpha \leq_{\mathcal{T}} \beta$ if and only if there exists an address c such that $\alpha \subseteq (\beta|_c)$ and $\alpha(\varepsilon) = (\beta|_c)(\varepsilon)$.
- The set \mathcal{L} is the set of all $@$'s in \mathcal{S} – the relation $\leq_{\mathcal{L}}$ is the identity relation on this set.
- The set \mathcal{V} is $\mathbf{B}(\mathcal{S}) \times \mathbf{B}(\mathcal{S})$ – the relation $\leq_{\mathcal{V}}$ is defined by $(\alpha_1, \alpha_2) \leq_{\mathcal{V}} (\beta_1, \beta_2)$ if and only if $\alpha_1 \subseteq \beta_1$ and $\alpha_2 \subseteq \beta_2$.
- The relation $\xrightarrow{\quad}$ is defined by $\alpha \xrightarrow{@_\phi} (\beta_1, \beta_2)$ if and only if $\alpha = @_\phi(\beta_1, \beta_2)$.
- The relation \vdash is the least relation satisfying the following condition: if $V = (\alpha_1, \alpha_2)$, $i \in \{1, 2\}$, $\beta_1, \dots, \beta_n \in \mathbf{B}_e(\mathcal{S})$ and $\alpha_i = *_{\bar{a}}(\beta_1, \dots, \beta_n)$, then $V \vdash \beta_j$ for each $j \in [1, \dots, n]$.

Note that the elements of \mathcal{V} are pairs of blueprints that may be rootless. However, if $V \vdash \beta$, the blueprint β is always a rooted blueprint, so the relation \vdash is indeed a subset of $\mathcal{V} \times \mathcal{T}$.

- (A) For all $\mathcal{T}' \subseteq \mathcal{T}$, the relation \subseteq is an AFR on \mathcal{T}' if and only if $\leq_{\mathcal{T}}$ is an AFR on \mathcal{T}' . Indeed, consider an arbitrary infinite sequence $\bar{\alpha}$ over \mathcal{T}' . This sequence contains an infinite subsequence $(\alpha)_{i \in \mathbb{N}}$ such that all $\alpha_i(\varepsilon)$ are equal. Clearly, $\alpha_i \subseteq \alpha_j$ implies $\alpha_i \leq_{\mathcal{T}} \alpha_j$. Conversely, if $\alpha_i \leq_{\mathcal{T}} \alpha_j$, there exists c such that $\alpha_i \subseteq \alpha_{j|_c}$ and $\alpha_i(\varepsilon) = \alpha_{j|_c}(\varepsilon) = \alpha_j(c)$. So $\alpha_i \subseteq \alpha_{j|_c} \uparrow \alpha_j$, hence $\alpha_i \subseteq \alpha_j$.
- (B) We now check that all axioms of Theorem 6.13 are satisfied:

- It is clear that Axiom I is satisfied.
- The set of elementary terms is the set of all blueprints consisting of single formulas of \mathcal{S} . The relation $\leq_{\mathcal{T}}$ is, of course, an AFR on the set of elementary terms, that is, Axiom II is satisfied.
- Axiom III is immediate.
- If $(\alpha_1, \alpha_2) \leq_{\mathcal{V}} (\beta_1, \beta_2)$, then $\alpha_1 \subseteq \beta_1$ and $\alpha_2 \subseteq \beta_2$, hence

$$@_v(\alpha_1, \alpha_2) \subseteq @_v(\beta_1, \beta_2),$$

a fortiori

$$@_v(\alpha_1, \alpha_2) \leq_{\mathcal{T}} @_v(\beta_1, \beta_2),$$

so Axiom IV-bis is satisfied.

- To prove that Axiom V is satisfied, we let $\mathcal{W} \subseteq \mathcal{V}$. Then, by definition,

$$\mathcal{W} \vdash = \{\beta \in \mathcal{T} \mid \exists (\alpha_1, \alpha_2) \in \mathcal{W}, (\alpha_1, \alpha_2) \vdash \beta\}.$$

[†] Mellies' result was actually established for a different list of axioms (numbered from I to VI), but Mellies mentions the possibility of dropping Axiom VI and replacing Axiom IV with Axiom IV-bis in a remark following his proof of the main theorem.

Assuming $\leq_{\mathcal{F}}$ is an AFR on \mathcal{W}_{\vdash} , we will prove that $\leq_{\mathcal{V}}$ is an AFR on \mathcal{W} . By part (A), the relation \subseteq is an AFR on $\mathcal{W}_{\vdash} \subseteq \mathbb{B}_e(\mathcal{S})$. Let

$$\mathcal{B} = \{*_a(\beta_1, \dots, \beta_n) \mid \forall i \in [1, \dots, n], \beta_i \in \mathcal{W}_{\vdash}\}.$$

By Lemma 6.11, the relation \subseteq is an AFR on \mathcal{B} . Moreover, $\mathcal{W} \subseteq \mathcal{B} \times \mathcal{B}$. By Proposition 6.3(2), the relation $\leq_{\mathcal{V}}$ is an AFR on $\mathcal{B} \times \mathcal{B}$, and is thus an AFR on \mathcal{W} . \square

Lemma 6.15. For each formula ϕ , the set of all compact ϕ -shadows is a finite set effectively computable from ϕ .

Proof. For each compact ϕ -shadow Ξ and each address a such that a is a leaf in Ξ , we define a *step-continuation at a* of Ξ to be any compact ϕ -shadow Ξ' such that

$$\text{dom}(\Xi') \subseteq \text{dom}(\Xi) \cup \{a \cdot (1), a \cdot (2)\}$$

and Ξ and Ξ' take the same values on $\text{dom}(\Xi)$. Let \rightsquigarrow be the relation defined by $\Xi \rightsquigarrow \Xi'$ if and only if Ξ' is a step continuation of Ξ . By Lemma 5.5 and the fact that the set of subformulas of ϕ is a finite set, for all Ξ , the set of all Ξ' such that $\Xi \rightsquigarrow \Xi'$ is a finite set effectively computable from Ξ . Let \mathcal{C} be the closure under \rightsquigarrow of $\{(\varepsilon \mapsto (\varepsilon, \emptyset_{\mathbb{B}}, \phi))\}$. The set of all compact ϕ -shadows is clearly equal to this set, so it is enough to prove that \mathcal{C} is a finite set. In order to show a contradiction, we assume that \mathcal{C} is infinite. By König's Lemma, there exists an infinite sequence $\Xi_0 \rightsquigarrow \Xi_1 \rightsquigarrow \dots$ over \mathcal{C} . The union $\Xi_{\infty} = \bigcup_{i \geq 0} \Xi_i$ is a tree of infinite domain. By König's Lemma again, there exists an infinite chain of addresses $a_1 < a_2 < \dots$ such that all a_i are nodes of Ξ_{∞} with the same arity and labelled with the same subformula of ϕ . If $i < j$ and a_i and a_j are labelled with $(\bar{\chi}_i, \gamma_i, \psi)$ and $(\bar{\chi}_j, \gamma_j, \psi)$, we cannot have $\gamma_i \subseteq \gamma_j$, since otherwise there would exist a k such that Ξ_k is not compact. A contradiction then follows from Lemma 6.14. \square

7. From the shadows to the light

Theorem 7.1. Ticket Entailment is decidable.

Proof. The following propositions are equivalent:

- The formula ϕ is provable in the logic T_{\rightarrow} .
- The formula ϕ is inhabited by a combinator within the basis $\text{BB}'\text{IW}$.
- The formula ϕ is Λ_{NF} -inhabited (Lemma 2.10).
- There exists a compact Λ_{NF} -inhabitant of ϕ (Lemma 4.9).
- There exists a compact ϕ -shadow with the same tree domain as a Λ_{NF} -inhabitant of ϕ (Lemmas 4.9 and 5.13).

By Lemma 6.15, the set of compact ϕ -shadows is effectively computable from ϕ . By the subformula property (Lemma 2.5), for each shadow Ξ in this set, up to the choice of bound variables, there are only a finite number of Λ_{NF} -inhabitant of ϕ with the same domain as Ξ . Moreover, this set of inhabitants is clearly computable from Ξ and ϕ . Hence the existence of a Λ_{NF} -inhabitant of ϕ is decidable. \square

Acknowledgments

This work could not have been achieved without countless helpful comments and invaluable support from Paweł Urzyczyn, Paul-André Melliès and Pierre-Louis Curien. I am also deeply indebted to the anonymous referees for their remarkably careful reading of the paper.

References

- Anderson, A. R. (1960) Entailment shorn of modality. *Journal of Symbolic Logic* **25** (4) 388.
- Anderson, A. R. and Belnap Jr, N. D. (1975) *Entailment: The Logic of Relevance and Necessity*, Volume 1, Princeton University Press.
- Anderson, A. R., Belnap Jr, N. D. and Dunn, J. M. (1990) *Entailment: The Logic of Relevance and Necessity*, Volume 2, Princeton University Press.
- Barendregt, H. (1984) *The Lambda Calculus: Its Syntax and Semantics* (Revised edition), Studies in Logic and the Foundations of Mathematics **103**, North Holland.
- Barwise, J. (ed.) (1977) *Handbook of Mathematical Logic*, Studies in Logic and Foundations of Mathematics, North-Holland.
- Bezem, M., Klop, J. W. and de Vrijer, R. (eds.)[†] (2003) Term Rewriting Systems. *Cambridge Tracts in Theoretical Computer Science* **55**, Cambridge University Press.
- Bimbó, K. (2005) Types of I-free hereditary right maximal terms. *Journal of Philosophical Logic* **34** (5-6) 607–620.
- Broda, S., Damas, L., Finger, M. and Silva e Silva, P. S. (2004) The decidability of a fragment of *BB'IW*-logic. *Theoretical Computer Science* **318** (3) 373–408.
- Bunder, M. W. (1996) Lambda Terms Definable as Combinators. *Theoretical Computer Science* **169** (1) 3–21.
- Higman, G. (1952) Ordering by divisibility in abstract algebra. *Proceedings of the London Mathematical Society* **3** (2) 326–336.
- Kripke, S. (1959) The problem of entailment. *Journal of Symbolic Logic* **24** (4) 324.
- Krivine, J.-L. (1993) *Lambda-calculus, types and models*, Masson.
- Kruskal, J. B. (1972) The theory of well-quasi-ordering: A frequently discovered concept. *Journal of Combinatorial Theory, Series A* **13** (3) 297–305.
- Melliès, P.-A. (1998) On a duality between Kruskal and Dershowitz theorems. In: Larsen, K. G, Skyum, S. and Winskel, G. (eds.) *ICALP. Springer-Verlag Lecture Notes in Computer Science* **1443** 518–529.
- Trigg, P., Hindley, J. R. and Bunder, M. W. (1994) Combinatory abstraction using *B*, *B'* and friends. *Theoretical Computer Science* **135** (2) 405–422.
- Urquhart, A (1984) The undecidability of entailment and relevant implication. *Journal of Symbolic Logic* **49** (4) 1059–1073.

[†] Also known as ‘Terese’.