HAL open science

Ticket Entailment is decidable

Vincent Padovani

To cite this version:

Vincent Padovani. Ticket Entailment is decidable. 2010. hal-00599342v6

HAL Id: hal-00599342

https://hal.science/hal-00599342v6

Preprint submitted on 5 Jul 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This version was accepted for publication on December 19, 2011. The final version (with a few minor corrections, e.g. improved English) is to appear in Math. Struct. in Comp. Science.

Ticket Entailment is decidable

VINCENT PADOVANI

Equipe Preuves, Programmes et Systèmes

Université Paris VII - Denis Diderot

Case 7014

75205 PARIS Cedex 13

padovani@pps.jussieu.fr

Received 19 June 2010; Revised 6 March 2012

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

The partial bases built upon the atomic combinators B, $\mathsf{B'}$, 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 for well over 70 years (Trigg *et al.* 1994). The partial basis $\mathsf{BB'IW}$ corresponds, via the types of its combinators, to the system $T_\to$ of *Ticket Entailment* introduced and motivated in (Anderson and Belnap 1975; Anderson *et al.* 1990). The system $T_\to$ consists of modus ponens and four axiom schemes that range over the following types for each atomic combinator:

— $\mathsf{B} : (\chi \to \psi) \to ((\phi \to \chi) \to (\phi \to \psi))$
— $\mathsf{B'} : (\phi \to \chi) \to ((\chi \to \psi) \to (\phi \to \psi))$
— $\mathsf{I} : \phi \to \phi$
— $\mathsf{W} : (\phi \to (\phi \to \chi)) \to (\phi \to \chi)$

The type inhabitation problem for $\mathsf{BB'IW}$ 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 $T_\to$.

Surprisingly, the question of the decidability of $T_\to$ has remained unsolved since it was raised in (Anderson and Belnap 1975), although 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 $R_\to$ of *Relevant Implication* (which corresponds to the basis BCIW) and the system $E_\to$ of *Entailment* (Anderson and Belnap 1975) are both decidable (Kripke 1959) whereas the extensions R, E, T of $R_\to$, $E_\to$, $T_\to$ to a larger set of connectives ($\to$, $\wedge$, $\vee$) 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 which every maximal negative subformula is of 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 $\mathsf{BB'IW}$-terms, accordingly the type inhabitation problem within the basis $\mathsf{BB'IW}$ is equivalent to the type inhabitation problem for HRM-terms.

We use in this paper the same approach as Broda, Dams, Finger and Silva e Silva's. We prove that the type inhabitation problem for HRM-terms is decidable, and conclude that the logic $T_{\rightarrow}$ is decidable[†].

Summary

In Section 1, we recall the definition of hereditarily right-maximal terms and the equivalence between the decidability of type inhabitation for $\mathsf{BB'IW}$ and the decidability of type inhabitation for HRM-terms. The principle of our proof is depicted on Figure 1.

In Sections 2 and 3 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 those inhabitants belong to two larger sets of terms, the set of *compact* and *locally compact* inhabitants of ϕ.

In Section 4 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 define for shadows the analogue of compactness for terms and prove that the shadow of a compact term is itself compact.

Finally, in Section 5, 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 that this set is effectively computable from ϕ. The proof appeals to Higman Theorem and Kruskal Theorem – more precisely, to Melliès' Axiomatic Kruskal Theorem.

The decidability of the type inhabitation problem for HRM-terms and the decidability of $T_{\rightarrow}$ follow from this last 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.

Preliminaries

The first section of this paper assumes some familiarity with pure and simply-typed lambda-calculus and with the usual notions of α-conversion, β-reduction and β-normal form (Barendregt 1984; Krivine 1993). The 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 briefly recall the definitions and results used in Section 1.

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

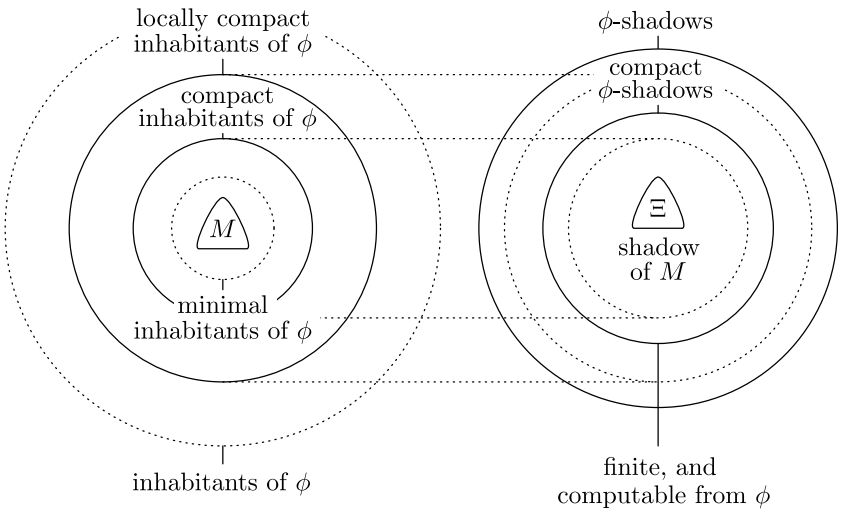

Fig. 1. Principle of the proof of decidability of type inhabitation for HRM-terms.

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

— 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.

Terms yielded by the second and third rules are called *abstractions* and *applications* respectively. The parentheses surrounding applications and abstractions are often omitted if unambiguous. We let $\lambda x_1 \ldots x_n.MN_1 \ldots N_p$ abbreviate $(\lambda x_1(\ldots(\lambda x_n(((MN_1)\ldots)N_p))\ldots))$. 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:

— $M = x$, or,
— $M = \lambda y.N$, $y \neq x$ and x is free in N, or,
— $M = NP$ 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 (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 $(M_1M_2)\langle x \leftarrow N\rangle = (M_1'M_2')$.

The *α-conversion* is defined as the least binary relation $\equiv_\alpha$ 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 $(M_1M_2) \equiv_\alpha (M_1'M_2')$.

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, however we will not follow this practice in our exposition.

The *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 *β-normal* – if there is no M' such that $M\beta M'$. A term M is *normalising* if there is a normal N – called *normal form* of M – such that $M\beta^*N$. It is *strongly normalising* if there is no infinite sequence $M = M_0\beta M_1\beta M_2 \ldots$

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 then its normal form is unique up to α-conversion.

The judgment "assuming $x_1, \ldots, x_n$ are of types $\psi_1, \ldots \psi_n$, the term M is of type ϕ", written $\{x_1 : \psi_1, \ldots, x_n : \psi_n\} \vdash M : \phi$, where $\psi_1, \ldots, \psi_n, \phi$ are formulas of propositional calculus and $x_1, \ldots, x_n$ 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 $\Gamma \vdash \lambda x.M : \phi \to \psi$.
— if $\Gamma \vdash M : \phi \to \psi$ 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:

1 (Strong normalisation) If $\Gamma \vdash M : \phi$, then M is strongly normalising.
2 (Subject reduction) If $\Gamma \vdash M : \phi$ and $M\beta N$, then $\Gamma \vdash N : \phi$.

1. From BB′IW to simply-typed lambda-calculus

The aim of this first 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 $\lambda fgx.f(gx)$, $\lambda fgx.g(fx)$, $\lambda x.x$, $\lambda hx.hxx$ in lambda-calculus, hence to each inhabited type ϕ in BB′IW 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 the ones that are reducts of translations of combinators within BB′IW.

The material and the results of this section are not new (Bunder 1996; Broda *et al.* 2004). The reader may as well skip the contents of Sections 1.3 and 1.4 entirely, accept Lemma 1.10 then go on with the study of stable parts and blueprints in Section 2.

The definition of hereditarily right-maximal terms is an adaptation of the definition given in (Bunder 1996). The proof of Lemma 1.6 (subject reduction for HRM-terms)

is similar to the proof of Property 2.4, p.375 in (Broda *et al.* 2004). The right-to-left implication of Lemma 1.10 can be deduced from Property 2.20, p.390 in (Broda *et al.* 2004), although our proof method seems to be simpler.

1.1. *Lambda-calculus*

Let $\mathcal{X}$ be a countably infinite set of variables $x, y, z \ldots$ 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 sequel, by *term* we always mean a term of lambda-calculus built over those variables. For each term M, we write $\mathsf{Free}(M)$ for the strictly increasing sequence of all free variables of M.

Terms are *not* identified modulo α-conversion - apart from Section 1, all considered terms will be in normal form, and the Greek letters α, β will be even used with new meaning at the beginning of Section 2. We adopt however the usual convention according to which two distinct λ's may not bound the same variable in a term, and no variable can be simultaneously free and bound in the same term.

1.2. *Hereditarily right-maximal terms*

Definition 1.1. 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 below (Lemma 1.6), 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 $\lambda fgx.f(gx)$, $\lambda fgx.g(fx)$, $\lambda x.x$, $\lambda hx.hxx$ are HRM, whereas $\lambda yz.zy$ is not, no matter if $y < z$ or $y > z$.

Definition 1.2. 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, letting $\Omega(x_1, \ldots, x_n) = (\Omega(x_1), \ldots, \Omega(x_n))$.

Definition 1.3. The judgment $M : \phi$, in words "M is of type ϕ w.r.t Ω", is defined by:
— if $\Omega(x) = \phi$, then $x : \phi$,
— if $x : \chi$, $M : \psi$ and $\lambda x.M$ is HRM, then $\lambda x.M : \chi \to \psi$,
— if $M : \chi \to \psi$, $N : \chi$ and (MN) is HRM, then $(MN) : \psi$.

The function Ω will remain fixed throughout our exposition. Accordingly the type of a

term M w.r.t 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 1.4. We write Λ_{NF} for the set of all typed terms in β-normal form. We call Λ_{NF}*-inhabitant of* ϕ every closed term $M \in \Lambda_{\mathrm{NF}}$ of type ϕ.

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

Lemma 1.5. (Subformula Property) Let M be a Λ_{NF}-inhabitant of ϕ. The types of the subterms of M are subformulas of ϕ.

1.3. *Subject reduction of hereditarily right-maximal terms*

Lemma 1.6. Suppose there exists a closed $M : \phi$. Then ϕ is Λ_{NF}-inhabited.

Proof. (1) We leave to the reader the proof of the fact that for every variable y and for every $N : \phi$, there exists $N' \equiv_\alpha N$ such that $N' : \phi$ and every bound variable of N' is strictly greater than y.

(2) We prove the following proposition by induction on P. Let P, Q be typed HRM-terms. Suppose:

— x and Q are of the same type,
— if Q is closed and $x \in \mathsf{Free}(P)$, then $x = \min(\mathsf{Free}(P))$
— if Q is not closed, then for all $z \in \mathsf{Free}(P)$:
  if $z < x$ then $z \leq \max(\mathsf{Free}(Q))$,
  if $x < z$ then $\max(\mathsf{Free}(Q)) < z$.
— if Q is not closed, then $\max(\mathsf{Free}(Q)) < z$ for all bound variables z of P.

Then $R = P\langle x \leftarrow Q\rangle$ is defined, HRM and of the same type as P. The proposition is clear if P is a variable.

Suppose $P = \lambda z.P'$. Then $\mathsf{Free}(P') = \mathsf{Free}(P) \cdot (z)$. By induction hypothesis $R' = P'\langle x \leftarrow Q\rangle$ is defined, HRM and of the same type as P'. The variable z is still the greatest free variable of R' and z is not free in Q, hence $R = \lambda z.R'$.

Suppose $P = (P_1 P_2)$. By induction hypothesis $R_i = P_i\langle x \leftarrow Q\rangle$ is defined, HRM and of the same type as P_i for each $i \in \{1, 2\}$. It remains to check that $R = (R_1 R_2)$ is HRM. Assume x is free in P and P_1 is not closed.

Suppose $\max(\mathsf{Free}(P_1)) > x$. Then $\max(\mathsf{Free}(P_1)) = \max(\mathsf{Free}(R_1)) \leq \max(\mathsf{Free}(P_2)) = \max(\mathsf{Free}(R_2))$.

Suppose $\max(\mathsf{Free}(P_1)) < x$. The term Q cannot be closed, and $\max(\mathsf{Free}(P_1)) = \max(\mathsf{Free}(R_1)) \leq \max(\mathsf{Free}(Q))$. We have either $\max(\mathsf{Free}(P_2)) = x$ and $\max(\mathsf{Free}(R_2)) = \max(\mathsf{Free}(Q))$, or $\max(\mathsf{Free}(P_2)) > x$ and $\max(\mathsf{Free}(P_2)) = \max(\mathsf{Free}(R_2))$.

Otherwise $\max(\mathsf{Free}(P_1)) = x$. Suppose $\max(\mathsf{Free}(P_2)) > x$. Then $\max(\mathsf{Free}(P_2)) = \max(\mathsf{Free}(R_2))$. If Q is closed, then $\mathsf{Free}(P_1) = (x)$ and R_1 is closed. Otherwise we have $\max(\mathsf{Free}(R_1)) = \max(\mathsf{Free}(Q)) \leq \max(\mathsf{Free}(P_2))$. The remaining case is $\max(\mathsf{Free}(P_2)) = x$. If Q is closed then $\mathsf{Free}(P_1) = \mathsf{Free}(P_2) = (x)$ and R_1, R_2 are closed. Otherwise $\max(\mathsf{Free}(R_1)) = \max(\mathsf{Free}(R_2) = \max(\mathsf{Free}(Q))$.

(3) Assume $N : \phi$ and N is not in normal form. We prove by induction on N the

existence of $N' : \phi$ such that $N\beta N'$. If $N = \lambda x.P$, or if $N = (N_1 N_2)$ 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 $P' \equiv_\alpha P$ 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 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. By (1), there exists $N' \equiv_\alpha N$ such that $N' : \phi$, that is, ϕ is Λ_{NF}-inhabited, □

1.4. *Equivalence between inhabitation in* BB′IW *and* Λ_{NF}*-inhabitation*

In the next three lemmas by $\phi_1 \ldots \phi_n \to \psi$ we mean the formula $(\phi_1 \to (\ldots(\phi_n \to \psi)\ldots))$ if $n > 0$, and otherwise the formula ψ. We write $\vdash_{\mathsf{BB'IW}} \phi$ for the judgment "there exists within the basis BB′IW a combinator of type ϕ".

Lemma 1.7. If $\vdash_{\mathsf{BB'IW}} \phi$, then ϕ is Λ_{NF}-inhabited.

Proof. If $f < g < x$ and $h < x$, then $\lambda x.x$, $\lambda fgx.f(gx)$, $\lambda fgx.g(fx)$ and $\lambda hx.hxx$ are HRM. For each type ϕ of an atomic combinator, the variables f, g, h, x can be chosen so that one of those terms is of 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 1.6, every such formula is Λ_{NF}-inhabited. □

Lemma 1.8. If $\vdash_{\mathsf{BB'IW}} \chi \to \psi$, then $\vdash_{\mathsf{BB'IW}} (\phi_1 \ldots \phi_n \to \chi) \to (\phi_1 \ldots \phi_n \to \psi)$ for all $\phi_1, \ldots, \phi_n$.

Proof. By induction on n, using left-applications of B. □

Lemma 1.9. Suppose $(i_1, \ldots, i_n)$, $(j_1, \ldots, j_m)$, $(k_1, \ldots, k_p)$ are strictly increasing sequences of integers, $\{k_1, \ldots, k_p\} = \{i_1, \ldots, i_n, j_1, \ldots, j_m\}$, $n = 0$ or ($n > 0$, $m > 0$, $i_n \leq j_m$). If

1 $\vdash_{\mathsf{BB'IW}} \omega_{i_1} \ldots \omega_{i_n} \to (\chi \to \psi)$,
2 $\vdash_{\mathsf{BB'IW}} \omega_{j_1} \ldots \omega_{j_m} \to \chi$,

then $\vdash_{\mathsf{BB'IW}} \omega_{k_1} \ldots \omega_{k_p} \to \psi$.

Proof. By induction on $n+m$. The proposition is true if $n = m = 0$. Assume $n+m > 0$. Then $m > 0$.

Suppose $n = 0$. Then $(j_i, \ldots, j_m) = (k_1, \ldots, k_p)$. We have:

(i) $\vdash_{\mathsf{BB'IW}} (\chi \to \psi) \to ((\omega_{j_m} \to \chi) \to (\omega_{j_m} \to \psi))$
(ii) $\vdash_{\mathsf{BB'IW}} (\omega_{j_m} \to \chi) \to (\omega_{j_m} \to \psi)$

where: (i) is a type for B; (ii) follows from (i), (1) and modus ponens. If $m = 1$ then

$\vdash_{\mathsf{BB'IW}} \omega_{j_1} \to \psi$ follows from (ii), (2) and modus ponens. Otherwise $\vdash_{\mathsf{BB'IW}} \omega_{j_1} \dots \omega_{j_m} \to \psi$ follows from (ii), (2) and the induction hypothesis.

We now assume $n > 0$. Suppose $m > 1$ and $i_n \leq j_{m-1}$. Then

(iii) $\vdash_{\mathsf{BB'IW}} (\chi \to \psi) \to ((\omega_{j_m} \to \chi) \to (\omega_{j_m} \to \psi))$

(iv) $\vdash_{\mathsf{BB'IW}} (\omega_{i_1} \dots \omega_{i_n} \to (\chi \to \psi)) \to (\omega_{i_1} \dots \omega_{i_n} \to ((\omega_{j_m} \to \chi) \to (\omega_{j_m} \to \psi)))$

(v) $\vdash_{\mathsf{BB'IW}} \omega_{i_1} \dots \omega_{i_n} \to ((\omega_{j_m} \to \chi) \to (\omega_{j_m} \to \psi))$

where: (iii) is a type for B; (iv) follows from (iii) and Lemma 1.8; (v) follows from (iv), (1) and modus ponens. We 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 $\vdash_{\mathsf{BB'IW}} \omega_{k_1} \dots \omega_{k_{p-1}} \to (\omega_{j_m} \to \psi)$ by (v), (2) and the induction hypothesis.

Suppose $m = 1$ or ($m > 1$ and $i_n > j_{m-1}$). Then

(vi) $\vdash_{\mathsf{BB'IW}} (\omega_{j_m} \to \chi) \to ((\chi \to \psi) \to (\omega_{j_m} \to \psi))$

(vii) $\vdash_{\mathsf{BB'IW}} (\omega_{j_1} \dots \omega_{j_m} \to \chi) \to (\omega_{j_1} \dots \omega_{j_{m-1}} \to ((\chi \to \psi) \to (\omega_{j_m} \to \psi)))$

(viii) $\vdash_{\mathsf{BB'IW}} \omega_{j_1} \dots \omega_{j_{m-1}} \to ((\chi \to \psi) \to (\omega_{j_m} \to \psi))$

(ix) $\vdash_{\mathsf{BB'IW}} \omega_{n_1} \dots \omega_{n_q} \to (\omega_{j_m} \to \psi)$

where: (vi) is a type for B'; (vii) follows from (vi) and Lemma 1.8; (viii) follows from (vii), (2) and modus ponens; $\{n_1, \dots, n_q\} = \{j_1, \dots, j_{m-1}, i_1, \dots, i_n\}$; (ix) follows from (viii), (1) and the induction hypothesis. If $j_m > i_n$, then $(n_1, \dots, n_q, j_m) = (k_1, \dots, k_p)$. Otherwise $j_m = i_n$, $n_q = i_n$, $(n_1, \dots n_q) = (k_1, \dots, k_p)$ and

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

(xi) $\vdash_{\mathsf{BB'IW}} (\omega_{i_n} \to (\omega_{i_n} \to \psi)) \to (\omega_{i_n} \to \psi)$

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

(xiii) $\vdash_{\mathsf{BB'IW}} \omega_{k_1} \dots \omega_{k_{p-1}} \to (\omega_{i_n} \to \psi)$

where: (x) is (ix); (xi) is a type for W; (xii) follows from (xi) and Lemma 1.8; (xiii) follows from (x), (xii) and modus ponens; (xiii) is $\vdash_{\mathsf{BB'IW}} \omega_{k_1} \dots \omega_{k_p} \to \psi$. □

Lemma 1.10. For every formula ϕ, we have $\vdash_{\mathsf{BB'IW}} \phi$ if and only if ϕ is Λ_{NF}-inhabited.

Proof. The left to right implication is Lemma 1.7. Using Lemma 1.9 when M is an application, an immediate induction on M shows that if $M : \psi$, $\mathsf{Free}(M) = (x_1, \dots, x_n)$ and $x_1 : \chi_1, \dots, x_n : \chi_n$, then $\vdash_{\mathsf{BB'IW}} \chi_1 \dots \chi_n \to \psi$ □

2. Stable parts and blueprints

The last lemma showed that the decidability of type inhabitation for $\mathsf{BB'IW}$ is equivalent to the decidability of Λ_{NF}-inhabitation. The sequel is devoted to the elaboration of a decision algorithm for the latter problem.

The problem we shall examine throughout Sections 2 and 3 is the following: 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 of the same type? This question is not easy because we are dealing with a lambda-calculus restricted with strong structural constraints (righ-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 of the same type or abstractions of the same type.
— $\mathsf{Free}(N) = X = (x_1, \ldots, x_n)$,
— $\mathsf{Free}(P) = Y = (y_0^1, \ldots, y_{p_1}^1, \ldots, y_0^n, \ldots, y_{p_n}^n)$
— $\Omega(X) = (\chi_1, \ldots, \chi_n)$,
— $\Omega(Y) = (\chi_0^1, \ldots, \chi_{p_1}^1, \ldots, \chi_0^n, \ldots, \chi_{p_n}^n)$,
— $\chi_j^i = \chi_i$ for each i, j.

Then 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' of the same size as P, of the same type and 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 of the same type but of strictly smaller size.

This simple property is far from being enough to characterise the minimal inhabitants of a formula: there are indeed formulas with inhabitants of abitrary 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 our 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 2.1, 2.2 and 2.2 we show how to build from any term $M \in \Lambda_{\mathrm{NF}}$ a partial tree labelled with formulas. This partial tree is called the *blueprint* of M. This blueprint can be seen as a synthesized 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 2.4 and 2.5 we introduce a rewriting relation on blueprints that allows one to "extract" sequences of formulas from a blueprint.
— In section 2.6 we prove our two key lemmas. Lemma 2.15 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, $\Omega(\mathsf{Free}(M))$) can always be extracted from its blueprint. Lemma 2.16 shows that for *every* sequence of formulas $\overline{\phi}$ 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 of the same type and with the same blueprint as M, and such that $\Omega(\mathsf{Free}(N)) = \overline{\phi}$.

As a continuation of our first example, let us 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 of the same type or abstractions of the same type. Suppose:

— the sequence $\Omega(\mathsf{Free}(N))$ can be extracted from the blueprint of P.

This situation is a generalization of the preceding one (in our first example $\Omega(X)$ could also be extracted from the blueprint of P, see Definition 2.10). The term M is still not of minimal size. Indeed, we may use the second key lemma to prove the existence of (non-uniform) renaming of the free variables of P that will produce a term P' of the same type as P such that $\mathsf{Free}(P') = \mathsf{Free}(N)$. The term N can be replaced with P' in M.

2.1. *Partial trees and trees*

Definition 2.1. Let $(\mathbb{A}, \leq)$ be the set of all finite sequences over the set $\mathbb{N}_+$ of natural numbers, ordered by prefix ordering. Elements of $\mathbb{A}$ are called *addresses*. We call *partial tree* every 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 $\{c \mid a \cdot c \in \mathrm{dom}(\pi)\}$.

Definition 2.2. For all partial trees π, π' and for every address a, we let $\pi[a \leftarrow \pi']$ denote the partial tree π'' such that $\pi''_{|a} = \pi'$ and $\pi''(b) = \pi(b)$ for all $b \in \mathrm{dom}(\pi)$ such that $a \not\leq b$.

Definition 2.3. 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 $j \in \{1, \ldots, i-1\}$. 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 call *tree* every function whose domain is a tree domain.

In the sequel 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; $(M_1 M_2)$ with the tree τ mapping ε to @ and such that $\tau_{|(i)}$ is the tree of M_i for each $i \in \{1, 2\}$.

2.2. *Blueprints*

Definition 2.4. 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 $@_\phi$ is of arity 2.

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

For each $\mathcal{S} \subseteq \mathfrak{S}$, we call $\mathcal{S}$*-blueprint* every 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 2.5. For every blueprint α and every address a, the *relative depth of a in α* is the number of $b \in \mathrm{dom}(\alpha)$ such that $b < a$. The *relative depth of* α is defined as 0 if α is of empty domain, the maximal relative depth of an address in α otherwise.

In the sequel the following notations will be used to denote blueprints (see Figure 2):

— $\emptyset_{\mathbb{B}}$ denotes the blueprint of empty domain.
— we abbreviate $\varepsilon \mapsto \phi$ as ϕ.
— $@_\phi(\alpha_1, \alpha_2)$ denotes the (rooted) blueprint α such that $\alpha(\varepsilon) = \phi$, $\alpha_{|(1)} = \alpha_1$, $\alpha_{|(2)} = \alpha_1$.
— for every sequence $\overline{a} = (a_1, \ldots, a_k)$ of pairwise incomparable addresses, $*_{\overline{a}}(\alpha_1, \ldots, \alpha_k)$ denotes the blueprint α of minimal domain such that $\alpha_{|a_i} = \alpha_i$ for each $i \in [1, \ldots, k]$.
— we let $*(\alpha_1, \ldots, \alpha_k)$ denote the blueprint $*_{\overline{a}}(\alpha_1, \ldots, \alpha_k)$ such that $\overline{a} = ((1), \ldots, (k))$.

For each blueprint α, the choice of $\overline{a}, \alpha_1, \ldots, \alpha_k$ such that $\alpha = *_{\overline{a}}(\alpha_1, \ldots, \alpha_k)$ is obviously not unique. The sequence $(\alpha_1, \ldots, \alpha_k)$ may contain an arbitrary number of empty blueprints, hence the sequence $\overline{a}$ may be of arbitrary length. Also, α can be roooted (if

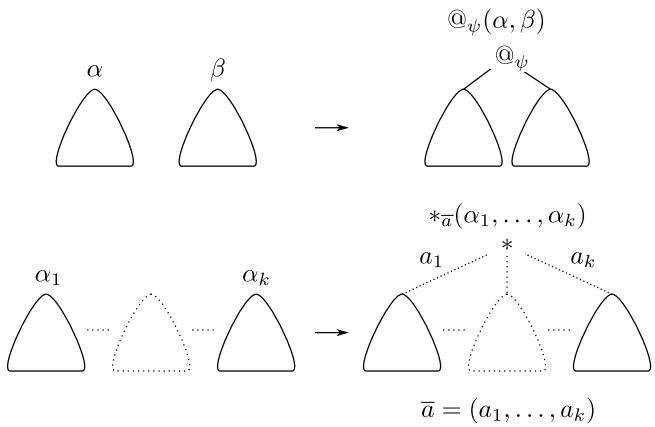

Fig. 2. Construction of blueprints, with the notations of Section 2.2. In the upper diagram, the blueprints α and β must be non-empty. Although $\alpha_1, \ldots, \alpha_k$ are displayed from left to right, the sequence $(a_1, \ldots, a_k)$ needs not to be lexicographically ordered.

$k = 1$, $a_1 = \varepsilon$ and α_1 is rooted) or empty (if $k = 0$ or $\alpha_1 = \ldots = \alpha_k = \emptyset_{\mathbb{B}}$). Those ambiguities will not be difficult to deal with, but they will require a few precautions in our proofs and definitions by induction on blueprints.

2.3. *Blueprint of a term*

Definition 2.6. For all $M \in \Lambda_{\mathrm{NF}}$, the *stable part* of M is the set of all $a \in \mathrm{dom}(M)$ such that $\mathsf{Free}(M_{|a}) \subseteq \mathsf{Free}(M)$ 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 is of empty stable part if and only if it is closed.

Definition 2.7. For all $M \in \Lambda_{\mathrm{NF}}$, we call *blueprint of* M 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 let $M \Vdash \alpha$ denote the judgment "M is of blueprint α" (Figure 3).

If $M = (M_1 M_2) \in \Lambda_{\mathrm{NF}}$, $M : \phi$, $M_1 \Vdash \alpha_1$, $M_2 \Vdash \alpha_2$, then each α_i is of non-empty domain and $(M_1 M_2) \Vdash @_\phi(\alpha_1, \alpha_2)$ – in other words the so-called blueprint of M is indeed a blueprint, provided so are the blueprints of M_1, M_2. When $M = \lambda x.M_1$ the blueprint of M is of the form $*(\alpha)$ – the relation between α and the blueprint of M_1 in that case will be clarified by Lemma 2.15.

Lemma 2.8. For all $M \in \Lambda_{\mathrm{NF}}$ and forall $a \cdot b \in \mathrm{dom}(M)$:

1 If $\mathsf{Free}(M_{|a \cdot b}) \subseteq \mathsf{Free}(M)$ then $\mathsf{Free}(M_{|a \cdot b}) \subseteq \mathsf{Free}(M_{|a})$.
2 If $M_{|a} \Vdash \alpha$ and $M_{|a \cdot b} \Vdash \beta$, then $\alpha_{|b} = \beta$.

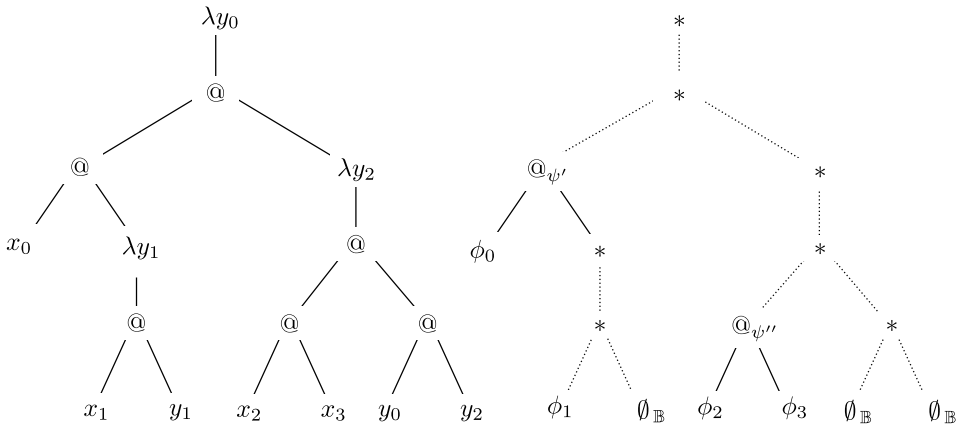

Fig. 3. An element of Λ_{NF} with its blueprint ($x_0 < x_1 < y_1$, $x_2 < x_3 < y_0 < y_2$, $x_1 < y_0 < y_2$).

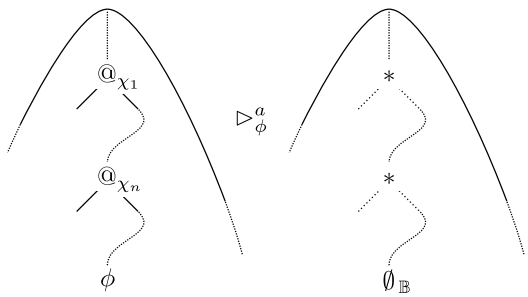

Fig. 4. Principle of blueprint reduction.

Proof. The first proposition is a consequence of our bound variable convention (see Section 1.1): if $\mathsf{Free}(M) = X$, $\mathsf{Free}(M_{|a}) = X' \cup Y$ where $X' \subseteq X$ and X, Y are disjoint, then every element of $\mathsf{Free}(M_{|a\cdot b})$ 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 proposition is equivalent to the first. □

2.4. *Extraction of the formulas of a blueprint*

Definition 2.9. The judgment "β is the blueprint obtained by extracting the formula ϕ at the address a in the blueprint α", written $\alpha \rhd^a_\phi \beta$, is inductively defined by:

1 $\phi \rhd^\varepsilon_\phi \emptyset_{\mathbb{B}}$,

2 if $\alpha \rhd^a_\phi \beta$, then $@_\psi(\gamma, \alpha) \rhd^{(2)\cdot a}_\phi *(\gamma, \beta)$

3 if $\alpha \rhd^a_\phi \beta$, then $*_{(b,c_1,\ldots,c_k)}(\alpha, \gamma_1, \ldots, \gamma_k) \rhd^{b\cdot a}_\phi *_{(b,c_1,\ldots,c_k)}(\beta, \gamma_1, \ldots, \gamma_n)$.

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.

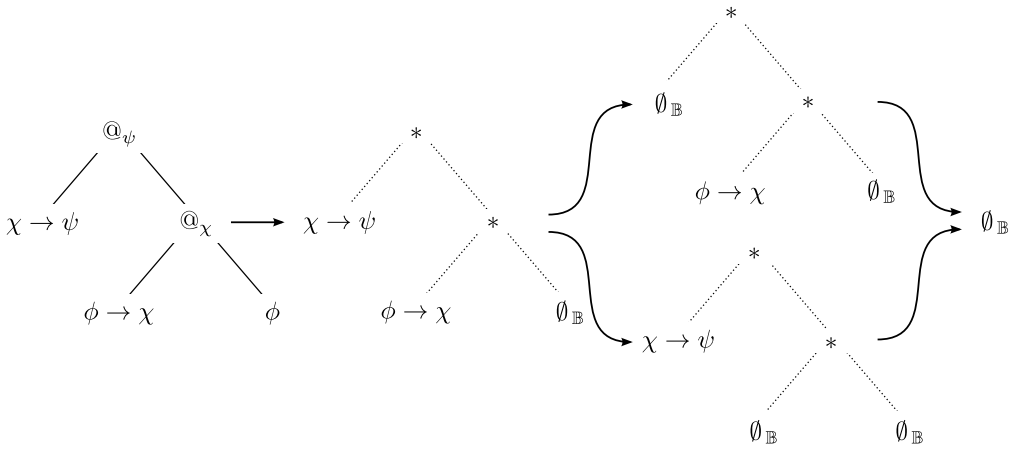

Fig. 5. Full reductions of $@_\psi(\chi \to \psi, @_\chi(\phi \to \chi, \phi))$ to $\emptyset_{\mathbb{B}}$.

For instance (Figure 5):

— $@_\psi(\chi \to \psi, @_\chi(\phi \to \chi, \phi)) \quad \rhd^{(2,2)}_{\phi} \quad *(\chi \to \psi, *(\phi \to \chi, \emptyset_{\mathbb{B}}))$
$\rhd^{(2,1)}_{\phi\to\chi} \quad *(\chi \to \psi, *(\emptyset_{\mathbb{B}}, \emptyset_{\mathbb{B}}))$
$\rhd^{(1)}_{\chi\to\psi} \quad *(\emptyset_{\mathbb{B}}, *(\emptyset_{\mathbb{B}}, \emptyset_{\mathbb{B}})) = \emptyset_{\mathbb{B}}$

— $@_\psi(\chi \to \psi, @_\chi(\phi \to \chi, \phi)) \quad \rhd^{(2,2)}_{\phi} \quad *(\chi \to \psi, *(\phi \to \chi, \emptyset_{\mathbb{B}}))$
$\rhd^{(1)}_{\chi\to\psi} \quad *(\emptyset_{\mathbb{B}}, *(\phi \to \chi, \emptyset_{\mathbb{B}}))$
$\rhd^{(2,1)}_{\phi\to\chi} \quad *(\emptyset_{\mathbb{B}}, *(\emptyset_{\mathbb{B}}, \emptyset_{\mathbb{B}})) = \emptyset_{\mathbb{B}}$

When $\alpha \rhd^a_\phi \beta$, 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 branch of @. 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.

2.5. *Sets of extractible sequences*

Definition 2.10. For each formula ϕ, let $\rhd_\phi$ be the relation defined by: $\alpha \rhd_\phi \beta$ if and only if there exists a such that $\alpha \rhd^a_\phi \beta$. We write $\rhd^+_\phi$ for the transitive closure of $\rhd_\phi$. For each α, we write $\mathbb{F}(\alpha)$ for the set of all sequences $(\phi_1, \ldots, \phi_n)$ such that $\alpha \rhd^+_{\phi_n} \ldots \rhd^+_{\phi_1} \emptyset_{\mathbb{B}}$.

The set $\mathbb{F}(\alpha)$ is what we called "set of extractible sequences of α" in the introduction of Section 2. Note that $\mathbb{F}(\emptyset_{\mathbb{B}}) = \{\varepsilon\}$. If $\alpha \neq \emptyset_{\mathbb{B}}$, then all elements of $\mathbb{F}(\alpha)$ are non-empty sequences. Note also that each $\rhd$-reduction strictly decreases the cardinality of the domain of a blueprint, therefore $\mathbb{F}(\alpha)$ is a finite set for all α. We now introduce the notion of *shuffle* which will allow us to characterise $\mathbb{F}(\alpha)$ depending on the structure of α.

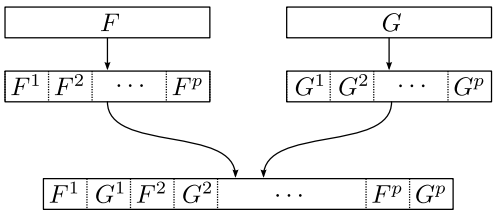

Fig. 6. Shuffling of two sequences. The chunks of F and G need not to be of the same size – some of them can be empty. Every contraction of the resulting sequence belongs to $\circledast(F, G)$. Each contraction belongs also to $\circledcirc(F, G)$ when F, G are non-empty and the last chunk G^p of G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

Definition 2.11. A *contraction* of a sequence F is either the sequence F or a sequence $G \cdot (f) \cdot H$ where $G \cdot (f) \cdot (f) \cdot H$ is a contraction of F.

Definition 2.12. For all finite sequences $F_1, \ldots, F_n$ we call *shuffle* of $(F_1, \ldots, F_n)$ every sequence $F_1^1 \cdot \ldots \cdot F_n^1 \cdot \ldots \cdot F_1^p \cdot \ldots \cdot F_n^p$ such that $F_i^1 \cdot \ldots \cdot F_i^p = F_i$ for each i. For each tuple of sets of finite sequences $(\mathcal{F}_1, \ldots, \mathcal{F}_n)$ we write $\circledast(\mathcal{F}_1, \ldots, \mathcal{F}_n)$ for the closure under contraction of the set of shuffles of elements of $\mathcal{F}_1 \times \ldots \times \mathcal{F}_n$.

Definition 2.13. Given two non-empty finite sequences F_1, F_2, we call *right-shuffle* of (F_1, F_2) every sequence $F_1^1 \cdot F_2^1 \cdot \ldots \cdot F_1^p \cdot F_2^p$ such that $F_i^1 \cdot \ldots F_i^p = F_i$ 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 $\circledcirc(\mathcal{F}_1, \mathcal{F}_2)$ 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 depicted on Figure 6. The following properties follow from our definitions and will be used without reference:

1 If $\alpha = \emptyset_{\mathbb{B}}$, then $\mathbb{F}(\alpha) = \{\varepsilon\}$.
2 If $\alpha = \phi$, then $\mathbb{F}(\alpha) = \{(\phi)\}$.
3 If $\alpha = *_{\overline{a}}(\alpha_1, \ldots, \alpha_k)$, then $\mathbb{F}(\alpha) = \circledast(\mathbb{F}(\alpha_1), \ldots, \mathbb{F}(\alpha_k))$.
4 If $\alpha = @_{\phi}(\alpha_1, \alpha_2)$, then $\mathbb{F}(\alpha) = \circledcirc(\mathbb{F}(\alpha_1), \mathbb{F}(\alpha_2))$.

2.6. *Abstraction vs. extraction*

Lemma 2.14. Suppose $\{a_1, \ldots, a_p\} = \{b_1, \ldots, b_p\}$, and:
— $\alpha \rhd_{\chi}^{a_1} \ldots \rhd_{\chi}^{a_p} \beta$,
— $\alpha \rhd_{\chi}^{b_1} \ldots \rhd_{\chi}^{b_p} \beta'$.
Then $\beta = \beta'$.

Proof. By an easy induction on α. □

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

The next lemma shows in particular that if $M, \lambda x.M \in \Lambda_{\mathrm{NF}}$, then M and $\lambda x.M$

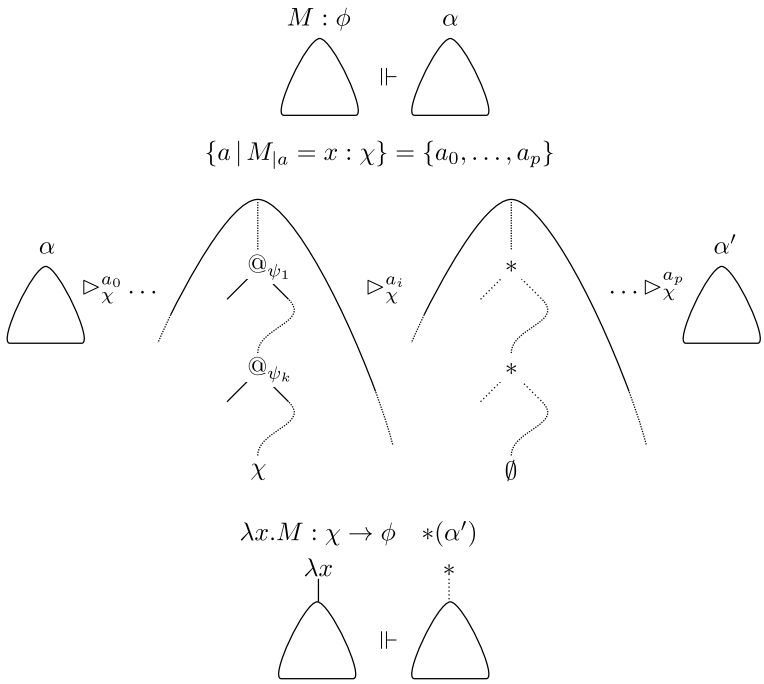

Fig. 7. How the blueprint of M evolves into the blueprint of $\lambda x.M$

are of blueprints α and β if and only if there exist $a_0, \ldots, a_p$ such that $\{a_0, \ldots, a_p\} = \{a \mid M_{|a} = x\}$, $\alpha \rhd^{a_0}_{\chi} \ldots \rhd^{a_p}_{\chi} \alpha'$ and $\beta = *(\alpha')$ (Figure 7).

Lemma 2.15. Suppose $M \in \Lambda_{\mathrm{NF}}$ is of blueprint α, with $\mathsf{Free}(M) = (x_1, \ldots, x_n)$ and $\Omega(x_1, \ldots, x_n) = (\chi_1, \ldots, \chi_n)$. For each $i \in [0, \ldots, n]$:

— let α_i be the restriction of α to $\mathrm{dom}(\alpha) \cap \{a \mid \mathsf{Free}(M_{|a}) \subseteq \{x_1, \ldots, x_i\}\}$.
— let β_i be the blueprint of $\lambda x_{i+1} \ldots x_n.M$,

Then:

1 For each $i \in [0, \ldots, n]$ we have $\mathrm{dom}(\beta_i) = \{1^{n-1} \cdot a \mid a \in \mathrm{dom}(\alpha_i)\}$ and $\beta_{i|1^{n-i}} = \alpha_i$.
2 For each $i \in]0, \ldots, n]$:
  (a) there exist $a^i_0, \ldots, a^i_{p_i}$ such that $\{a^i_0, \ldots, a^i_{p_i}\} = \{a \mid M_{|a} = x_i\}$ and $\alpha_i \rhd^{a^i_0}_{\chi_i} \ldots \rhd^{a^i_{p_i}}_{\chi_i} \alpha_{i-1}$,
  (b) if $\{b_0, \ldots, b_{p_i}\} = \{a \mid M_{|a} = x_i\}$ and $\alpha_i \rhd^{b_0}_{\chi_i} \ldots \rhd^{b_{p_i}}_{\chi_i} \alpha'$ then $\alpha' = \alpha_{i-1}$.
3 We have $(\chi_1, \ldots, \chi_n) \in \mathbb{F}(\alpha)$.

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 (2.a). Property (2.b) follows from Property (2.a) and Lemma 2.14. As to prove (2.a) we introduce the following notations.

For each $N \in \Lambda_{\mathrm{NF}}$, 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 $\rho_N(Y,\gamma) = \delta$, $\{b \mid N_{|b} = y\} = \{b_0, \ldots, b_m\}$ and $\delta \rhd_\chi^{b_0} \ldots \rhd_\chi^{b_m} \delta'$, then $\rho_M((y)\cdot Y, \gamma) = \delta'$. By Lemma 2.14, if $\{b \mid N_{|b} = y\} = \{b_0, \ldots, b_m\} = \{c_0, \ldots, c_m\}$, $\delta \rhd_\chi^{b_0} \ldots \rhd_\chi^{b_m} \delta'$ and $\delta \rhd_\chi^{c_0} \ldots \rhd_\chi^{c_m} \delta''$, then $\delta' = \delta''$, thus 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 $\mathrm{dom}(\gamma) \cap \{b \mid \mathsf{Free}(N_{|b}) \subseteq Y'\}$.

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

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

thus (2.a) holds. The case $X' = \varepsilon$ is immediate, hence we may as well assume that X' is a non-empty suffix of $\mathsf{Free}(M)$. The case of M equal to a variable follows immediately from our definitions.

Suppose $M = (M_1 M_2)$, $M_1 \Vdash \gamma_1$ and $M_2 \Vdash \gamma_2$. There exist X_1, X_2, X_1', X_2' such that: $X_1 \cup X_2 = X$; $X_1' \cup X_2' = X'$; $\mathsf{Free}(M_j) = X_j \cdot X_j'$ for each $j \in \{1, 2\}$. 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 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 hence the last elements of X', X_2' are equal. Assume $X' = X'' \cdot (x)$ and $X_2' = X_2'' \cdot (x)$. If x is not the last element of X_1' then:

$$\begin{aligned}
\rho_M(X', \alpha) &= \rho_M(X'' \cdot (x), @_\psi(\gamma_1, \gamma_2)) \\
&= \rho_M(X_1' \cup X_2'', *(\gamma_1, \rho_{M_2}((x), \gamma_2))) \\
&= *(\rho_{M_1}(X_1', \gamma_1), \rho_{M_2}(X_2'', \rho_{M_2}((x), \gamma_2))) \\
&= *(\rho_{M_1}(X_1', \gamma_1), \rho_{M_2}(X_2'' \cdot (x), \gamma_2)) \\
&= *(\rho_{M_1}(X_1', \gamma_1), \rho_{M_2}(X_2', \gamma_2))
\end{aligned}$$

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

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

Suppose $M = \lambda x.M_1$, $M_1 \Vdash \gamma_1$. By induction hypothesis $\mu_{M_1}(X, \gamma_1) = \rho_{M_1}(X' \cdot (x), \gamma_1) = \rho_{M_1}(X', \rho_{M_1}((x), \gamma_1)) = \rho_{M_1}(X', \mu(X \cdot X', \gamma_1)) = \rho_{M_1}(X', \alpha_{|(1)})$. Moreover $\mu_{M_1}(X, \gamma_1) = \mu_{M_1}(X, \mu_{M_1}(X \cdot X', \gamma_1)) = \mu_{M_1}(X, \alpha_{|(1)})$. Hence $\mu_{M_1}(X, \alpha_{|(1)}) = \rho_{M_1}(X', \alpha_{|(1)})$, therefore $\mu_{M_1}(X, \alpha) = \rho_{M_1}(X', \alpha)$. □

Thus the full sequence of the types of the free variables of M can be extracted from

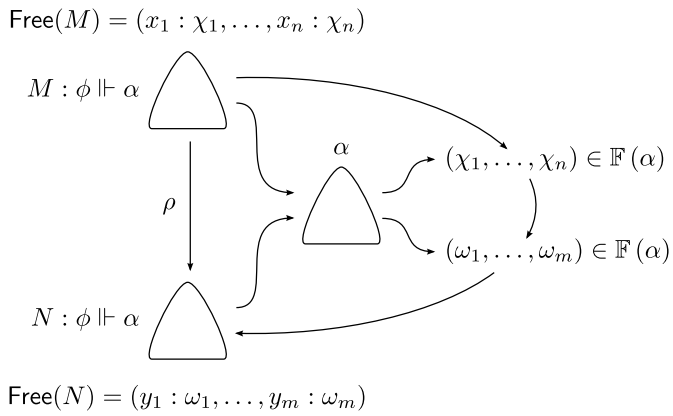

Fig. 8. A non-uniform renaming of the variables of M, based on an alternate extraction of the formulas of its blueprint.

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

Lemma 2.16. Let $M \in \Lambda_{\mathrm{NF}}$ be a term of blueprint α. Suppose

$$\alpha \rhd^{b_0^m}_{\omega_m} \ldots \rhd^{b_{p_m}^m}_{\omega_m} \quad \ldots \quad \rhd^{b_0^1}_{\omega_1} \ldots \rhd^{b_{p_1}^1}_{\omega_1} \emptyset_{\mathbb{B}}$$

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

Proof. By induction on M. The proposition is clear if M is a variable. The case of $M = (M_1 M_2)$ follows easily from the induction hypothesis. Suppose $M = \lambda x.M_1 : \phi \to \psi$ with $M_1 \Vdash \gamma$. Let $Y' = (y_1, \ldots, y_m, x)$. By Lemma 2.15.(2.a) there exist $a_1, \ldots, a_p$ such that $\{a_1, \ldots, a_p\} = \{a \mid M_{|a} = x\}$ and $\gamma \rhd^{a_0}_{\phi} \ldots \rhd^{a_p}_{\phi} \gamma' = \alpha_{|1}$. Now

$$\alpha \rhd^{b_0^m}_{\omega_m} \ldots \rhd^{b_{p_m}^m}_{\omega_m} \quad \ldots \quad \rhd^{b_0^1}_{\omega_1} \ldots \rhd^{b_{p_1}^1}_{\omega_1} \emptyset_{\mathbb{B}}$$

hence each b_j^i is of the form $(1) \cdot c_j^i$. Furthermore

$$\gamma \rhd^{a_0}_{\phi} \ldots \rhd^{a_p}_{\phi} \rhd^{c_0^m}_{\omega_m} \ldots \rhd^{c_{p_m}^m}_{\omega_m} \quad \ldots \quad \rhd^{c_0^1}_{\omega_1} \ldots \rhd^{c_{p_1}^1}_{\omega_1} \emptyset_{\mathbb{B}}$$

By induction hypothesis there exists N_1 with the same domain, blueprint and of the same type as M_1 such that $\mathsf{Free}(N_1) = Y'$, $\{a \mid N_{1|a} = x\} = \{a_0, \ldots, a_p\}$ and $\{c \mid N_{1|c} = y_i\} = \{c_0^i, \ldots, c_{p_i}^i\}$ for each i. By Lemma 2.15.(2.b) we have $\lambda x.N_1 \Vdash \alpha$, hence we may take $N = \lambda x.N_1$. □

3. Vertical compressions and compact terms

The aim of this section is to provide a partial characterisation of minimal inhabitants. Section 3.1 is just a simple remark on the relative depths of their blueprints, and an easy

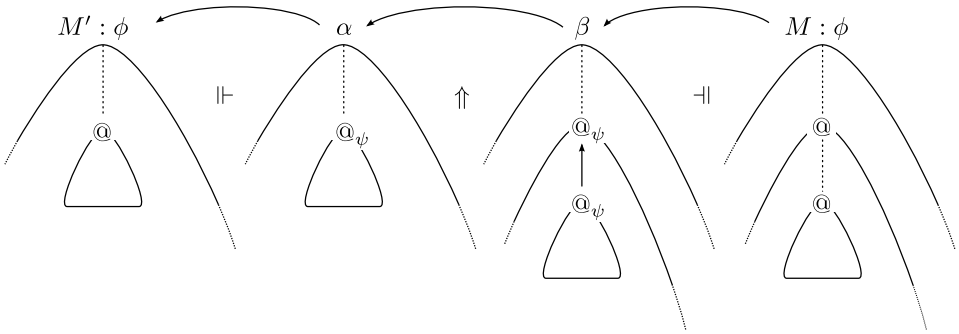

Fig. 9. How the compression of terms is able to follow the compression of blueprints.

consequence of the subformula property (Lemma 1.5): if M is a minimal Λ_{NF}-inhabitant of ϕ, then for all addresses a in M the blueprint of $M_{|a}$ is of 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 call *locally compact* every Λ_{NF}-inhabitant satisfying this condition. In Section 3.2 we introduce the notion of *vertical compression* of a blueprint. A (strict) vertical compression of β is obtained by taking any address b in β, then by 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 of more times. The key property of those compressions is the following (see Figure 9):

— If M is of 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 of the same type as M. Moreover M' and M are applications of the same type or abstractions of the same type.

Let us again consider a Λ_{NF}-inhabitant M and two addresses a, b such that $a < b$, $M_{|a}$ and $M_{|b}$ are applications of the same type or abstractions of the same type. Suppose:

— there exists a vertical compression α' of the blueprint of $M_{|b}$ such that the sequence $\Omega(\mathsf{Free}(M_{|a}))$ can be extracted from α'.

This situation is a generalisation of the last example in the introduction of Section 2 (in which α' was equal to the blueprint of $M_{|b}$, 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 of the same type or abstractions of the same type. By Lemma 2.16, there exists a term P of the same type and with the same domain as N such that $\mathsf{Free}(P) = \mathsf{Free}(M_{|a})$. The graft of P at a yields an inhabitant of strictly smaller size.

We will call *compact* all inhabitants in which the preceding situation does not occur. All inhabitants of minimal size are of course compact. As we shall see in Section 5, 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 in the exhaustive computation of their domains.

3.1. *Depths of the blueprints of minimal inhabitants*

Definition 3.1. Two terms $M, M' \in \Lambda_{\mathrm{NF}}$ are *of the same kind* if and only if they are both variables, or both applications, or both abstractions, and if they are of the same type.

Definition 3.2. For all formulas ϕ, we write $\mathsf{Sub}(\phi)$ for the set of all subformulas of ϕ.

Definition 3.3. Let $M \in \Lambda_{\mathrm{NF}}$. Let a be any address in M. Let $(a_1, \ldots, a_m)$ be the strictly increasing sequence of all prefixes of a. Let $(\lambda x_1, \ldots, \lambda x_k)$ be the subsequence of $(M(a_1), \ldots, M(a_m))$ consisting of all labels of the form λx. We write $\Lambda(M, a)$ for $(x_1, \ldots, x_k)$.

Definition 3.4. 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 $M_{|a}$ is of relative depth at most $|\Lambda(M, a)| \times |\mathsf{Sub}(\phi)|$.

Lemma 3.5. 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 $\mathsf{Free}(M_{|b}) = \mathsf{Free}(M_{|b'})$. Moreover, M is not a Λ_{NF}-inhabitant of ϕ of minimal size.

Proof. For each address a in $\mathrm{dom}(M)$, let α_a be the blueprint of $M_{|a}$ and let $X_a = \mathsf{Free}(M_{|a})$. Assume the existence of an α_a of relative depth $n > |\Lambda(M, a)| \times |\mathsf{Sub}(\phi)|$. There exist $b_1, \ldots, b_{n+1} \in \mathrm{dom}(\alpha_a)$ such that $b_1 < \ldots < b_n < b_{n+1}$. By Lemma 2.8.(1) we have $X_{a \cdot b_n} \subseteq \ldots \subseteq X_{a \cdot b_1} \subseteq \Lambda(M, a)$. By Lemma 1.5, each $\phi_{a \cdot b_i}$ is a subformula of ϕ. Hence there exist i, j such that $i < j$ and $(X_{a \cdot b_i}, \phi_{a \cdot b_i}) = (X_{a \cdot b_j}, \phi_{a \cdot b_j})$, that is, $M_{|a \cdot b_i}$ and $M_{|a \cdot b_j}$ are applications of the same type and with the same free variables (Figure 10). Now, let $M' = M[a \cdot b_i \leftarrow M_{|a \cdot b_j}]$. The term M' is a Λ_{NF}-inhabitant of ϕ of strictly smaller size. $\square$

3.2. *Vertical compression of a blueprint*

Definition 3.6. We let $\Uparrow$ be the least reflexive and transitive binary relation on blueprints satisfying the following: if $a, b \in \mathrm{dom}(\beta)$, $a < b$ and $\beta(a) = \beta(b)$, then $\beta[a \leftarrow \beta_{|b}] \Uparrow \beta$.

Lemma 3.7. Suppose $M \in \Lambda_{\mathrm{NF}}$, $M : \phi$, $M \Vdash \beta$ and $\alpha \Uparrow \beta$. There exists a term $M' \in \Lambda_{\mathrm{NF}}$ of the same kind as M, of blueprint α and such that $|\mathrm{dom}(M')| \leq |\mathrm{dom}(M)|$.

Proof. It suffices to consider the case of $\alpha = \beta[a \leftarrow \beta_{|b}]$ with $a, b \in \mathrm{dom}(\beta)$, $a < b$ and $\beta(a) = \beta(b)$. We prove the existence of M' by induction on the length of a. If $a = \varepsilon$ then M is necessarily an application and $\beta(\varepsilon) = \beta(b) = @_\phi$, hence $M_{|b}$ is an application of type ϕ, and we can take $M' = M_{|b}$. Assume $a \neq \varepsilon$.

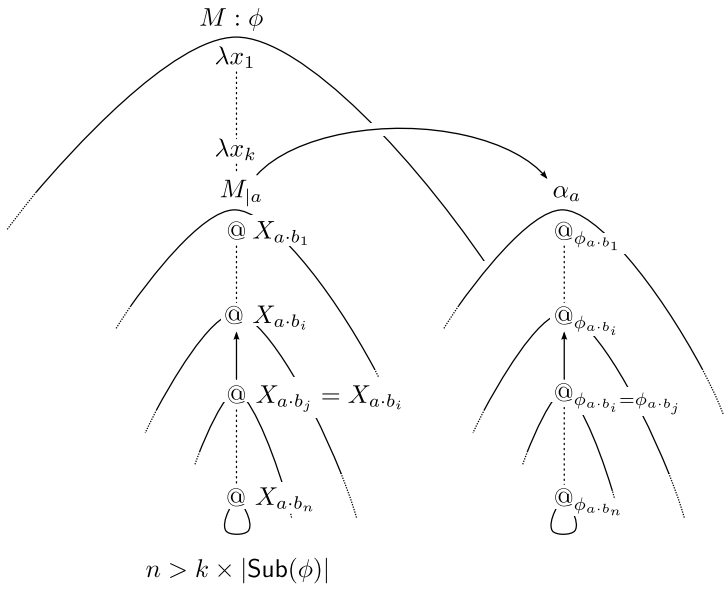

Fig. 10. Proof of Lemma 3.5.

(1) Suppose $M = (M_1 M_2)$, $M_1 \Vdash \beta_1$, $M_2 \Vdash \beta_2$, $a = (i)\cdot a_i$ and $b = (i)\cdot b_i$. By induction hypothesis there exists M_i' of blueprint $\alpha_i = \beta_i[a_i \leftarrow \beta_{i|b_i}] = \beta_i[a_i \leftarrow \beta_{|b}]$, of the same kind as M_i and such that $\mathrm{dom}(M_i') \leq \mathrm{dom}(M_i)$. Let $j = 1$ if $i = 2$, otherwise let $j = 2$. Let $(M_j', \alpha_j) = (M_j, \beta_j)$. Let $X = (x_1, \ldots, x_n)$ be the strictly increasing sequence of all variables free or bound in M_2'. Let $Y = (y_1, \ldots, y_n)$ be a strictly increasing sequence of variables such that $\Omega(X) = \Omega(Y)$ and y_1 is greater that 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 $M' = (M_1' M_2'')$.

(2) Suppose $M = \lambda x.M_1$, $M_1 \Vdash \beta_1$, $x : \chi$, $a = (1)\cdot a_1$ and $b = (1)\cdot b_1$. As $a, b \in \mathrm{dom}(\beta)$, we have also $a_1, b_1 \in \mathrm{dom}(\beta_1)$. By induction hypothesis there exists M_1' of the same kind as M_1, of blueprint $\alpha_1 = \beta_1[a_1 \leftarrow \beta_{1|b_1}]$ and such that $\mathrm{dom}(M_1') \leq \mathrm{dom}(M_1)$. By Lemma 2.15.(2.a) there exist $\gamma_1, c_0, \ldots, c_p$ such that $\{c_0, \ldots, c_p\} = \{c \mid M_{|c} = x\}$, $\beta_1 \rhd_\chi^{c_0} \ldots \rhd_\chi^{c_p} \gamma_1$ and $\beta = *(\gamma_1)$. Since $a, b \in \mathrm{dom}(\alpha)$, a_1 and c_i are incomparable addresses for all i. Hence $\alpha_1 = \beta_1[a_1 \leftarrow \beta_{1|b_1}] \rhd_\chi^{c_0} \ldots \rhd_\chi^{c_p} \gamma_1[a_1 \leftarrow \beta_{1|b_1}] = \beta[a \leftarrow \beta_{|b}]_{|(1)} = \alpha_{|1}$. By Lemma 2.16 there exists a term M_1'' of the same type and with the same domain as M_1' such that the greatest variable y free in M_1'' is of type χ and $\{c \mid M_{1\,|c}'' = y\} = \{c_0, \ldots, c_p\}$. By Lemma 2.15.(2.b) we have $\lambda y.M_1'' \Vdash \alpha$, hence we may take $M' = \lambda y.M_1''$. □

Definition 3.8. A term $M \in \Lambda_{\mathrm{NF}}$ is *compact* when there are no a, b, α' such that $a < b$, $M_{|a}$ and $M_{|b}$ are of the same kind, $M_{|b} \Vdash \alpha_b$, $\alpha' \Uparrow \alpha_b$ and $\Omega(\mathsf{Free}(M_{|a})) \in \mathbb{F}(\alpha')$.

Lemma 3.9. Every Λ_{NF}-inhabitant of minimal size is compact. Every compact Λ_{NF}-inhabitant of ϕ is locally compact.

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

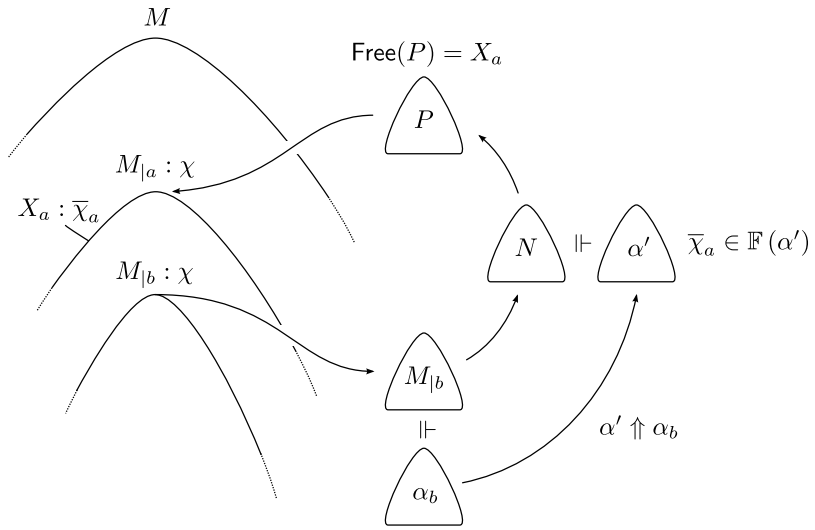

Fig. 11. Proof of Lemma 3.9, part (1).

(1) Assume M is not compact. Let a, b be such that $a < b$, $M_{|a}$ and $M_{|b}$ are of the same kind, $M_{|b} \Vdash \alpha_b$, $\alpha' \Uparrow \alpha_b$, $\mathsf{Free}(M_{|a}) = X_a$ and $\Omega(X_a) \in \mathbb{F}(\alpha')$ (see Figure 11). By Lemma 3.7 there exists a term $N \in \Lambda_{\mathrm{NF}}$ of blueprint α', of the same kind as $M_{|b}$ and such that $|\mathrm{dom}(N)| \leq |\mathrm{dom}(M_{|b})|$. By Lemma 2.16 there exists $P \in \Lambda_{\mathrm{NF}}$ of blueprint α', of the same kind as N, such that $\mathrm{dom}(P) = \mathrm{dom}(N)$ and $\mathsf{Free}(P) = X_a$. 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 3.5. Let $\alpha_{b'}$ be the blueprint of $M_{|b'}$. By Lemma 2.15.(3) we have $\Omega(\mathsf{Free}(M_{|b})) = \Omega(\mathsf{Free}(M_{|b'}) \in \mathbb{F}(\alpha_{b'})$. Since the relation $\Uparrow$ is reflexive, M is not compact. $\square$

4. Shadows

So far we have isolated two properties shared by all minimal inhabitants (Lemma 3.9). We shall now exploit these properties so as to design a decision method for the inhabitation problem.

In Section 4.1 and 4.2 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 $(\overline{\chi}_a, \gamma_a, \phi_a)$ where ϕ_a is the type of $M_{|a}$, the sequence $\overline{\chi}_a$ is $\Omega(\mathsf{Free}(M_{|a}))$, and γ_a is a "transversal compression" of the blueprint α_a of $M_{|a}$ (Definitions 4.1 and 4.2). Recall that $\overline{\chi}_a \in \mathbb{F}(\alpha_a)$ (by Lemma 2.15.(3)). The blueprint γ_a can be seen as a synthesized version of α_a of the same relative depth but of smaller "width", and such that $\overline{\chi}_a \in \mathbb{F}(\gamma_a) \subseteq \mathbb{F}(\alpha_a)$.

Each tree prefix of the shadow of M belongs to a finite set effectively computable from ϕ and the domain of this prefix. In particular, one can compute all possible values for

its labels, regardless of the full knowledge of M – or even without the 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) only from the knowledge of $\overline{\chi}_a, \phi_a, \gamma_b, \phi_b$ and the arity of the nodes at a and b. Indeed, suppose $a < b$, $\phi_a = \phi_b$ and the nodes at a, b are of the same arity (1, or 2). Now, assume:

— there exists $\gamma' \Uparrow \gamma_b$ such that $\overline{\chi}_a \in \mathbb{F}(\gamma')$.

Then $M_{|a}$ and $M_{|b}$ are of the same kind and there exists $\alpha' \Uparrow \alpha_b$ such that $\overline{\chi}_a = \Omega(\mathsf{Free}(M_{|a})) \in \mathbb{F}(\gamma') \subseteq \mathbb{F}(\alpha')$, therefore M is not compact.

In Section 4.2, what we call a *shadow* is merely a tree $a \mapsto (\overline{\chi}_a, \gamma_a, \phi_a)$ of a certain shape, no matter if 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 5 we will prove that for every formula ϕ, the set of shadows of compact inhabitants of ϕ is a finite set 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.

4.1. *Blueprint equivalence and transversal compression*

Definition 4.1. 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 $@_\phi(\alpha_1, \alpha_1) \equiv @_\phi(\beta_1, \beta_2)$,
4 if $|\overline{a}| = |\overline{b}| = n$ and $\alpha_i \equiv \beta_i$ for each $i \in [1, \ldots, n]$, then $*_{\overline{a}}(\alpha_1, \ldots, \alpha_n) \equiv *_{\overline{b}}(\beta_1, \ldots, \beta_n)$.

In (3), we assume $\alpha_1, \alpha_2, \beta_1, \beta_2$ non-empty. In (4), we assume that the elements of each sequence $\overline{a}$, $\overline{b}$ are pairwise incomparable addresses. As to avoid circularity we assume also $a \neq \varepsilon$ 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 $*_{\overline{a}}(\alpha_1, \ldots, \alpha_n) \equiv *(\alpha_1, \ldots, \alpha_n) \equiv *(\alpha_n, \ldots, \alpha_1)$, also $*(*(\alpha, \beta), \gamma) \equiv *(\alpha, \beta, \gamma) \equiv *(\alpha, *(\beta, \gamma))$, etc. It is easy to check that $\alpha \equiv \beta$ implies $\mathbb{F}(\alpha) = \mathbb{F}(\beta)$ – this property will be used without reference.

Definition 4.2. For each $m \in \mathbb{N}$, we let $\curvearrowleft_m$ be the least binary relation such that:

1 if $\gamma_1 \equiv \ldots \equiv \gamma_m \equiv \gamma_{m+1} \not\equiv \emptyset_{\mathbb{B}}$, then $*_{\overline{a}}(\gamma_1, \ldots, \gamma_m) \curvearrowleft_m *_{\overline{a}\cdot(b)}(\gamma_1, \ldots, \gamma_m, \gamma_{m+1})$,
2 if $\alpha = *_{\overline{a}}(\alpha_1, \ldots, \alpha_n)$, $\beta = *_{\overline{b}}(\beta_1, \ldots, \beta_p)$ and $\alpha \curvearrowleft_m \beta$, then:
   (a) $@_\phi(\alpha, \gamma) \curvearrowleft_m @_\phi(\beta, \gamma)$,
   (b) $@_\phi(\gamma, \alpha) \curvearrowleft_m @_\phi(\gamma, \beta)$,
   (c) $*_{\overline{a}\cdot(c)}(\alpha_1, \ldots, \alpha_n, \gamma) \curvearrowleft_m *_{\overline{b}\cdot(c)}(\beta_1, \ldots, \beta_p, \gamma)$.

We call *m-compression of* β every α such that $\alpha \curvearrowleft_m \beta$. The *width* of β is defined as the least $m \in \mathbb{N}$ for which there is no α such that $\alpha \curvearrowleft_m \beta$.

Again the elements of $\overline{a}\cdot(b)$, $\overline{a}\cdot(c)$ and $\overline{b}\cdot(c)$ must be pairwise incomparable addresses, and α,β,γ must be non-empty. Note that for all non-empty β, we have $\emptyset_{\mathbb{B}} \curvearrowleft_0 \beta$, hence the empty blueprint is the only blueprint of null width. If β is of width $m>0$, then for all addresses a, for $\beta_{|a} = *_{\overline{a}}(\gamma_1,\ldots,\gamma_k)$ and for each $\gamma_i \neq \emptyset_{\mathbb{B}}$, the sequence $(\gamma_1,\ldots,\gamma_k)$ contains no more than m blueprints $\equiv$-equivalent to γ_i. For instance, if ϕ,ψ,χ are distinct formulas, $*(\phi,\phi,\phi,\psi,\psi,\chi)$ is of width 3, $*(\omega,@_\omega(*(\phi,\psi),\phi),@_\omega(*(\psi,\phi),\phi))$ is of width 2, etc.

Definition 4.3. For each $m\in\mathbb{N}$, we write $\sqsubseteq_m$ for the reflexive and transitive closure of the union of $\equiv$ and $\curvearrowleft_m$. We let $\sqsubseteq_m^{\max}$ 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 $j\in[1,\ldots,m]$ and clearly, $\alpha \curvearrowleft_m \beta$ implies $|\mathrm{dom}(\alpha)| < |\mathrm{dom}(\beta)|$, therefore $\curvearrowleft_m$ is well-founded.

Definition 4.4. For all $\mathcal{S}\subseteq\mathfrak{S}$, for all $d\in\mathbb{N}$ and for all $m\in\mathbb{N}$:

— we let $\mathbb{B}(\mathcal{S},d,\infty)$ be the set of $\mathcal{S}$-blueprints of 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 4.5. For all finite $\mathcal{S}\subseteq\mathfrak{S}$, for all $d\in\mathbb{N}$ and for 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 $\mathbb{B}_\varepsilon(\mathcal{S},d,m)/_\equiv$ is a finite set and 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 $(\alpha_1,\ldots,\alpha_k)$ be an enumeration of $\mathbb{R}_\varepsilon(\mathcal{S},d,m)$. Let Σ_d be the set of all functions from $\{1,\ldots,k\}$ to $\{0,\ldots,m\}$. For each $\beta\in\mathbb{B}(\mathcal{S},d,m)$ there exist $\beta_1,\ldots,\beta_n\in\mathbb{B}_\varepsilon(\mathcal{S},d,m)$ and $\overline{b}$ such that $\beta = *_{\overline{b}}(\beta_1,\ldots,\beta_n)$. We let σ_β be the function mapping each $i\in\{1,\ldots,k\}$ to the number of occurrences of an element $\equiv$-equivalent to α_i in the sequence $(\beta_1,\ldots,\beta_n)$. 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 $\sigma_\beta=\sigma_{\beta'}$, hence $\mathbb{B}(\mathcal{S},d,m)$ is a finite set.

For each $\tau\in\Sigma_d$, let $\rho_\tau = *(\alpha_1^1,\ldots,\alpha_1^{\tau(1)},\ldots,\alpha_k^1,\ldots,\alpha_k^{\tau(k)})$ 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 $\{\rho_\tau \mid \tau\in\Sigma_d\}$.

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 lemma follows by induction on d, using (1) and the facts that: $\mathbb{B}_\varepsilon(\mathcal{S},0,0)$ is

empty (hence $\mathbb{B}(\mathcal{S}, d, 0) = \{\emptyset_{\mathbb{B}}\}$ for all d); if $m \in \mathbb{N}_+$, then $\mathbb{B}_\varepsilon(\mathcal{S}, 0, m)$ is the finite set of all formulas of $\mathcal{S}$. □

4.2. *Shadow of a term*

Definition 4.6. Let ϕ be a formula. Let $\mathcal{S}_\phi$ be the union of $\mathsf{Sub}(\phi)$ (Definition 3.2) and the set of all $@_\psi$ such that $\psi \in \mathsf{Sub}(\phi)$. For each integer k, for each formula ϕ, we let $\mathfrak{R}(\phi, k) = \mathbb{R}\,(\mathcal{S}_\phi, k \times |\mathsf{Sub}(\phi)|, k)$, where $\mathbb{R}$ is the function introduced in Lemma 4.5.(2).

Definition 4.7. A *shadow* is a finite tree in which each node is of arity at most 2 and is labelled with a triple of the form $(\overline{\chi}, \gamma, \psi)$, where $\overline{\chi}$ is a sequence of formulas, γ is a blueprint and ψ is a formula.

We call *ϕ-shadow* every shadow Ξ satisfying the following conditions. We have $\Xi(\varepsilon) = (\varepsilon, \emptyset_{\mathbb{B}}, \phi)$. For each $a \in \mathrm{dom}(\Xi)$, let k_a be the number of $b < a$ such that the node of Ξ at b is unary, and let $(\overline{\chi}_a, \gamma_a, \psi_a) = \Xi(a)$. Then:

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

Definition 4.8. Let M be a locally compact Λ_{NF}-inhabitant of ϕ. For each $a \in \mathrm{dom}(M)$:

— let $\overline{\chi}_a = \Omega(\mathsf{Free}(M_{|a}))$,
— let α_a be the blueprint of $M_{|a}$,
— let $\gamma_a \in \mathfrak{R}(\phi, |\Lambda(M, a)|)$ be such that $\gamma_a \sqsubseteq^{\max}_{|\Lambda(M,a)|} \alpha_a$,
— let ϕ_a be the type of $M_{|a}$.

The tree Ξ mapping each $a \in \mathrm{dom}(M)$ to $(\overline{\chi}_a.\gamma_a, \phi_a)$ will be called *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 $M_{|a}$ is of relative depth at most $|\Lambda(M, a)| \times |\mathsf{Sub}(\phi)|$. 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 possibly not unique (although it is, since $\mathbb{R}$ is a selector and one can actually prove that $\gamma \sqsubseteq^{\max}_m \alpha$ and $\gamma' \sqsubseteq^{\max}_m \alpha$ implies $\gamma \equiv \gamma'$, but this property is irrelevant to our discussion). We assume that *some* γ_a is chosen for each address a in M.

Obviously 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 satisfies also the third.

4.3. *Compact shadows and compact inhabitants*

Definition 4.9. A shadow Ξ is *compact* if and only if there are no a, b such that: $a < b$, the nodes of Ξ at a, b are of the same arity, $\Xi(a) = (\overline{\chi}_a, \gamma_a, \psi)$, $\Xi(b) = (\overline{\chi}_b, \gamma_b, \psi)$ and there exists $\gamma' \Uparrow \gamma_b$ such that $\overline{\chi}_a \in \mathbb{F}\,(\gamma')$.

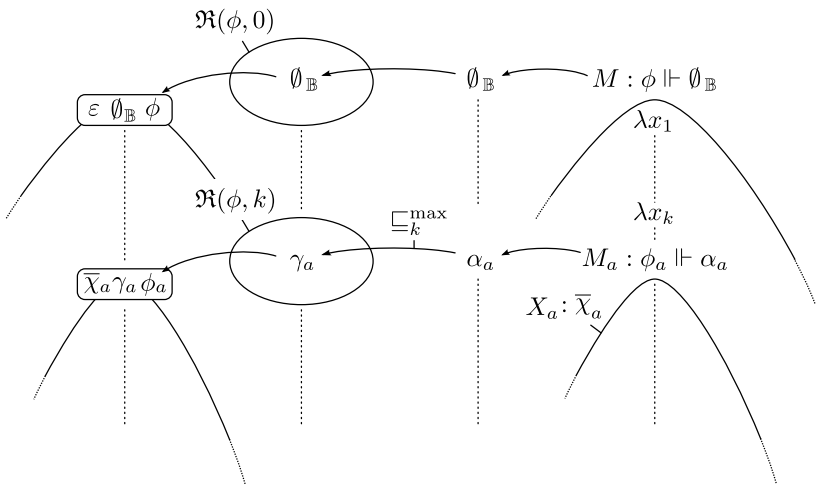

Fig. 12. A compact inhabitant and its shadow.

Compare this definition with the definition of compactness for term (Definition 3.8). With the help of three auxiliary lemmas, we now prove the key lemma of Section 4: if M is a compact inhabitant – a fortiori locally compact by Lemma 3.9 – then the shadow of M is a compact ϕ-shadow.

Lemma 4.10. 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 $|\mathrm{dom}(\beta')|$ 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 $|\mathrm{dom}(\beta')|$ shows that if $\alpha \Uparrow \beta \curvearrowleft_1 \beta'$, then there exists α' such that $\alpha \curvearrowleft_1 \alpha' \Uparrow \beta'$. The only non trivial case is $\alpha = *_{(a_1)}(\alpha_1)$, $\beta = *_{(a_1)}(\beta_1)$ with $\alpha_1 \Uparrow \beta_1$ and $\beta' = *_{(a_1,a_2)}(\beta_1, \beta_2)$ with $\beta_1 \equiv \beta_2$. 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 $\alpha = *_{(a_1)}(\alpha_1) \curvearrowleft_1 *_{(a_1,a_2)}(\alpha_1, \alpha_2) \Uparrow *_{(a_1,a_2)}(\beta_1, \beta_2) = \beta'$.

(3) Using (1) and (2), the lemma follows by induction on the length of an arbitrary sequence $(\beta_0, \ldots, \beta_n)$ such that $\beta_0 = \beta$, $\beta_n = \beta'$ and $\beta_{i-1} \equiv \beta_i$ or $\beta_{i-1} \curvearrowleft_1 \beta_i$ for each $i \in [1, \ldots, n]$. □

Lemma 4.11. If $\alpha \sqsubseteq_1 \beta$, then $\mathbb{F}(\alpha) \subseteq \mathbb{F}(\beta)$.

Proof. By induction on $|\mathrm{dom}(\beta)|$. Since $\gamma \equiv \gamma'$ implies $\mathbb{F}(\gamma) = \mathbb{F}(\gamma')$ and $|\mathrm{dom}(\gamma)| = |\mathrm{dom}(\gamma')|$, it suffices 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. The remaining cases follow easily from the induction hypothesis. □

Lemma 4.12. If $\alpha \sqsubseteq_m \beta$, then the set of all elements of $\mathbb{F}(\beta)$ of length at most m is a subset of $\mathbb{F}(\alpha)$.

Proof. By induction on $|\mathrm{dom}(\beta)|$. Again, we examine only the case $\alpha \curvearrowleft_m \beta$. The

proposition is trivially true if $m = 0$. Suppose $m > 0$. The only non-trivial case is $\alpha \equiv *_{\overline{a}}(\gamma_1, \ldots, \gamma_m)$ and $\beta \equiv *_{\overline{a}}(\gamma_1, \ldots, \gamma_m, \gamma_{m+1})$ with $\gamma_i \equiv \gamma$ for all i. Let $\Phi = \mathbb{F}(\gamma)$. For each integer k, let $\Phi^{(k)} = \circledast(\Phi_1, \ldots, \Phi_k)$ where $\Phi_i = \mathbb{F}(\gamma)$ for each i. Let $\overline{\phi} = (\phi_1, \ldots, \phi_p) \in \mathbb{F}(\beta)$ be such that $p \leq m$. We have to prove that $\overline{\phi} \in \mathbb{F}(\alpha)$. For each $J \subseteq \{1, \ldots, p\}$, let $(j_1, \ldots, j_q)$ be the strictly increasing enumeration of all elements of J and let $f(J) = (\phi_{j_1}, \ldots, \phi_{j_q})$. We have $\overline{\phi} \in \mathbb{F}(\beta) = \Phi^{(m+1)}$, hence there exist $J_1, \ldots, J_{m+1}$ such that $J_1 \cup \ldots \cup J_{m+1} = \{1, \ldots, p\}$, and $f(J_i) \in \mathbb{F}(\gamma)$ for each $i \in \{1, \ldots, m+1\}$. For each $j \in \{1, \ldots, p\}$, let k_j be any element of $\{1, \ldots, m+1\}$ such that $j \in J_{k_j}$. Then $J_{k_1} \cup \ldots \cup J_{k_p} = \{1, \ldots, p\}$, so $\overline{\phi} \in \circledast(\{f(J_{k_1})\}, \ldots, \{f(J_{k_p})\}) \subseteq \Phi^{(p)} \subseteq \Phi^{(m)} = \mathbb{F}(\alpha)$. $\square$

Lemma 4.13. Let M be a locally compact Λ_{NF}-inhabitant of ϕ. The shadow of M is a ϕ-shadow. If M is compact, then this shadow is also compact.

Proof. For each address a in M, the sequence $\overline{\chi}_a = \Omega(\mathsf{Free}(M_{|a}))$ is a subsequence of $\Omega(\Lambda(M, a))$, hence the first proposition follows from the definition of the shadow of M, Lemma 1.5, Lemma 2.15.(3) and Lemma 4.12. Let Ξ be shadow of M. Assume Ξ is not compact. There exist $a, b \in \mathrm{dom}(\Xi) = \mathrm{dom}(M)$ such that $\Xi(a) = (\overline{\chi}_a, \gamma_a, \psi)$, $\Xi(b) = (\overline{\chi}_b, \gamma_b, \psi)$, the nodes at a,b in Ξ are of the same arity, and there exists $\gamma' \Uparrow \gamma_b$ such that $\overline{\chi}_a \in \mathbb{F}(\gamma')$. We have $M_{|a}$, $M_{|b}$ of the same kind. Let α_a, α_b be the blueprints of $M_{|a}$, $M_{|b}$. Since $\gamma_b \sqsubseteq^{\max}_{|\Lambda(M,a\cdot b)|} \alpha_b$, we have $\gamma' \Uparrow \gamma_b \sqsubseteq_1 \alpha_b$. By Lemma 4.10 there exists α' such that $\gamma' \sqsubseteq_1 \alpha' \Uparrow \alpha_b$. By Lemma 4.11, we have $\overline{\chi}_a \in \mathbb{F}(\gamma') \subseteq \mathbb{F}(\alpha')$, hence M is not compact. $\square$

5. Finiteness of the set of compact ϕ-shadows

Our last aim will be to prove that for each formula ϕ, the set of all compact ϕ-shadows is a finite set effectively computable from ϕ.

In definition 5.1, we introduce a last binary relation $\Subset$ on blueprints. The key lemma of this section (Lemma 5.14) 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 $\Subset$ is an almost full relation (Bezem, Klop and de Vrijer 2003) on the set of all $\mathcal{S}$-blueprints: for every infinite sequence $\gamma_1, \gamma_2, \ldots$ over $\mathbb{B}(\mathcal{S})$, there exists i, j such that $i < j$ and $\gamma_i \Subset \gamma_j$. This result will be proven with the help of Melliès' Axiomatic Kruskal Theorem (Melliès 1998). The finiteness of the set of compact ϕ-shadows follows from this key lemma with the help of König's Lemma (Lemma 5.15). The ability to compute these shadows follows directly from their definition.

By Lemma 4.13, a consequence of this result is also 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 5.13 is non-constructive and that it gives no information about the complexity of our proof-search method – this question might be itself another open problem.

5.1. *Almost full relations and Higman Theorem*

Definition 5.1. We let $\Subset$ be the relation on blueprints defined by $\alpha \Subset \beta$ if and only if for all $\overline{\chi} \in \mathbb{F}(\alpha)$, there exists $\gamma \Uparrow \beta$ such that $\overline{\chi} \in \mathbb{F}(\gamma)$.

Definition 5.2. 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 5 will be to prove the last key lemma from which we will easily infer the decidability of Λ_{NF}-inhabitation: for each finite $\mathcal{S} \subseteq \mathfrak{S}$, the relation $\Subset$ is an AFR on $\mathbb{B}(\mathcal{S})$.

Proposition 5.3.

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

Definition 5.4. Let $\mathcal{U}$ be a set, let $\ll$ be a binary relation. We let $\mathbb{S}(\mathcal{U})$ denote the set of all finite sequences over $\mathcal{U}$. The relation $\ll_{\mathbb{S}}$ *induced by* $\ll$ *on* $\mathbb{S}(\mathcal{U})$ is defined by $(U_1, \ldots, U_n) \ll_{\mathbb{S}} (V_1, \ldots, V_m)$ if and only if there exists a strictly monotone function $\eta : \{1, \ldots, n\} \to \{1, \ldots, m\}$ such that $U_i \ll V_{\eta(i)}$ for each $i \in \{1, \ldots, n\}$.

Theorem 5.5. (Higman) If $\ll$ is an AFR on $\mathcal{U}$, then $\ll_{\mathbb{S}}$ is an AFR on $\mathbb{S}(\mathcal{U})$.

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

5.2. *From rooted to unrooted blueprints*

Melliès' Axiomatic Kruskal Theorem allows one to conclude that a relation is an AFR (a "well binary relation" in (Melliès 1998)) as long as it satisfies a set of five properties or "axioms" (six in the original version of the theorem – see the remarks of Melliès at the end of its proof explaining why five axioms suffice). The details of those axioms will be given in Section 5.3.

Four of those five axioms are relatively easy to check. The remaining axiom is more problematical. This rather technical section is entirely devoted to the proof of Lemma 5.11, which will ensure that this last axiom is satisfied. We want to prove the following proposition:

Let $\mathcal{S}$ be a finite subset of $\mathfrak{S}$. Let $\mathcal{B}_\varepsilon$ be a subset of $\mathbb{B}_\varepsilon(\mathcal{S})$.
Let $\mathcal{B} = \{*_{\overline{a}}(\beta_1, \ldots, \beta_n) \mid \forall i \in [1, \ldots, n], \beta_i \in \mathcal{B}_\varepsilon\}$.
If $\Subset$ is an AFR on $\mathcal{B}_\varepsilon$, then $\Subset$ is an AFR on $\mathcal{B}$.

Recall that $\mathbb{B}_\varepsilon(\mathcal{S})$ stands for the set of all rooted $\mathcal{S}$-blueprints. We want to be able to

extend the property that $\Subset$ is an AFR on a given set of rooted blueprints to the set all blueprints that have those rooted blueprints at their minimal addresses.

Higman Theorem suffices to show that $\Subset_{\mathbb{S}}$ (Definition 5.4) is an AFR on the set of finite sequences over $\mathcal{B}_\varepsilon$. However, if one considers an infinite sequence $(\beta_i)_{i\in\mathbb{N}}$ over $\mathcal{B}$ and transforms each $\beta_i = *_{\overline{a}_i}(\beta^i_1, \ldots, \beta^i_{n_i})$ where $\beta^i_1, \ldots, \beta^i_{n_i} \in \mathcal{B}_\varepsilon$ into $\sigma(\beta_i) = (\beta^i_1, \ldots, \beta^i_{n_i})$, the theorem will only provide two integers i, j and strictly monotone function η such that $i < j$ and $\beta^i_k \Subset \beta^j_{\eta(k)}$ for each $k \in \{1, \ldots, n_i\}$. This is sufficient to ensure that $\beta_i = *_{\overline{a}_i}(\beta^i_1, \ldots, \beta^i_{n_i}) \Subset *_{\overline{b}}(\beta^j_{\eta(1)}, \ldots, \beta^j_{\eta(n_i)})$, but not in general $\beta_i \Subset \beta_j$.

To bypass this difficulty we show how for each blueprint $\beta \in \mathbb{B}(\mathcal{S})$, one can extract from the set of all vertical compressions of β a complete set of "followers" of β of minimal size (Lemma 5.7). This set $\{\alpha_1, \ldots, \alpha_p\}$ has the property that for each $\overline{\phi} \in \mathbb{F}(\beta)$, there exists at least one α_i such that $\mathbb{F}(\alpha_i)$ contains a *subsequence* of $\overline{\phi}$ – but not necessarily $\overline{\phi}$ itself. The relative depth of each α_i does not depend on the relative depth on β, but only on $\mathcal{S}$: it is at most $\Sigma_{i=1}^{1+|\mathcal{S}_@|}\, i$, where $\mathcal{S}_@$ is the set of all binary symbols in $\mathcal{S}$. The lemma in proven in four steps.

First, observe that the set of all $\alpha \Uparrow \beta$ of relative depth at most $\Sigma_{i=1}^{1+|\mathcal{S}_@|}\, i$ 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.

Second, 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 $\overline{\phi} \in \mathbb{F}(\beta)$ there exists an $\mathcal{S}$-residual γ of β such that $\mathbb{F}(\gamma)$ contains a subsequence of $\overline{\phi}$ (Lemma 5.9).

Third, we prove that if $\beta = *_{\overline{a}}(\beta_1, \ldots, \beta_n)$, $\beta' = *_{\overline{b}}(\beta'_1, \ldots, \beta'_n, \beta'_{n+1}, \ldots, \beta'_{n+k})$ are such that $\beta_i \Subset \beta'_i$ for each $i \in [1, \ldots, n]$, and if furthermore β, β' have the same set of $\mathcal{S}$-residuals, then $\beta \Subset \beta'$ (Lemma 5.10).

The last step is the proof of the lemma itself. The set of $\mathcal{S}$-residuals is finite up to $\equiv$ (Lemma 4.5), 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 follows from the third step and Higman Theorem.

Definition 5.6. For every $\mathcal{S} \subseteq \mathfrak{S}$, we let $\mathcal{S}_@$ denote the set of all binary symbols in $\mathcal{S}$.

Lemma 5.7. Let $\mathcal{S}$ be a finite subset of $\mathfrak{S}$. For all $\beta \in \mathbb{B}(\mathcal{S})$, for all $\overline{\psi} \in \mathbb{F}(\beta)$, there exists α of relative depth at most $\Sigma_{i=1}^{1+|\mathcal{S}_@|}\, i$ such that $\alpha \Uparrow \beta$ and such that $\mathbb{F}(\alpha)$ contains a subsequence of $\overline{\psi}$.

Proof. Call $\mathcal{S}$*-linearisation* every pair $(\gamma, \overline{\chi})$ such that $\gamma \in \mathbb{B}(\mathcal{S})$ and $\overline{\chi} \in \mathbb{F}(\gamma)$. Call *starting address for* $(\gamma, \overline{\chi})$ every address b for which there exist ϕ, γ' such that $\gamma \rhd^b_\phi \gamma'$ and $\overline{\chi} \in \circledcirc(\mathbb{F}(\gamma'), (\phi))$. Call *path to* b *in* γ the maximal sequence $(b_1, \ldots, b_n, b_{n+1})$ over $\mathrm{dom}(\gamma)$ such that $b_1 < \ldots < b_n < b_{n+1} = b$.

Given an arbitrary $\mathcal{S}$-linearisation $(\beta, \overline{\psi})$, we prove simultaneously by induction on $|\mathrm{dom}(\beta)|$ the following properties:

1 There exists an $\mathcal{S}$-linearisation $(\gamma, \overline{\chi})$ such that:

(a) $\gamma \Uparrow \beta$ and $\overline{\chi}$ is a subsequence of $\overline{\psi}$,

(b) γ is of relative depth at most $1 + \Sigma_{i=1}^{|\mathcal{S}_{@}|} i$.

2 There exists an $\mathcal{S}$-linearisation $(\alpha, \overline{\phi})$ such that:

(a) $\alpha \Uparrow \beta$, $\overline{\phi}$ is a subsequence of $\overline{\psi}$,
and if $\psi \neq \varepsilon$, then the last elements of $\overline{\phi}, \overline{\psi}$ are equal,

(b) for each starting address b for $(\alpha, \overline{\phi})$ and for $(b_1, \ldots, b_n, b_{n+1})$ equal to the path to b in α, the values $\alpha(b_1), \ldots, \alpha(b_n)$ are pairwise distinct,

(c) for all c incomparable with each starting address for $(\alpha, \overline{\phi})$,
$(\alpha_{|c})$ is of relative depth at most $1 + \Sigma_{i=1}^{|\mathcal{S}_{@}|} i$.

Note that the conjunction of (2.b) and (2.c) implies that every address d in α is of relative depth at most $|S_{@}|+1+\Sigma_{i=1}^{|\mathcal{S}_{@}|} i = \Sigma_{i=1}^{1+|\mathcal{S}_{@}|} i$. Indeed, suppose d is of maximal relative depth and not a starting address for $(\alpha, \overline{\phi})$. Then d must be incomparable with each starting address for $(\alpha, \overline{\phi})$. Let e be the shortest prefix of d in $\mathrm{dom}(\alpha)$ that is incomparable with each starting address for $(\alpha, \overline{\phi})$. The address e is of relative depth at most $|\mathcal{S}_{@}|$ in α – otherwise there would exist in $\mathrm{dom}(\alpha)$ an address $f < e$ of relative depth $|\mathcal{S}_{@}|$ and a starting adress for $(\alpha, \overline{\phi})$ of the form $f \cdot f'$, of relative depth strictly greater than $|\mathcal{S}_{\alpha}|$, 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 cases $\beta = \emptyset_{\mathbb{B}}$ is immediate. If $\beta = *_{\overline{\alpha}}(\beta_1, \ldots, \beta_n)$, $i \neq j$ and $\beta_i, \beta_j \neq \emptyset_{\mathbb{B}}$, then the conclusion follows easily from the induction hypothesis. Suppose $\beta = @_{\psi}(\beta_1, \beta_2)$.

(1) Let d be an address of maximal length in $\beta^{-1}(@_{\psi})$. Let $\delta = @_{\psi}(\delta_1, \delta_2) = \beta_{|d}$. By assumption ε is the only element of $\delta^{-1}(@_{\psi})$. As $\overline{\psi} \in \mathbb{F}(\beta)$, there exist $\overline{\psi}_0 \in \mathbb{F}(\delta)$, $\overline{\psi}_1 \in \mathbb{F}(\delta_1)$, $\overline{\psi}_2 \in \mathbb{F}(\delta_2)$ such that $\overline{\psi}_0$ is a subsequence $\overline{\psi}$ and $\overline{\psi}_0 \in \circledcirc(\{\overline{\psi}_1\}, \{\overline{\psi}_2\})$. By induction hypothesis there exists an $(\mathcal{S} - \{@_{\psi}\})$-linearisation $(\gamma_1, \overline{\chi}_1)$ satisfying conditions (1.a), (1.b) w.r.t $(\delta_1, \overline{\psi}_1)$, and an $(\mathcal{S} - \{@_{\psi}\})$-linearisation $(\gamma_2, \overline{\chi}_2)$ satisfying conditions (2.a), (2.b), (2.c) w.r.t $(\delta_2, \overline{\psi}_2)$. Let $\gamma = @_{\psi}(\gamma_1, \gamma_2)$. We have $\gamma \Uparrow \delta$ and $\beta(\varepsilon) = \delta(\varepsilon) = \gamma(\varepsilon)$, hence $\gamma \Uparrow \beta$. The blueprint γ_1 is of relative depth at most $1 + \Sigma_{i=1}^{|\mathcal{S}_{@}|-1} i \leq \Sigma_{i=1}^{|\mathcal{S}_{@}|} i$. The blueprint γ_2 is of relative depth at most $|\mathcal{S}_{@}|+\Sigma_{i=1}^{|\mathcal{S}_{@}|-1} = \Sigma_{i=1}^{|\mathcal{S}_{@}|} i$. Therefore γ is of relative depth at most $1 + \Sigma_{i=1}^{|\mathcal{S}_{@}|} i$. Now $\overline{\chi}_2$ is a subsequence of $\overline{\psi}_2$ with the same last element, so there exists in $\circledcirc(\{\overline{\chi}_1\}, \{\overline{\chi}_2\}) \subseteq \mathbb{F}(@_{\psi}(\gamma_1, \gamma_2))$ a subsequence $\overline{\chi}$ of $\overline{\psi}_0$. Thus $(\gamma, \overline{\chi})$ satisfies (1.a) and (1.b) w.r.t $(\beta, \overline{\psi})$.

(2) As $\overline{\psi} \in \mathbb{F}(\beta)$, there exist $\overline{\psi}_1 \in \mathbb{F}(\beta_1), \overline{\psi}_2 \in \mathbb{F}(\beta_2)$ such that $\overline{\psi} \in \circledcirc(\{\overline{\psi}_1\}, \{\overline{\psi}_2\})$. By induction hypothesis there exists an $\mathcal{S}$-linearisation $(\alpha_1, \overline{\phi}_1)$ satisfying conditions (1.a), (1.b) w.r.t $(\beta_1, \overline{\psi}_1)$, and an $\mathcal{S}$-linearisation $(\alpha_2, \overline{\phi}_2)$ satisfying conditions (2.a), (2.b), (2.c) w.r.t $(\beta_2, \overline{\psi}_2)$.

Let $\alpha_0 = @_{\psi}(\alpha_1, \alpha_2)$. We have $\alpha_0 \Uparrow \beta$. The last elements of $\overline{\phi}_2$, $\overline{\psi}_2$ are equal and $\circledcirc(\{\overline{\phi}_1\}, \{\overline{\phi}_2\}) \subseteq \mathbb{F}(\alpha_0)$. Hence there exists in $\mathbb{F}(\alpha_0)$ a subsequence $\overline{\phi}_0$ of $\overline{\psi}$ with the same last element as $\overline{\psi}$. Thus $(\alpha_0, \overline{\phi}_0)$ satisfies (2.a).

For all c incomparable with each starting address for $(\alpha_0, \overline{\phi}_0)$, either $c = (1) \cdot c'$ and $c' \in \mathrm{dom}(\alpha_1)$, or $c = (2) \cdot c''$ and $c'' \in \mathrm{dom}(\alpha_2)$ is incomparable with each starting address in α_2. As a consequence, the choice of α_1, α_2 ensures that $(\alpha_0, \overline{\phi}_0)$ satisfies (2.c).

If $(\alpha_0, \overline{\phi}_0)$ satisfies (2.b), then we may take $(\alpha, \overline{\phi}) = (\alpha_0, \overline{\phi}_0)$. Otherwise some starting

address b for $(\alpha_0, \overline{\phi}_0)$ does not satisfy condition (2.b). Let $(b_1, \ldots, b_n, b_{n+1})$ 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 $(d_2, \ldots, d_{n+1})$ is then a path to $d = d_{n+1}$ in α_2, and d is a starting address for $(\alpha_2, \overline{\phi}_2)$. The values $\alpha_2(d_2), \ldots, \alpha_2(d_n)$ are pairwise distinct, so there must exist $i > 1$ such that $\alpha(b_i) = @_\psi$. Since b_i is in the path to b, there exists in $\mathbb{F}(\alpha_{2|d_i})$ a subsequence $\overline{\phi}'_0$ of $\overline{\phi}_0$ with the same last element as $\overline{\phi}_0$. For $\alpha'_0 = \alpha_0[\varepsilon \leftarrow \alpha_{2|d_i}]$, we have $\alpha'_0 \Uparrow \beta$, $\overline{\phi}'_0 \in \mathbb{F}(\alpha'_0)$ and the last elements of $\overline{\phi}'_0, \overline{\phi}_0, \overline{\psi}$ are equal. By induction hypothesis there exists an $\mathcal{S}$-linearisation $(\alpha, \overline{\phi})$ satisfying (2.a), (2.b), (2.c) w.r.t $(\alpha'_0, \overline{\phi}'_0)$. The pair $(\alpha, \overline{\phi})$ satisfies also those conditions w.r.t $(\beta, \overline{\psi})$. □

Definition 5.8. Let $\mathcal{S}$ be a finite subset of $\mathfrak{S}$. For all $\beta \in \mathbb{B}(\mathcal{S})$, for all $\alpha \Uparrow \beta$ of relative depth at most $\Sigma_{i=1}^{1+|\mathcal{S}_@|} i$, we call *$\mathcal{S}$-residual* of β every α_0 such that $\alpha_0 \sqsubseteq_1^{\max} \alpha$.

Note that the set of $\mathcal{S}$-residuals of β is $\{\emptyset_{\mathbb{B}}\}$ if $\beta = \emptyset_{\mathbb{B}}$. 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 4.3).

Lemma 5.9. Let $\mathcal{S}$ be a finite subset of $\mathfrak{S}$. For all $\beta \in \mathbb{B}(\mathcal{S})$ and for all $\overline{\psi} \in \mathbb{F}(\beta)$, there exists an $\mathcal{S}$-residual α_0 of β such that $\mathbb{F}(\alpha_0)$ contains a subsequence of $\overline{\psi}$.

Proof. (1) Let γ, δ be arbitrary blueprints. Suppose $\gamma \curvearrowleft_1 \delta$. We prove by induction on δ that for all $\overline{\phi} \in \mathbb{F}(\delta)$, there exists in $\mathbb{F}(\gamma)$ a subsequence of $\overline{\phi}$. In order to deal with the case $\delta = @_\phi(\delta_1, \delta_2)$, we need to prove a slightly more precise property: for all $\overline{\phi} \in \mathbb{F}(\delta)$, there exists in $\mathbb{F}(\gamma)$ a subsequence $\overline{\psi}$ of $\overline{\phi}$ such that the last elements of $\overline{\phi}$, $\overline{\psi}$ 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 prove the lemma. By Lemma 5.7 and by definition of an $\mathcal{S}$-residual, there exist α_0, α such that $\alpha_0 \sqsubseteq_1 \alpha \Uparrow \beta$, $\mathbb{F}(\alpha)$ contains a subsequence of $\overline{\psi}$ and α_0 is an $\mathcal{S}$-residual. It follows from (1) that $\mathbb{F}(\alpha_0)$ contains a subsequence of $\overline{\psi}$. □

Lemma 5.10. Let $\mathcal{S}$ be a finite subset of $\mathfrak{S}$. Suppose:

— $\beta = *_{\overline{a}}(\beta_1, \ldots, \beta_n) \in \mathbb{B}(\mathcal{S})$,
— $\beta' = *_{\overline{b}}(\beta'_1, \ldots, \beta'_n, \beta'_{n+1}, \ldots, \beta'_{n+k}) \in \mathbb{B}(\mathcal{S})$,
— $\beta_i \Subset \beta'_i$ for each $i \in \{1, \ldots, n\}$,
— the sets of $\mathcal{S}$-residuals of β and β' are equal.

Then $\beta \Subset \beta'$.

Proof. Let $\overline{\psi} \in \mathbb{F}(\beta)$. There exists for each $i \in [1, \ldots, n]$ a sequence $\overline{\psi}_i \in \mathbb{F}(\beta_i)$ such that $\overline{\psi} \in \circledast(\{\overline{\psi}_1\}, \ldots, \{\overline{\psi}_n\})$. By assumption there exists for each $i \in [1, \ldots, n]$ an $\alpha_i \Uparrow \beta'_i$ such that $\overline{\psi}_i \in \mathbb{F}(\alpha_i)$. As a consequence $\overline{\psi} \in \mathbb{F}(*(\alpha_1, \ldots, \alpha_n))$.

By Lemma 5.9 there exists an $\mathcal{S}$-residual α_0 of β such that $\mathbb{F}(\alpha_0)$ contains a subsequence $\overline{\phi}$ of $\overline{\psi}$. By assumption α_0 is also an $\mathcal{S}$-residual of β', hence there exist $\alpha'_1, \ldots, \alpha'_{n+k}$, $\overline{b}$ such that $\alpha_0 \sqsubseteq_1 *_{\overline{b}}(\alpha'_1, \ldots, \alpha'_{n+k}) \Uparrow \beta'$. By Lemma 4.11, we have $\overline{\phi} \in \mathbb{F}(*_{\overline{b}}(\alpha'_1, \ldots, \alpha'_{n+k}))$. Hence for each $i \in [1, \ldots, n+k]$, there exists in $\mathbb{F}(\alpha'_i)$ a subsequence of $\overline{\phi}$, which is also a subsequence of $\overline{\psi}$. Now, let $\alpha = *_{\overline{b}}(\alpha_1, \ldots, \alpha_n, \alpha'_{n+1}, \ldots, \alpha'_{n+k})$. Then $\alpha \Uparrow \beta'$,

$\overline{\psi} \in \mathbb{F}(*(\alpha_1, \dots \alpha_n))$, and for each $j \in [1, \dots, k]$ there exists in $\mathbb{F}(\alpha'_{n+j})$ a subsequence of $\overline{\psi}$. As a consequence $\overline{\psi} \in \mathbb{F}(\alpha)$. □

Lemma 5.11. Let $\mathcal{S}$ be a finite subset of $\mathfrak{S}$. Let $\mathcal{B}_\varepsilon$ be a subset of $\mathbb{B}_\varepsilon(\mathcal{S})$. Let $\mathcal{B} = \{*_{\overline{a}}(\beta_1, \dots, \beta_n) | \forall i \in [1, \dots, n], \beta_i \in \mathcal{B}_\varepsilon\}$. If $\Subset$ is an AFR on $\mathcal{B}_\varepsilon$, then $\Subset$ is an AFR on $\mathcal{B}$.

Proof. Let $\mathcal{R} = \mathbb{B}(\mathcal{S}, \Sigma_{i=1}^{1+|\mathcal{S}_@|} i, 1)$ (see Definition 4.4). 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 $\equiv$ (as $\equiv$ is a subset of $\sqsubseteq_1$, see Definition 4.3), that is, $\rho(\beta)$ is a union of the elements of a subset of $\mathcal{R}/_\equiv$. By Lemma 4.5.(1) the latter is a finite set, therefore $\{\rho(\beta) \,|\, \beta \in \mathcal{B}\}$ is a finite set.

For each $\beta = *_{\overline{a}}(\beta_1, \dots, \beta_n) \in \mathcal{B}$ where $\overline{a}$ is increasing w.r.t 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 $\overline{a} = \varepsilon$, $n = 0$ if $\beta = \emptyset_{\mathbb{B}}$, and $\overline{a} = (\varepsilon)$, $n = 1$ if β is a rooted blueprint. Since $\{\rho(\beta) \,|\, \beta \in \mathcal{B}\}$ 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 $\Subset$ is an AFR on $\mathcal{B}_\varepsilon$. By Theorem 5.5, $\Subset_{\mathbb{S}}$ is an AFR on $\{\sigma(\beta) \,|\, \beta \in \mathcal{B}\}$.

Thus for every infinite sequence $(\beta_i)_{i \in \mathbb{N}}$ over $\mathcal{B}$ there exist i, j such that $i < j$, $\sigma(\beta_i) \Subset_{\mathbb{S}} \sigma(\beta_j)$ and β_i, β_j have the same set of residuals. For $\sigma(\beta_i) = (\beta_1^i, \dots, \beta_n^i)$ and $\sigma(\beta_j) = (\beta_1^j, \dots, \beta_{n+k}^j)$, there exists a subsequence $(\beta_{l_1}^i, \dots, \beta_{l_n}^i)$ of $\sigma(\beta_j)$ such that $\beta_1^i \Subset \beta_{l_1}^i, \dots, \beta_n^i \Subset \beta_{l_n}^i$. There exist also $l_{n+1}, \dots, l_{n+k}$ and two sequences $\overline{a}$ and $\overline{b}$ such that $\beta_i = *_{\overline{a}}(\beta_1^i, \dots, \beta_n^i)$ and $\beta_j = *_{\overline{b}}(\beta_{l_1}^j, \dots, \beta_{l_n}^j, \beta_{l_{n+1}}^j, \dots, \beta_{l_{n+k}}^j)$. By Lemma 5.10 we have $\beta_i \Subset \beta_j$. □

5.3. *Axiomatic Kruskal Theorem and main key lemma*

The following definition is borrowed from (Melliès 1998):

Definition 5.12. An *abstract decomposition system* is an 8-tuple

$$(\mathcal{T}, \mathcal{L}, \mathcal{V}, \preceq_{\mathcal{T}}, \preceq_{\mathcal{L}}, \preceq_{\mathcal{V}}, \dot{\longrightarrow}, \vdash)$$

where:

— $\mathcal{T}$ is a set of *terms* noted $t, u, \dots$ equipped with a binary relation $\preceq_{\mathcal{T}}$,
— $\mathcal{L}$ is a set of *labels* noted $f, g, \dots$ equipped with a binary relation $\preceq_{\mathcal{L}}$,
— $\mathcal{V}$ is a set of *vectors* noted $T, U, \dots$ equipped with a binary relation $\preceq_{\mathcal{V}}$,
— $\dot{\longrightarrow}$ is a relation on $\mathcal{T} \times \mathcal{L} \times \mathcal{V}$, *e.g.* $t \xrightarrow{f} T$
— $\vdash$ is a relation on $\mathcal{V} \times \mathcal{T}$, *e.g.* $T \vdash t$.

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

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

An *elementary term* t is a term minimal w.r.t $\rhd_{\mathcal{T}}$, that is, a term for which there exists no u such that $t \rhd_{\mathcal{T}} u$.

Theorem 5.13. (Melliès) Suppose $(\mathcal{T}, \mathcal{L}, \mathcal{V}, \preceq_{\mathcal{T}}, \preceq_{\mathcal{L}}, \preceq_{\mathcal{V}}, \dot{\longrightarrow}, \vdash)$ satisfies the following properties:

— (Axiom I) There is no infinite chain $t_1 \rhd_{\mathcal{T}} t_2 \rhd_{\mathcal{T}} \dots$

— (Axiom II) The relation $\preceq_{\mathcal{T}}$ is an AFR on the set of elementary terms.
— (Axiom III) For all t, u, u',
if $t \preceq_{\mathcal{T}} u'$ and $u \rhd_{\mathcal{T}} u'$, then $t \preceq_{\mathcal{T}} u$.
— (Axiom IV-bis) For all t, u, f, g, T, U,
if $t \xrightarrow{f} T$ and $u \xrightarrow{g} U$ and $f \preceq_{\mathcal{L}} g$ and $T \preceq_{\mathcal{V}} U$, then $t \preceq_{\mathcal{T}} u$.
— (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 $\preceq_{\mathcal{T}}$ is an AFR on $\mathcal{W}_{\vdash}$, then $\preceq_{\mathcal{V}}$ is an AFR on $\mathcal{W}$.

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

Proof. See (Melliès 1998). Mellies' result is actually established for an alternate list of axioms (numbered from I to VI). The possibility to drop Axiom VI and to replace Axiom IV with Axiom IV-bis is a remark that follows the proof of the main theorem. □

Lemma 5.14. For each finite $\mathcal{S} \subseteq \mathfrak{S}$, the relation $\Subset$ is an AFR on $\mathbb{B}(\mathcal{S})$.

Proof. According to Lemma 5.11 it is sufficient to prove that $\Subset$ is an AFR on $\mathbb{B}_{\varepsilon}(\mathcal{S})$. Let $(\mathcal{T}, \mathcal{L}, \mathcal{V}, \preceq_{\mathcal{T}}, \preceq_{\mathcal{L}}, \preceq_{\mathcal{V}}, \dot{\longrightarrow}, \vdash)$ be the abstract decomposition system defined as follows.

— The set $\mathcal{T}$ is $\mathbb{B}_{\varepsilon}(\mathcal{S})$; we let $\alpha \preceq_{\mathcal{T}} \beta$ if and only if there exists an address c such that $\alpha \Subset (\beta_{|c})$ and $\alpha(\varepsilon) = (\beta_{|c})(\varepsilon)$.
— The set $\mathcal{L}$ is the set of all @'s in $\mathcal{S}$, the relation $\preceq_{\mathcal{L}}$ is the identity relation on this set.
— The set $\mathcal{V}$ is $\mathbb{B}(\mathcal{S}) \times \mathbb{B}(\mathcal{S})$.
The relation $\preceq_{\mathcal{V}}$ is defined by $(\alpha_1, \alpha_2) \preceq_{\mathcal{V}} (\beta_1, \beta_2)$ if and only if $\alpha_1 \Subset \beta_1$ and $\alpha_2 \Subset \beta_2$.
— The relation $\dot{\longrightarrow}$ 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 $V = (\alpha_1, \alpha_2)$, $i \in \{1, 2\}$, $\beta_1, \ldots, \beta_n \in \mathbb{B}_{\varepsilon}(\mathcal{S})$ and $\alpha_i = *_{\overline{a}}(\beta_1, \ldots, \beta_n)$, then $V \vdash \beta_j$ for each $j \in [1, \ldots, n]$.

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

(A) For all $\mathcal{T}' \subseteq \mathcal{T}$, the relation $\Subset$ is an AFR on $\mathcal{T}'$ if and only if $\preceq_{\mathcal{T}}$ is an AFR on $\mathcal{T}'$. Indeed, consider an arbitrary infinite sequence $\overline{\alpha}$ over $\mathcal{T}'$. This sequence contains an infinite subsequence $(\alpha)_{i \in \mathbb{N}}$ such that all $\alpha_i(\varepsilon)$ are equal. Clearly $\alpha_i \Subset \alpha_j$ implies $\alpha_i \preceq_{\mathcal{T}} \alpha_j$. Conversely, if $\alpha_i \preceq_{\mathcal{T}} \alpha_j$, then there exists c such that $\alpha_i \Subset \alpha_{j_{|c}}$ and $\alpha_i(\varepsilon) = \alpha_j(\varepsilon) = \alpha_j(c)$. So $\alpha_i \Subset \alpha_{j_{|c}} \Uparrow \alpha_j$, hence $\alpha_i \Subset \alpha_j$.

(B) We now check that all axioms of Theorem 5.13 are satisfied. Axiom I is clear. The set of elementary terms is the set of all blueprints consisting of single formulas of $\mathcal{S}$. The relation $\preceq_{\mathcal{T}}$ is of course an AFR on the set of elementary terms, that is, axiom II is satisfied. Axiom III is immediate. If $(\alpha_1, \alpha_2) \preceq_{\mathcal{V}} (\beta_1, \beta_2)$ then $\alpha_1 \Subset \beta_1$ and $\alpha_2 \Subset \beta_2$, hence $@_{\psi}(\alpha_1, \alpha_2) \Subset @_{\psi}(\beta_1, \beta_2)$, a fortiori $@_{\psi}(\alpha_1, \alpha_2) \preceq_{\mathcal{T}} @_{\psi}(\beta_1, \beta_2)$, hence Axiom IV-bis is satisfied. It remains to prove that Axiom V is satisfied. Let $\mathcal{W} \subseteq \mathcal{V}$. By definition $\mathcal{W}_{\vdash} = \{\beta \in \mathcal{T} \mid \exists (\alpha_1, \alpha_2) \in \mathcal{W}, (\alpha_1, \alpha_2) \vdash \beta\}$. Assuming $\preceq_{\mathcal{T}}$ is an AFR on $\mathcal{W}_{\vdash}$, we prove that $\preceq_{\mathcal{V}}$ is an AFR on $\mathcal{W}$. By (A) the relation $\Subset$ is an AFR on $\mathcal{W}_{\vdash} \subseteq \mathbb{B}_{\varepsilon}(\mathcal{S})$. Let $\mathcal{B} = \{*_{\overline{a}}(\beta_1, \ldots, \beta_n) \mid \forall i \in [1, \ldots, n], \beta_i \in \mathcal{W}_{\vdash}\}$. By Lemma 5.11 the relation $\Subset$ is an AFR

on $\mathcal{B}$. Moreover $\mathcal{W} \subseteq \mathcal{B} \times \mathcal{B}$. By Proposition 5.3.(2) the relation $\preceq_{\mathcal{V}}$ is an AFR on $\mathcal{B} \times \mathcal{B}$, therefore an AFR on $\mathcal{W}$. □

Lemma 5.15. 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 for each address a such that a is a leaf in Ξ, call *step-continuation at* a of Ξ every compact ϕ-shadow Ξ' such that $\mathrm{dom}(\Xi') \subsetneq \mathrm{dom}(\Xi) \cup \{a \cdot (1), a \cdot (2)\}$ and Ξ, Ξ' take the same values on $\mathrm{dom}(\Xi)$. Let $\rightsquigarrow$ be the relation defined by $\Xi \rightsquigarrow \Xi'$ if and only if Ξ' is a step continuation of Ξ. By Lemma 4.5 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, hence it suffices to prove that $\mathcal{C}$ is a finite set. Assume by way of contradiction that $\mathcal{C}$ is infinite. By König's Lemma there exists an infinite sequence $\Xi_0 \rightsquigarrow \Xi_1 \rightsquigarrow \ldots$ over $\mathcal{C}$. The union $\Xi_\infty = \cup_{i \geq 0} \Xi_i$ is a tree of infinite domain. By König's Lemma again, there exists an infinite chain of addresses $a_1 < a_2 < \ldots$ 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, a_j are labelled with $(\overline{\chi}_i, \gamma_i, \psi)$, $(\overline{\chi}_i, \gamma_j, \psi)$, then we cannot have $\gamma_i \Subset \gamma_j$, otherwise there would exist a k such that Ξ_k is not compact. A contradiction follows from Lemma 5.14. □

6. From the shadows to the light

Theorem 6.1. 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 $\mathsf{BB'IW}$,
— the formula ϕ is Λ_{NF}-inhabited (Lemma 1.10),
— there exists a compact Λ_{NF}-inhabitant of ϕ (Lemma 3.9)
— there exists a compact ϕ-shadow with the same tree domain as a Λ_{NF}-inhabitant of ϕ (Lemmas 3.9 and 4.13).

By Lemma 5.15, the set of compact ϕ-shadows is effectively computable from ϕ. By the subformula property (Lemma 1.5), 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. □

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.

References

Anderson, A. R., and Belnap Jr, N. D. (1975) *Entailment: The Logic of Relevance and Necessity, Vol. 1.* Princeton University Press.

Anderson, A. R. (1960) Entailment shorn of modality. *J. Symb. Log.* **25** (4), 388.

Anderson, A. R., Belnap Jr, N. D., and Dunn, J. M. (1990) *Entailment: The Logic of Relevance and Necessity, Vol. 2.* Princeton University Press.

Barendregt, H., *The Lambda Calculus: Its Syntax and Semantics*, Studies in Logic and the Foundations of Mathematics, 103 (Revised ed.), North Holland.

Handbook of Mathematical Logic (1977). Edited by Barwise, J., Studies in Logic and Foundations of Mathematics, North-Holland.

Bezem, M., Klop, J.,W., de Vrijer, R., ("Terese") (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. *Theor. Comput. Sci.* **318** (3), 373–408.

Bunder, M.,W., (1996) Lambda Terms Definable as Combinators. *Theor. Comput. Sci.* **169** (1), 3–21.

Higman, G. (1952) Ordering by divisibility in abstract algebra. *Proc. London Math. Soc.* **3** (2), 326–336.

Kripke, S (1959) The problem of entailment. *J. Symb. Log.* **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. *J. Comb. Theory, Ser. A* **13** (3), 297–305.

Melliès, P.-A. (1998) On a duality between Kruskal and Dershowitz theorems. In: Larsen, K. G, Skyum, S., Winskel, G. (Eds.), ICALP, *Lecture Notes in Computer Science* **1443**, 518–529, Springer-Verlag.

Trigg, P., Hindley, J. R., and Bunder, M. W. (1994) Combinatory abstraction using B, B' and friends. *Theor. Comput. Sci.* **135** (2), 405–422.

Urquhart, A (1984) The undecidability of entailment and relevant implication. *J. Symb. Log.* **49** (4), 1059–1073.