UNIVERSITÀ DEGLI STUDI DI MILANO

Dipartimento di Scienze dell'Informazione

RAPPORTO INTERNO N° 304-05

A Comprehensive Framework for Combined Decision Procedures

Silvio Ghilardi, Enrica Nicolini, Daniele Zucchelli

A Comprehensive Framework for Combined Decision Procedures

Silvio Ghilardi[1], Enrica Nicolini[2], and Daniele Zucchelli[1]
[1]Dipartimento di Scienze dell'Informazione
[2]Dipartimento di Matematica
Università degli Studi di Milano - Italy

October 12, 2005

Abstract

We define a general notion of a fragment within higher order type theory; a procedure for constraint satisfiability in combined fragments is outlined, following Nelson-Oppen schema. The procedure is in general only sound, but it becomes terminating and complete when the shared fragment enjoys suitable noetherianity conditions and allows an abstract version of a 'Keisler-Shelah like' isomorphism theorem. We show that this general decidability transfer result covers as special cases, besides applications which seem to be new, the recent extension of Nelson-Oppen procedure to non-disjoint signatures [28] and the fusion transfer of decidability of consistency of A-Boxes with respect to T-Boxes axioms in local abstract description systems [10]; in addition, it reduces decidability of modal and temporal monodic fragments [58] to their extensional and one-variable components.*

*This Technical Report is the extended version of an invited talk at *FroCoS'05*. The present version includes few additions and improvements resulting from post-conference discussions.

Contents

1 Introduction

Decision procedures for fragments of various logics and theories play a central role in many applications of logic in computer science, for instance in formal methods and in knowledge representation. Within these application domains, relevant data appears to be heterogeneously structured, so that modularity in combining and re-using both algorithms and concrete implementations becomes crucial. This is why the development of meta-level frameworks, accepting as input specialized devices, turns out to be strategic for future advances in building powerful, fully or partially automatized systems. In this paper, we shall consider one of the most popular and simple schemata (due to Nelson-Oppen) for designing a cooperation protocol among separate reasoners; we shall plug it into a higher order framework and show how it can be used to deal with various classes of combination problems, often quite far from the originally intended application domain.

This technical report is the extended version of a paper [30] presented at FroCoS ’05; being a full version, proofs details as well as various kinds of remarks and observations have been included. The conference version presents our settings and our results in a more synthetic way and we recommend it to readers willing to have a quicker look to the material presented here. In this introductory section, we shall also give an intuitive but sufficiently detailed account of our approach, so that readers may get a first global overview.

1.1 Nelson-Oppen Method

Nelson-Oppen method [48], [50], [56] was originally designed in order to combine decision procedures for the universal fragment of first-order theories; this is the kind of problems arising in software verification, although it should be noted that most standard topics in computational algebra [41] concerns equality in finitely presented algebras and hence they can also be equivalently reformulated as decision problems for universal fragments of first-order equational theories.

The basic feature of Nelson-Oppen method is quite simple: constraints involving mixed signatures are purified into equisatisfiable pure constraints and then the specialized reasoners try to share all the information they can acquire concerning constraints in the common subsignature, till an inconsistency is detected or till a saturation state is reached.

There are mainly two big problems that must be adequately addressed in this (broadly intended) Nelson-Oppen approach, namely termination and completeness of the proposed combined procedure. *Termination* can be guaranteed in case the total amount of exchangeable information is finite, i.e. in case there are only finitely many ‘representative’ constraints in the common subsignature; although this is the most frequent method to enforce termination, less brutal requirements might be sufficient (for instance, we shall see in the paper that it is sufficient to assume a weaker 'noetherianity' condition, related to finiteness of properly ascending chains of positive constraints). *Completeness* is however the most serious problem: Nelson-Oppen method was guaranteed to be complete only for disjoint signatures and stably infinite theories, till quite recently, when it was realized [27], [28] that stable infiniteness is just a special case of a compatibility notion, which is related to model completions of shared sub-theories.[1]

The above extension of Nelson-Oppen method to combination of theories operating over non disjoint signatures lead to various applications to decision problems in modal logics: such applications (sometimes involving non trivial extensions of the method as well as integration with other work) concerned transfer of decidability of global consequence relation to fusions [28], [29] and to $\mathcal{E}$-connections [4], [5], as well as transfer of decidability of local consequence relation to fusions [11], [9]. The latter result was rather remarkable and also surprising, not only because the generality of the formulation solved an open problem in modal logic, but also because Nelson-Oppen method is not designed itself to solve word problems [7] and conditions for decidability transfer of word problems were previously formulated in a completely different way for non-disjoint signatures (see [8], [21]).

Thus, most of previously existing decidability results on fusions of modal logics (for instance those in [59]) were recaptured and sometimes also improved by general automated reasoning methods based on Nelson-Oppen ideas. However, this is far from exhausting all the potentialities of such ideas and further extensions are possible. In fact, the standard approach to decision problems in modal/temporal/description logics is directly based on Kripke models (see for instance [10], [24]), without the intermediation of an algebraic formalism, whereas the intermediation of the formalism of Boolean algebras with operators is essential in the approach of papers like [28], [29], [11], [4]. The appeal to the algebraic formulation of decision problems on one side produces proofs which are much smoother and which apply also to semantically incomplete propositional logics, but on the other side it limits the method to the cases in which such a purely algebraic counterpart of semantic decision problems can be identified.

1.2 Towards Higher Order Logic

One of the main reasons for avoiding first-order formalisms in favor of propositional modal logic-style languages lies in the better computational performances of the latter. However, from a purely declarative point of view, first-order formalisms are essential in order to specify

[1] In [26], these results from [27], [28] are included in a general schema for combination of decision procedures for first-order theories in which every decision procedure is formalized as an inference system, and their combination is formalized through the so-called inference modules.

in a semantically meaningful language the relevant decision problems. This goal is mainly achieved in the case of modal logic through first-order translations, the role of such translations being simply that of codifying the intended semantics (and not necessarily that of providing computational tools). The simplest first-order translation known from the literature [57] is the so-called standard translation $ST(\varphi, w)$ of a propositional modal formula φ: this is the mechanism that translates propositional variables x to atomic formulae $X(w)$, boolean connectives to themselves and $\Diamond\varphi$ to $\exists v(R(w, v) \wedge ST(\varphi, v))$ (where R is a binary relation associated with the possibility operator $\Diamond$).

If a semantic class $\mathcal{S}$ of frames (i.e. of sets endowed with a binary relation) is given, relevant decision problems are formulated as satisfiability problems for standard translations of propositional modal formulae. More precisely, the local $\mathcal{S}$-satisfiability problem is the problem of finding a structure in $\mathcal{S}$ in which a first-order formula of the kind $ST(\varphi, w)$ is satisfiable; by contrast, the global $\mathcal{S}$-satisfiability problem is the problem of finding a structure in $\mathcal{S}$ in which a first-order formula of the kind $\forall v ST(\psi, v) \wedge ST(\varphi, w)$ is satisfiable. Clearly, once problems are formulated in this way, the unary predicates occurring in standard translations are considered as *second order variables*[2] (whereas the variables v, w occurring in them are still first order). Let us express the same observation from a slightly different point of view: propositional modal formulae are inquired for (local or global) satisfiability in a Kripke model based on $\mathcal{S}$, where such a Kripke model is a frame $(W, R) \in \mathcal{S}$ together with an assignment for propositional variables. Now it is evident that such a propositional assignment gets converted into an assignment for second order variables, after taking standard translation.

The role played by second order variables becomes even more evident if we analyze the way in which standard translations of modal formulae in fusions are obtained from standard translations of formulae in the component languages. For instance, $ST(\Diamond_1\Diamond_2\mathrm{x}, w)$ is obtained by substituting into

$$ST(\Diamond_1\mathrm{y}, w) = \exists v(R_1(w, v) \wedge Y(v))$$

the 'abstracted' second order term

$$\{v \mid ST(\Diamond_2\mathrm{x}, v)\} = \{v \mid \exists z(R_2(v, z) \wedge X(z))\}$$

for Y (a β-conversion should follow the replacement in order to get as normal form precisely $ST(\Diamond_1\Diamond_2\mathrm{x}, w)$).

Thus, even if we do not 'computationally' trust first-order logic (and consequently not even higher order logic, for much stronger reasons), it makes nevertheless sense to *analyze combination problems in the framework where they arise*, that is in the framework which is

[2]This is also the reason why, if propositional modal formulae are seen as axiom schemata for a logic, correspondence theory [57] associates with them a second order sentence.

the most natural for them. This analysis will tend, as a first effect, to give some higher order version of procedures which were already shown to be successful in the plain first-order case: Nelson-Oppen is among such procedures.

1.3 Constraint Solving in Fragments

We choose Church's type theory as our framework for higher order logic: thus our syntax deals with types and terms, terms being endowed with a (codomain) type. Types can be built up from primitive sorts by using function type constructor, whereas terms can be built up from typed variables and constants, λ-abstraction and function evaluation. Types include the truth-value type Ω and constants include symbols for boolean connectives and for equality over each type. Formulae are treated as special terms, namely terms of type Ω and quantifiers can be introduced through explicit definitions (see Section 2 for more details). If φ has type Ω and x is a variable (of type, say, τ), we write $\{x \in \tau \mid \varphi\}$ (or simply $\{x \mid \varphi\}$) for the term $(\lambda x)\varphi$ of type $\tau \to \Omega$.

The first task to be accomplished in this framework is that of finding a definition of what we mean by a fragment (only relatively small fragments can indeed have a chance to be decidable here). Clearly, a fragment should be a pair consisting of a signature for type theory and on a set of terms in that signature; however, we want the set of terms to enjoy some minimal properties that make it suitable for our combination purposes. These properties are fixed in the notion of an *algebraic fragment* (see Definition 3.2 below): basically in an algebraic fragment $\Phi = \langle \mathcal{L}, T \rangle$, the set of terms T in the signature $\mathcal{L}$ must be closed under composition (i.e. under substitution) and must contain all variables whose type is a Φ-type (a Φ-type is the type of a term in T or of a variable occurring free in some $t \in T$).

Clearly simply typed λ-calculus, first-order terms or formulae, standard translations of modal formulae, etc. are algebraic fragments in the above sense.[3] Notice, however, that a honest 'naively intended' fragment can be formally turned into an algebraic fragment in many ways and this flexibility is crucial when taking *combined* algebraic fragments (these are taken to be the minimum algebraic fragments extending two given ones, see the definition at the beginning of Section 4). For instance, in the modal case, if we take the fragment consisting of terms $\beta\eta$-equivalent to terms of the kind $\{w \mid ST(\varphi, w)\}$, then the only type in this fragment is $W \to \Omega$ and consequently this fragment produces (up to β-conversions) the λ-abstraction of standard translations of modal formulae in fusion languages, if it is combined with another fragment of the same kind. By contrast, if we take the fragment consisting of

[3] Sometimes, one needs to include certain variables in the set of terms of the intended fragment to match precisely the above definition, but the addition of such variables is usually harmless for the kind of problems we consider in this paper.

terms of the kind $ST(\varphi, w)$, then the combination with another fragment of the same kind does not produce anything really interesting (just variables can be used to try to form combined terms by replacement); however now there are different types in the fragment (namely W and Ω), which means that the fragment is ready for non completely trivial combinations with fragments containing for instance non-variable first-order terms of type W. Of course, we can also consider a third version of a modal fragment, consisting both of terms of the kind $\{w \mid ST(\varphi, w)\}$ and of terms of the kind $ST(\varphi, w)$: this leaves both combination possibilities open. From the point of view of decision problems, notice that whenever the λ-abstracted terms $\{w \mid ST(\varphi, w)\}$ are in the fragment, constraint satisfiability problems as defined below cover satisfiability of a concept description wrt a given T-Box (in the terminology of description logics); if we want they to cover satisfiability of an A-Box, we need in the fragment, besides terms of the kind $ST(\varphi, w)$, also the further terms $R(v, w)$, where R is the binary constant corresponding to a diamond operator and v, w are variables of type W. Hence, for satisfiability of an A-Box wrt a T-Box, all terms of the kind $\{w \mid ST(\varphi, w)\}, ST(\varphi, w)$ and $R(v, w)$ must be included in the fragment. The moral of this discussion is that an algebraic fragment must be appropriately designed, taking into consideration both actual expressivity power and potential combination opportunities.

An algebraic fragment $\Phi = \langle \mathcal{L}, T \rangle$ is *interpreted* when a class $\mathcal{S}$ of (ordinary set-theoretic) models of $\mathcal{L}$ is fixed ($\mathcal{S}$ is assumed to be closed under isomorphisms). All the algebraic fragments we consider are interpreted and if $\mathcal{S}$ is not specified, it is intended to be the class of all $\mathcal{L}$-structures.

Given an algebraic fragment $\Phi = \langle \mathcal{L}, T \rangle$, a Φ-atom is an equation like $t_1 = t_2$ for terms $t_1, t_2 \in T$ (having the same type), a Φ-literal is a Φ-atom or the negation of a Φ-atom, a Φ-clause is a disjunction of Φ-literals and a Φ-constraint is a conjunction of Φ-literals. Now the *word problem* for an interpreted algebraic fragment $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ is the problem of deciding unsatisfiability of a negative Φ-literal in all $\mathcal{A} \in \mathcal{S}$, whereas the *constraint satisfiability problem* for Φ is the problem of deciding the satisfiability of a Φ-constraint in some $\mathcal{A} \in \mathcal{S}$.

Notice that there are genuinely higher order interpreted algebraic fragments whose word problem is decidable (see for instance Friedman theorem for simply typed λ-calculus [22]) and also whose constraint satisfiability problem is decidable (see Rabin results on monadic second order logic [51]). Literature on modal/temporal/description logics is full of examples of interpreted algebraic fragments having decidable word or constraint satisfiability problems (the former usually corresponds to local satisfiability, whereas the latter to global satisfiability, but the precise relationship depends on the specific version of the fragment that is adopted, see the remarks above). The whole literature on computational algebra deals with examples of interpreted algebraic fragments consisting of first-order individual terms and having decidable word or conditional word problems (the latter correspond to our constraint satisfiability problems, given the convexity of equational theories).

1.4 Why Nelson-Oppen may happen to succeed

Our main goal is the combination of decision procedures for constraint satisfiability. The definition of an algebraic fragment we gave is sufficient to substantially reproduce Nelson-Oppen purification steps in our framework, hence we may freely suppose that combined constraints can be splitted into sets of equisatisfiable pure constraints. The interpreted algebraic fragments Φ_1, Φ_2 to be combined share some interpreted algebraic subfragment Φ_0 in the common subsignature and a Nelson-Oppen style fair exchange protocol can now be activated. In Subsection 4.2 we describe it in a detailed way, assuming only the availability of two abstractly axiomatized sound 'residue enumerators' for the positive Φ_0-clauses which are entailed by the current Φ_1- (resp. Φ_2)-constraints. The procedure terminates in case certain noetherianity conditions are satisfied (this is always the case whenever Φ_0 is locally finite, namely whenever there are only finitely many non Φ_0-equivalent terms in a fixed finite number of variables).

The main problem we are faced now is the completeness of the procedure: can we infer that the input combined constraint is satisfiable in case the system halts with a saturation message? In general this *cannot* be the case: in fact, it can be shown for instance that the (λ-abstractions of) first-order translations of modal formulae in a product modal logic like $K \times K$ form a combined interpreted algebraic fragment, whose components have decidable constraint satisfiability problems and share also a locally finite common subfragment. However, it is well known that $K \times K$ has undecidable global consequence relation [24] and in fact our combined procedure fails in proving also basic commutativity principles. The point is that a 'saturation-found' run of our procedure means at the semantic level (see Proposition 4.11 below) the existence of Φ_i-structures ($i = 1, 2$), whose Φ_0-reducts are only $\Phi_0(\underline{c}_0)$-equivalent (i.e. such Φ_0-reducts satisfy the same closed Φ_0-atoms, in a language augmented with finitely many free constants $\underline{c}_0$): this is far from being enough to build a combined structure satisfying the union of the pure constraints.

To analyze the reason why, under certain hypotheses, Nelson-Oppen approach nevertheless succeeds, let us turn to the original Nelson-Oppen case. Here, signatures are one-sorted, first order and disjoint; if we are given *infinite* structures $\mathcal{A}_i$ for the input languages[4] whose Φ_0-reducts satisfy the same equations among the free constants in $\underline{c}_0$, then it is indeed possible to build a structure for the combined language out of them. The reasons are the following: (i) the theory of an infinite set is the model completion of the theory of equality, hence from the fact

[4] We may assume that $\mathcal{A}_1, \mathcal{A}_2$ are infinite because the hypothesis for the standard Nelson-Oppen procedure is precisely that the input theories are stably infinite.

that $\mathcal{A}_1$ and $\mathcal{A}_2$ are $\Phi_0(\underline{c}_0)$-equivalent if follows that they are also elementarily equivalent, as far as the language $\mathcal{L}_0(\underline{c}_0)$ of $\Phi_0(\underline{c}_0)$ is concerned; (ii) by Keisler-Shelah theorem [17], applying ultrapowers to $\mathcal{A}_1, \mathcal{A}_2$, we can make them $\mathcal{L}_0(\underline{c}_0)$-isomorphic and this is clearly sufficient to build the combined structure.[5]

Roughly speaking (see Section 5 for details), let us call *isomorphism theorem* a theorem saying that the application of a certain semantic operation makes two $\Phi_0(\underline{c}_0)$-equivalent structures $\mathcal{L}_0(\underline{c}_0)$-isomorphic. The above outlined argument says that the existence of a suitable isomorphism theorems is sufficient for our combined procedure to be complete: this is formally stated in our main result, namely Theorem 5.6.

Of course, isomorphisms theorems are quite peculiar and rare. However, another isomorphism theorem justifies the completeness of the procedure in the case of fusions of modal logics.[6] In that case, the fragment $\Phi_0(\underline{c}_0)$ is included into the one-variable fragment of first-order classical logic. Here a structure is specified up to $\Phi_0(\underline{c}_0)$-equivalence once we are given the information about emptiness of subsets definable through Boolean combinations of the finitely many unary predicates $\underline{c}_0$. If $\mathcal{A}_1, \mathcal{A}_2$ are $\Phi_0(\underline{c}_0)$-equivalent, then they can be made $\mathcal{L}_0(\underline{c}_0)$-isomorphic by taking disjoint unions, provided the set indexing the cardinality of the disjoint copies is sufficiently large. Now if the input modal fragments are interpreted in a semantic class closed under disjoint unions, Theorem 5.6 applies (notice that we gave in this way also an explanation of the reasons why closure under disjoint unions is the crucial hypothesis in fusion transfer decidability results).

Since in the above argument involving disjoint unions, there is nothing very specific to fragments obtained through standard translations of propositional modal formulae, it is evident that we can get analogous transfer results starting from interpreted guarded and packed guarded fragments [1], [33], [46], which are also preserved under taking disjoint copies of the same structure. Of course, guarded and packed guarded fragments should be designed in such a way that all predicate symbols, except unary ones, are taken to be constants: in other words, second order variables should be used just for unary predicates and not for relations (in this way the shared fragment is still contained into the monadic fragment of first-order classical logic). But further possible combinations arise, for instance because first order formulae without equality are also preserved under taking disjoint copies, provided taking disjoint copies at the level of structures is defined in a proper alternative way (that we will call 'conglomera-

[5]Instead of this argument, a more elementary argument, based on Robinson's joint consistency theorem, is used in [28].

[6]This argument applies more generally to relativized satisfiability in abstract description systems in the sense of [10]. Notice that, since full A-Boxes are involved, this application of our procedure goes slightly beyond the cases (analyzed for instance in [28], [29], [11], [9], [4]) where a reformulation of satisfiability problems in the algebraic framework of Boolean algebras with operators is needed.

tion' in Subsection 5.3). Theorem 5.21 summarize these new decidability transfer results in a unique statement, referring to the notion of a *monadically suitable* fragment. We also show in Theorem 5.11 how to get decidability transfer results for the combination of an A-Box (in the sense of description logics) and of a stably infinite first order theory, operating on disjoint signatures.

As a final application, we shall see how to analyze monodic modal fragments (in the sense of [58], [24]) as combinations of extensional first-order fragments and standard translations of one variable modal fragments. Since a suitable isomorphism theorem (based on disjoint copies and fiberwise disjoint copies) holds here too, our procedure is complete and justifies new rather general decidability transfer results. Notice that in order to apply Theorem 5.6 to this case, we make some little use (very simple, but important from a conceptual point of view) of ideas coming from *descent theory*, see Subsection 5.4 for details.

In conclusion, our higher order logic approach to combination problems seems to be fruitful, in the sense that it encompasses relevant known results and suggests new applications. From a completely different perspective, we hope it might also contribute to the integration of fully automatized specialized reasoners into higher order proof assistants.

2 Type-Theoretic Languages

We fix our notation for higher order syntax; we adopt a type theory in Church's style (see [2], [3], [39] for introductions to the subject).

2.1 Signatures

We use letters $S_1, S_2, \dots$ to indicate *sorts* (also called *primitive types*) of a signature. Formally, sorts are a set $\mathcal{S}$ and *types over* $\mathcal{S}$ are built inductively as follows:

- every sort $S \in \mathcal{S}$ is also a type;
- Ω is a type (this is called the *truth-values* type):
- if τ_1, τ_2 are types, so is $(\tau_1 \to \tau_2)$.

As usual external brackets are omitted; moreover, we shorten the expression $\tau_1 \to (\tau_2 \to \dots (\tau_n \to \tau))$ into $\tau_1 \dots \tau_n \to \tau$ (in this way, every type τ has the form $\tau_1 \dots \tau_n \to \tau$, where $n \geq 0$ and τ is a sort or it is Ω). In the following, we use the notation $\mathcal{T}(\mathcal{S})$ or simply $\mathcal{T}$ to indicate a *types set*, i.e. the totality of types that can be built up from the set of sorts $\mathcal{S}$. In this way, $\mathcal{S}$ is sometimes left implicit in the notation, however we always reserve to sorts the letters $S_1, S_2, \dots$ (as opposed to the letters τ, υ, etc. which are used for arbitrary types).

A *signature* (or a *language*) is a triple $\mathcal{L} = \langle \mathcal{T}, \Sigma, a\rangle$, where $\mathcal{T}$ is a types set, Σ is a set of constants and a is an *arity map*, namely a map $a : \Sigma \longrightarrow \mathcal{T}$; we write $f : \tau_1 \dots \tau_n \to \tau$ to express that f is a *constant of type* $\tau_1 \dots \tau_n \to \tau$, i.e. that $a(f) = \tau_1 \dots \tau_n \to \tau$. According to the above observation, we can assume that τ is a sort or that $\tau = \Omega$; in the latter case, we say that f is a *predicate* or a *relational* symbol (predicate symbols are preferably indicated with the letters $P, Q, \dots$).

We require the following *special constants* to be always present in a signature:

- $\top$ and $\bot$ of type Ω;
- $\neg$ of type $\Omega \to \Omega$;
- $\vee$ and $\wedge$ of type $\Omega\,\Omega \to \Omega$;
- $=_\tau$ of type $\tau\,\tau \to \Omega$ for each type $\tau \in \mathcal{T}$ (we usually write it as '=' without specifying the subscript τ).

The *proper symbols* of a signature are its sorts and its non special constants.

A signature is *one-sorted* iff its set of sorts is a singleton. A signature $\mathcal{L}$ is *first-order* if for any proper $f \in \Sigma$, we have that $a(f) = S_1 \dots S_n \to \tau$, where τ is a sort or it is Ω. A first-order signature is called *relational* iff any proper $f \in \Sigma$ is a relational constant, that is $a(f) = S_1 \dots S_n \to \Omega$. By contrast, a first order signature is called *functional* iff any proper $f \in \Sigma$ has arity $S_1 \dots S_n \to S$.

Let $\mathcal{L}_1 = \langle \mathcal{T}_1, \Sigma_1, a_1\rangle$ and $\mathcal{L}_2 = \langle \mathcal{T}_2, \Sigma_2, a_2\rangle$ be two signatures; we say that $\mathcal{L}_1$ is a sub-signature of $\mathcal{L}_2$ (written $\mathcal{L}_1 \subseteq \mathcal{L}_2$) if $\mathcal{T}_1 \subseteq \mathcal{T}_2$, $\Sigma_1 \subseteq \Sigma_2$ and $a_1 \subseteq a_2$. Furthermore, given $\mathcal{L}_1 = \langle \mathcal{T}_1, \Sigma_1, a_1\rangle$ and $\mathcal{L}_2 = \langle \mathcal{T}_2, \Sigma_2, a_2\rangle$, in case a_1 and a_2 coincide[7] on $\Sigma_1 \cap \Sigma_2$, we define the *union signature* $\mathcal{L}_1 \cup \mathcal{L}_2$ to be (let $\mathcal{T}_1$ be $\mathcal{T}(\mathcal{S}_1)$ and $\mathcal{T}_2$ be $\mathcal{T}(\mathcal{S}_2)$) $\langle \mathcal{T}(\mathcal{S}_1 \cup \mathcal{S}_2), \Sigma_1 \cup \Sigma_2, a_1 \cup a_2\rangle$ and the *intersection signature* $\mathcal{L}_1 \cap \mathcal{L}_2$ to be $\langle \mathcal{T}_1 \cap \mathcal{T}_2, \Sigma_1 \cap \Sigma_2, a_1 \cap a_2\rangle$.

2.2 Terms

Given a signature $\mathcal{L} = \langle \mathcal{T}, \Sigma, a\rangle$ and a type $\tau \in \mathcal{T}$, we define the notion of an $\mathcal{L}$-*term* (or just *term*) of type τ, written $t : \tau$, as follows (for the definition we need, for every type $\tau \in \mathcal{T}$, a countable supply V_τ of variables of type τ):

- $x : \tau$ (for $x \in V_\tau$) is an $\mathcal{L}$-term of type τ;
- $c : \tau$ (for $c \in \Sigma$ and $a(c) = \tau$) is an $\mathcal{L}$-term of type τ;

[7]Modulo renaming some elements of Σ_1, we can assume that this condition is always satisfied, so that union and intersection signatures are always defined.

– if $t : \upsilon \to \tau$ and $u : \upsilon$ are $\mathcal{L}$-terms of types $\upsilon \to \tau$ and υ, respectively, then $val_\upsilon(t, u) : \tau$ (also written as $t(u) : \tau$) is an $\mathcal{L}$-term of type τ ;

– if $t : \tau$ is an $\mathcal{L}$-term of type τ and $x \in V_\upsilon$ is a variable of type υ, $\lambda x^\upsilon\, t : \upsilon \to \tau$ is an $\mathcal{L}$-term of type $\upsilon \to \tau$.

In the following, we consider the notation x^τ (c^τ) equivalent to $x : \tau$ ($c : \tau$), where x (resp. c) is a variable (resp. a constant); if it can be deduced from the context, *the specification of the type of a term may be omitted*. Moreover, a term of type τ is also called a τ*-term* and terms of type Ω are also called *formulae*. Given a formula φ, we write $\{x \mid \varphi\}$ for $\lambda x\, \varphi$.

We shorten $val_{\upsilon_n}(\cdots(val_{\upsilon_1}(t, u_1), \cdots), u_n)$ to $t(u_1, \ldots, u_n)$ where u_i is a term of type υ_i ($i \in \{1, \ldots, n\}$) and t is a term of type $\upsilon_1 \ldots \upsilon_n \to \tau$.

For each term φ of type Ω, we define the Ω-terms $\forall x^\upsilon\, \varphi$ and $\exists x^\upsilon\, \varphi$ as $\{x^\upsilon \mid \varphi\} = \{x^\upsilon \mid \top\}$ and as $\neg\, \forall x^\upsilon\, \neg\varphi$, respectively (the latter can also be defined differently, in an intuitionistically acceptable way, see [39]). For terms φ_1, φ_2 of type Ω, the terms $\varphi_1 \to \varphi_2$ and $\varphi_1 \leftrightarrow \varphi_2$ of type Ω are classically defined by $\neg\varphi_1 \vee \varphi_2$ and by $(\varphi_1 \to \varphi_2) \wedge (\varphi_2 \to \varphi_1)$, respectively (but notice that $\varphi_1 \leftrightarrow \varphi_2$ can be defined in a semantically equivalent way also as $\varphi_1 = \varphi_2$).

By the above definitions, first-order formulae *can be considered as a subset* of the higher order formulae defined in this section. More specifically, when we speak of *first-order terms*, we mean variables $x : S$, constants $c : S$ and terms of the kind $f(t_1, \ldots, t_n) : S$, where $t_1, \ldots, t_n$ are (inductively given) first-order terms and $a(f) = S_1 \cdots S_n \to S$. Now *first-order formulae* are obtained from formulae of the kind $\top : \Omega, \bot : \Omega, P(t_1, \ldots, t_n) : \Omega$ (where $t_1, \ldots, t_n$ are first-order terms and $a(P) = S_1 \cdots S_n \to \Omega$) by applying $\exists x^S, \forall x^S, \wedge, \vee, \neg, \to, \leftrightarrow$.

2.3 Substitutions and Conversions

An occurrence of a variable x in a term t is *bound* if it appears in a subterm of t of the kind $\lambda x\, u$, otherwise it is said to be *free*. A variable x occurs free in a term t if and only if at least one occurrence of x in t is free; by $fvar(t)$ we mean the set of the variables that occur free in t (whereas $fvar_\tau(t)$ is the set of variables of type τ that occur free in t). If Γ is a set of terms, $fvar(\Gamma)$ means $\bigcup_{t \in \Gamma} fvar(t)$. We often use notations like $\underline{x}, \underline{y}$ to mean tuples of distinct free variables.

A term without free variables is called a *closed* term and a formula without free variables is called a *sentence*. The notation $t[x_1, \ldots, x_n]$ (resp. $\Gamma[x_1, \ldots, x_n]$) means that $fvar(t) \subseteq \{x_1, \ldots, x_n\}$ (resp. $fvar(\Gamma) \subseteq \{x_1, \ldots, x_n\}$).

Two terms are said to be equivalent modulo α*-conversion* iff they differ only by a bound variables renaming; in the following, we shall identify α-equivalent terms, i.e. we consider terms as representatives of their equivalence class modulo α-conversion.

Let $\mathcal{V}$ be the disjoint union of the sets of variables V_τ ($\tau \in \mathcal{T}$). We define the notion of substitution as usual: a *substitution* is a map $\sigma : \mathcal{V} \to T$ (from the set $\mathcal{V}$ of the variables into the set T of the terms) that respects types (i.e. if $x \in V_\tau$ then $x\sigma$ is a term of type τ) and such that the set $\{x \mid x \not\equiv x\sigma\}$ is finite.[8] The set $dom(\sigma) := \{x \mid \not\equiv x\sigma\}$ is called the *domain* of the substitution σ. A substitution σ will be written as $x_1 \mapsto x\sigma_1, \ldots, x_n \mapsto x_n\sigma$, or equivalently as $x_1\sigma/x_1, \ldots, x_n\sigma/x_n$, where $dom(\sigma) \subseteq \{x_1, \ldots, x_n\}$. A substitution is a *renaming* iff it is a variable permutation.

Substitutions can be extended in the domain from variables to all terms in the usual way; notice however that, when defining inductively the term $t\sigma$, it might happen that α-conversions must be applied before actual replacements, in order to avoid clashes. If $\sigma = \{x_1 \mapsto u_1, \ldots, x_n \mapsto u_n\}$ and $fvar(t) \subseteq \{x_1, \ldots, x_n\}$, the term $t\sigma$ can also be written as $t[u_1, \ldots, u_n]$. Given two substitutions σ_1 and σ_2, the *composite* substitution $\sigma_1\sigma_2$ is the substitution that maps the variable x to $(x\sigma_1)\sigma_2$.

The notion of $\beta\eta$-equivalence between terms is introduced through the following previous inductive definition of the relation $\rhd^1_{\beta\eta}$ (we follow [19]):

- (β) $\quad val(\lambda x\, t, u) \rhd^1_{\beta\eta} t\sigma$, where $\sigma : \{x \mapsto u\}$;
- (η) $\quad \lambda x\, val(t, x) \rhd^1_{\beta\eta} t$, if x is not free in t;
- (μ) $\quad val(t, u) \rhd^1_{\beta\eta} val(t, u')$, if $u \rhd^1_{\beta\eta} u'$;
- (ν) $\quad val(t, u) \rhd^1_{\beta\eta} val(t', u)$, if $t \rhd^1_{\beta\eta} t'$;
- (ξ) $\quad \lambda x\, t \rhd^1_{\beta\eta} \lambda x\, t'$, if $t \rhd^1_{\beta\eta} t'$.

The *$\beta\eta$-equivalence relation* $\sim_{\beta\eta}$ is now the reflexive, symmetric and transitive closure of the relation $\rhd^1_{\beta\eta}$. By definition, $\sim_{\beta\eta}$ is an equivalence relation compatible with the term constructors. It is known that the $\beta\eta$-reduction relation $\rhd_{\beta\eta}$ obtained from the transitive closure of $\rhd^1_{\beta\eta}$, gives a rewrite system that is strongly normalizable and confluent [13], [31], i.e. each term has *a unique $\beta\eta$-normal form* modulo α-conversion.

Sometimes, however, it is preferable to use the so-called *long-$\beta\eta$-normal form* of a term $t : \tau$ (instead of the $\beta\eta$-normal form of t). This is defined as follows: suppose $\tau = \tau_1 \cdots \tau_n \to \upsilon$, consider the $\beta\eta$-normal form

$$\lambda x_1 \cdots \lambda x_m\; y(u_1, \ldots, u_p)$$

of t (here $m \leq n$ and y is a variable or a constant) and then take

$$\lambda x_1 \cdots \lambda x_m \lambda x_{m+1} \cdots \lambda x_n\; y(u'_1, \ldots, u'_p, x'_{m+1}, \ldots, x'_n)$$

[8] Since the equality symbol '=' is present in the object language, we prefer to use '$\equiv$' in the metalanguage for coincidence of syntactic expressions.

to be the long-$\beta\eta$-normal forms of t (where $u'_1, \ldots, u'_p, x'_{m+1}, \ldots, x'_n$ are the long-$\beta\eta$-normal forms of $u_1, \ldots, u_p, x_{m+1}, \ldots, x_n$, respectively). Thus, for instance, the long $\beta\eta$-normal form of a predicate constant $P : \tau \to \Omega$ is $\{x \mid P(x)\}$.

2.4 Models

In order to introduce our computational problems, we need to recall the notion of an interpretation of a type-theoretic language. Formulae of higher order type theory which are valid in ordinary set-theoretic models do not form an axiomatizable class, as it is well-known from classical limitative results. Hence, in order to re-gain axiomatizability, one has to use Henkin models [2] or to take interpretations into elementary toposes [39]. However, we shall confine ourselves to standard set-theoretic models, because we are not interested in the whole type theoretic language (nor in any calculus for it). On the other hand, the generalization to more powerful semantics of the definitions given in this subsections is well-known and can be found for instance in the above mentioned textbooks.

If we are given a map that assigns to every sort $S \in \mathcal{S}$ a set $[\![S]\!]$, we can inductively extend it to all types over $\mathcal{S}$, by taking $[\![\tau \to \upsilon]\!]$ to be the set of functions from $[\![\tau]\!]$ to $[\![\upsilon]\!]$.

Given a language $\mathcal{L} = \langle \mathcal{T}, \Sigma, a \rangle$, a *$\mathcal{L}$-structure* (or just a *structure*) $\mathcal{A}$ is a pair $\langle [\![-]\!]_{\mathcal{A}}, \mathcal{I}_{\mathcal{A}} \rangle$, where:

(i) $[\![-]\!]_{\mathcal{A}}$ is a function assigning to a sort $S \in \mathcal{T}$, a set $[\![S]\!]_{\mathcal{A}}$;

(ii) $\mathcal{I}_{\mathcal{A}}$ is a function assigning to a constant $c \in \Sigma$ of type τ, an element $\mathcal{I}_{\mathcal{A}}(c^\tau) \in [\![\tau]\!]_{\mathcal{A}}$ (here $[\![-]\!]_{\mathcal{A}}$ has been extended from sorts to types as explained above).

In every structure $\mathcal{A}$, we finally require also that $[\![\Omega]\!]_{\mathcal{A}} = \{0, 1\}$, that $\mathcal{I}_{\mathcal{A}}(\bot) = 0$, that $\mathcal{I}_{\mathcal{A}}(\top) = 1$, that $\mathcal{I}_{\mathcal{A}}(=_\tau)$ is the characteristic function of the identity relation on $[\![\tau]\!]_{\mathcal{A}}$, and that $\mathcal{I}_{\mathcal{A}}(\neg)$, $\mathcal{I}_{\mathcal{A}}(\vee)$, $\mathcal{I}_{\mathcal{A}}(\wedge)$ are the usual truth tables functions (notice that, in these and similar passages, we implicitly use the isomorphisms $(X^Y)^Z \simeq X^{Y \times Z}$ in order to treat in the natural way curryfied binary function symbols).

We do not exclude, in principle, that in a structure $\mathcal{A}$ we can have $[\![\tau]\!]_{\mathcal{A}} = \emptyset$ for some type τ (in fact, the use of *finitary* assignments below is compatible with empty domains);[9] however, when dealing with one-sorted signatures $\mathcal{L}$, we shall implicitly assume, for simplicity, that in $\mathcal{L}$-structures *the unique sort is always interpreted into a non empty domain.*

[9]Usual (total) assignments are inadequate if one wants tautological sentences to be satisfiable in a structure in which some sort is interpreted into the empty set. Finite assignments eliminate this inconvenient, however empty domains cause further problems on the syntactic side, if one wants to formulate suitable calculi. These questions do not concern the present paper, however we recall that there is a simple well-known solution to them, namely the explicit indication of the variables involved in a proof (see [39]).

Given a $\mathcal{L}$-structure $\mathcal{A} = \langle [\![-]\!]_{\mathcal{A}}, \mathcal{I}_{\mathcal{A}} \rangle$, let $\mathcal{L}^{\mathcal{A}}$ be the language enriched by a constant $\bar{a}$ of type τ for every $a \in [\![\tau]\!]_{\mathcal{A}}$; $\mathcal{A}$ can be canonically considered as a $\mathcal{L}^{\mathcal{A}}$-structure once $\mathcal{I}_{\mathcal{A}}$ is extended to the new constants by stipulating that $\mathcal{I}_{\mathcal{A}}(\bar{a}) := a$. By induction, it is now possible to extend $\mathcal{I}_{\mathcal{A}}$ to all closed $\mathcal{L}^{\mathcal{A}}$-terms t as follows:

- $\mathcal{I}_{\mathcal{A}}(val(t, u)) = \mathcal{I}_{\mathcal{A}}(t)(\mathcal{I}_{\mathcal{A}}(u))$ (this is set-theoretic functional application);
- $\mathcal{I}_{\mathcal{A}}(\lambda x^{\tau}\, t)$ is the function that maps each element $a \in [\![\tau]\!]_{\mathcal{A}}$ into $\mathcal{I}_{\mathcal{A}}(t[\bar{a}/x])$.

From now on, we shall not distinguish for simplicity between a and its name $\bar{a}$.

A $\mathcal{L}^{\mathcal{A}}$-sentence φ is *true* in $\mathcal{A}$ (in symbols $\mathcal{A} \models \varphi$) iff $\mathcal{I}_{\mathcal{A}}(\varphi) = 1$. Notice that, according to the above definition of universal quantification, we have that $\mathcal{A} \models \forall x^{\tau}\, \varphi$ if and only if for each $a \in [\![\tau]\!]_{\mathcal{A}}$, we have $\mathcal{I}_{\mathcal{A}}(\varphi[a/x]) = 1$.

To introduce the notion of satisfiability we use finite assignments. Let $\mathcal{A} = \langle [\![-]\!]_{\mathcal{A}}, \mathcal{I}_{\mathcal{A}} \rangle$ be a $\mathcal{L}$-structure and let $\underline{x}$ be a finite set of variables; an $\underline{x}$*-assignment* (or simply an assignment if $\underline{x}$ is clear from the context) α is a map associating with every variable $x^{\tau} \in \underline{x}$ an element $\alpha(x) \in [\![\tau]\!]_{\mathcal{A}}$. An $\mathcal{L}$-formula φ is *satisfied* in $\mathcal{A}$ under the $\underline{x}$-assignment α (where $\underline{x} \supseteq fvar(\varphi)$) iff $\mathcal{I}^{\alpha}_{\mathcal{A}}(\varphi) = 1$, where $\mathcal{I}^{\alpha}_{\mathcal{A}}(\varphi)$ is the $\mathcal{L}^{\mathcal{A}}$-sentence obtained by replacing in φ the variables $x \in \underline{x}$ by (the names of) $\alpha(x)$. We usually write $\mathcal{A} \models_{\alpha} \varphi$ for $\mathcal{I}^{\alpha}_{\mathcal{A}}(\varphi) = 1$.

A formula is *satisfiable* iff it is satisfied under some assignment and a set of formulae Γ (containing altogether only finitely many variables) is satisfiable iff for some assignment α we have that $\mathcal{A} \models_{\alpha} \varphi$ holds for each $\varphi \in \Gamma$ (of course, for this to make sense, α must be an $\underline{x}$-assignment for some $\underline{x} \supseteq fvar(\Gamma)$ - and one can even assume $\underline{x} = fvar(\Gamma)$ without loss of generality).

For signature inclusions $\mathcal{L}_0 \subseteq \mathcal{L}$, there is an obvious *taking reduct* operation mapping an $\mathcal{L}$-structure $\mathcal{A}$ to an $\mathcal{L}_0$-structure $\mathcal{A}_{|\mathcal{L}_0}$; we can similarly take the $\mathcal{L}_0$-reduct of an assignment, by ignoring the values assigned to variables whose types are not in $\mathcal{L}_0$ (we leave the reader to define these notions properly).

Two $\mathcal{L}$-structures $\mathcal{A}_1 = \langle [\![-]\!]_{\mathcal{A}_1}, \mathcal{I}_{\mathcal{A}_1} \rangle$ and $\mathcal{A}_2 = \langle [\![-]\!]_{\mathcal{A}_2}, \mathcal{I}_{\mathcal{A}_2} \rangle$ are said to be *isomorphic* iff there are bijections $\iota_{\tau} : [\![\tau]\!]_{\mathcal{A}_1} \longrightarrow [\![\tau]\!]_{\mathcal{A}_2}$ (varying $\tau \in \mathcal{T}$) such that $\iota_{\tau}(\mathcal{I}_{\mathcal{A}_1}(c)) = \mathcal{I}_{\mathcal{A}_2}(c)$ holds for all $c : \tau \in \Sigma$ and such that $\iota_{\tau \to \upsilon}(h) = \iota_{\upsilon} \circ h \circ \iota_{\tau}^{-1}$ holds for all $h \in [\![\tau \to \upsilon]\!]_{\mathcal{A}_1}$.[10] Isomorphic structures are in fact indistinguishable (in particular, the same sentences are true in them).

[10] Thus, once again, to give an isomorphism it is sufficient to specify the bijections ι_S for all sorts S.

3 Fragments

General type theory is very hard to attack from a computational point of view, this is why we are basically interested only in more tractable fragments and in combinations of them. Fragments are defined as follows:

Definition 3.1. *A* fragment *is a pair* $\langle \mathcal{L}, T \rangle$ *where* $\mathcal{L} = \langle \mathcal{T}, \Sigma, a \rangle$ *is a signature and* T *is a recursive set of* $\mathcal{L}$*-terms.*

3.1 Algebraic Fragments

We want to use fragments as ingredients of larger and larger combined fragments: a crucial notion in this sense is that of an algebraic fragment.

Definition 3.2. *A fragment* $\langle \mathcal{L}, T \rangle$ *is said to be an* algebraic fragment *iff* T *satisfies the following conditions:*

(i) T is closed under composition*, i.e. if* $u[x_1, \ldots, x_n] \in T$*, then* $u\,\sigma \in T$*, where* $\sigma : \{x_1 \mapsto t_1, \ldots, x_n \mapsto t_n\}$ *is a substitution such that* $t_i \in T$ *for all* $i = 1, \ldots, n$*;*

(ii) T contains domain variables*, i.e. if* τ *is a type such that some variable of type* τ *occurs free in a term* $t \in T$*, then every variable of type* τ *belongs to* T*;*

(iii) T contains codomain variables*, i.e. if* $t : \tau$ *belongs to* T*, then every variable of type* τ *belongs to* T*.*

Observe that from the above definition it follows that T is closed under renamings, i.e. that if $t \in T$ and σ a renaming, then $t\sigma \in T$. The role of Definition 3.2(i) is that of making fragment combinations non trivial, whereas the other conditions of Definition 3.2 will be needed in order to apply preprocessing purification steps to combined constraints.

Quite often, one is interested in interpreting the terms of a fragment not in the class of all possible structures for the language of the fragment, but only in some selected ones (e.g. when checking satisfiability of some temporal formulae, one might be interested only in checking satisfiability in particular flows of time, those which are for instance discrete or continuous). This is the reason for 'interpreting' fragments:

Definition 3.3. *An* interpreted algebraic fragment *(to be shortened as i.a.f.) is a triple* $\Phi{=}\langle \mathcal{L}, T, \mathcal{S} \rangle$*, where* $\langle \mathcal{L}, T \rangle$ *is an algebraic fragment and* $\mathcal{S}$ *is a class of* $\mathcal{L}$*-structures closed under isomorphisms.*

The set of terms T in an i.a.f. $\Phi{=}\langle \mathcal{L}, T, \mathcal{S} \rangle$ is called the set of Φ-*terms* and the set of types τ such that $t : \tau$ is a Φ-term for some t is called the set of Φ-*types*. A Φ-*variable* is a variable

x^τ such that τ is a Φ-type (or equivalently, a variable which is a Φ-term). It is also useful to *identify a (non-interpreted) algebraic fragment* $\langle \mathcal{L}, T\rangle$ with the interpreted algebraic fragment $\Phi{=}\langle \mathcal{L}, T, \mathcal{S}\rangle$, where $\mathcal{S}$ is taken to be the class of all $\mathcal{L}$-structures.

Definition 3.4. *Given an i.a.f. fragment* Φ*, a* Φ-atom *is an equation* $t_1 = t_2$ *between* Φ*-terms* t_1, t_2 *of the same type; a* Φ-literal *is a* Φ*-atom or a negation of a* Φ*-atom, a* Φ-constraint *is a finite conjunction of* Φ*-literals, a* Φ-clause *is a finite disjunction of* Φ*-literals. Infinite sets of* Φ*-literals (representing an infinite conjunction) are called* generalized Φ-constraints *(provided they contain altogether only finitely many free variables).*

Some Conventions. Without loss of generality, we may assume that $\top$ *is a* Φ*-atom* in every i.a.f. Φ (in fact, to be of any interest, a fragment should at least contain one term t and we can let $\top$ to be $t = t$). As a consequence, $\perp$ will always be a Φ-literal; by convention, however, we shall *include* $\perp$ *among* Φ*-atoms* (hence a Φ-atom is either an equation among Φ-terms - $\top$ included - or it is $\perp$). Since we have $\perp$ as an atom, there is no need to consider the empty clause as a clause, so clauses will be disjunctions of *at least one* literal. The reader should keep in mind these slightly non standard conventions for the whole paper.

A Φ-clause is said *positive* if only Φ-atoms occur in. A Φ-atom $t_1 = t_2$ is closed if and only if t_i is closed ($i \in \{1, 2\}$); the definition of closed Φ-literals, -constraints and -clauses is analogous. For a finite set $\underline{x}$ of variables and an i.a.f. Φ, a $\Phi(\underline{x})$-atom (-term, -literal, -clause, -constraint) is a Φ-atom (-term, -literal, -clause, -constraint) A such that $fvar(A) \subseteq \underline{x}$.

We deal in this paper mainly with the *constraint satisfiability problem* for an interpreted algebraic fragment $\Phi{=}\langle \mathcal{L}, T, \mathcal{S}\rangle$: this is *the problem of deciding whether a* Φ*-constraint is satisfiable in some structure* $\mathcal{A} \in \mathcal{S}$. On the other hand, the *word problem* for Φ is the problem of deciding if the universal closure of a given Φ-atom is true in every structure $\mathcal{A} \in \mathcal{S}$.

The literature on fragments and on decision procedures for fragments is extremely large (in a sense, one may argue that mathematical logic itself consists of studying the various fragments and their syntactic and semantic properties). Notice however that our definition of a fragment refers to an embedding into a higher order typed language: the consequence of this approach is that a given well-known fragment (in the naive sense) can formally be turned into a fragment in our sense *in many ways* and the differences among such ways are crucial when concretely applying the definition of a combined fragment to be given in Section 4.

The reason is the following: although we have not yet given the relevant definition, the reader may imagine that combining algebraic fragments means, roughly speaking, taking the smallest algebraic fragment containing some given ones. Now, when *defining* the set of Φ-terms of a fragment, it does not actually matter whether certain symbols are treated as free constants or as free variables: since every variable is existentially quantified in the definition of a satisfiable constraint, then one may indifferently use free constants or variables in Φ-terms. However, constants are not good to be used as placeholders when defining the composition (=substitution) of terms, so *the (ab)use of free constants reduces the expressive richness of the combined fragments that can be build over the given one.*

Another opportunity in defining the set of Φ-terms of a fragment, is that of taking a final λ-abstraction in order to get rid of free variables.[11] Clearly the choice of closing by λ-abstraction the Φ-terms of a fragment changes the nature of the satisfiability of the resulting constraints (e.g. it makes the difference between local and global satisfiability, in the case of the standard translation of modal propositional formulae). However such a choice has another relevant and more hidden effect: having taken λ-abstraction, we produced higher order terms which are now ready to be *substituted for* higher order variables when taking combined fragments.

Sometimes the above options cannot be used together: for instance, the set of prenex first-order formulae having a certain given prefix shape are *not* an algebraic fragment if the predicate letters in them are treated as second order variables and if first-order variables are λ-abstracted (closure under substitutions of Φ-terms for variables fails).

The moral of this discussion is that our framework is quite general and flexible, but just for this reason, *it needs to be handled with some care.* In next subsection we shall give examples of algebraic fragments (the reader now knows why we will apparently make 'many different copies' of seemingly the same fragment).

We would like to draw the reader's attention to the fact that in Definition 3.2, when formulating the closure under composition requirement for the set of the terms T of an algebraic fragment, we asked that if $t[x_1, \dots, x_n] \in T$ and $u_1, \dots, u_n \in T$, then *precisely* the term $t[u_1/x_1, \dots, u_n/x_n]$ belongs to T (and not just some other term which is $\beta\eta$-equivalent to it, like for instance its $\beta\eta$-normal form). The reason for this strict requirement is that we want a term belonging to a combined i.a.f. to be effectively decomposable into some iterated composition of pure terms (see Subsection 4.1). With the present version of Definition 3.2, there is an evident algorithm for computing such a decomposition. Of course, we did not eliminate the $\beta\eta$-conversion problem in this way, but we simply left to the user of our combination procedure the responsibility of effectively certifying that the terms forming the constraints he is interested to decide for satisfiability really belong (maybe up to $\beta\eta$-equivalence) to the combined fragment to which he is going to apply the procedure. For instance, before claiming that our procedure decides relativized satisfiability in fusions of modal logics, we shall have

[11] For Φ-terms $\varphi[\underline{x}]$ of type Ω, this usually has the effect of taking universal closure in Φ-atoms: the term $\{\underline{x} \mid \top\}$ is usually in Φ, so that taking the Φ-atom $\{\underline{x} \mid \top\} = \{\underline{x} \mid \varphi\}$ amounts to consider the universal closure of φ.

to produce such a certificate (this is the content of Lemma 5.13 below).

Before closing this subsection, we make a little *digression* (not relevant for the comprehension of the remaining part of the paper) about the choice of the word 'algebraic' in order to name our fragments (this digression may also help future development in a more conceptual setting).

It is well-known [39] that higher order (intuitionistic) type theories correspond to elementary toposes; thus our signatures must in particular be related to free toposes. Algebraic fragments in this context are cartesian subcategories of such toposes (we use the name 'cartesian category' for 'category with finite products'). In fact, algebraic fragments are closed under compositions and contains projections, namely variables. Now it is well-known from Lawvere functorial semantics [40] that cartesian categories correspond to equational theories in a similar way as toposes correspond to higher order type theories. Thus, our algebraic fragments, if considered *from ouside*, are just equational (i.e. algebraic) theories. In fact, if we take for instance the algebraic fragment given by the standard translation of modal propositional formulae, if we consider it as a cartesian category by itself and if we then come back to a presentation of it as a first-order equational theory, we get the theory of modal algebras (=Boolean algebras plus meet-preserving operators), namely the theory of the algebras which are used as the standard algebraic semantics for modal logic.

However, the embedding of an algebraic fragment into a higher order language (i.e. the consideration of a specific topos in which a cartesian category is embedded - we recall that one such always exists) gives new information on the *internal* structure of the fragment itself and we want this information *to be part of our data*. In fact, when interpreting an algebraic fragment, we consider not just set-valued cartesian functors having as a domain the cartesian category corresponding to the fragment itself, but just those such functors which are restrictions of set-valued logical functors defined on the bigger recipient topos (in the case of the above presentation of the theory of modal algebras, for instance, this means that we are considering Kripke models, not just arbitrary algebraic models). Secondly, the specification of the recipient topos seems to influence the construction of our combined larger fragments. Thirdly, the internal information on the fragment is useful to identify certain ad hoc operations on the models of the recipient topos and to exploit specific preservation properties (with respect to the formulae in the fragment) of such operations: these preservation properties will be essential ingredients for justifying completeness of combined decision procedures.

3.2 Examples

We give here a list of examples of i.a.f.'s; we shall mainly concentrate on those examples which will play a central role in the positive results of the paper. In all cases, the proof that the properties of Definition 3.2 are satisfied is just sketched or entirely left to the reader (such proofs are all immediate or they reduce to easy inductive arguments based on standard information from Subsection 2.3).

Example 3.5 (Simply Typed λ-Calculus)**.** This is the i.a.f. Φ that one gets by keeping only the terms that can be built by 'omitting any reference to the type Ω'. According to Friedman theorem [22], this i.a.f. has decidable word problem,[12] because $\beta\eta$-normalization can decide equality of Φ-terms in all interpretations. However, constraint satisfiability problem is no longer decidable.

Example 3.6 (First-order equational fragments)**.** Let us consider a first-order language $\mathcal{L} = \langle \mathcal{T}, \Sigma, a \rangle$ (for simplicity, we also assume that $\mathcal{L}$ is one-sorted). Let T be the set of the first-order $\mathcal{L}$-terms and let $\mathcal{S}$ consists of the $\mathcal{L}$-structures which happen to be models of a certain first-order theory in the signature $\mathcal{L}$. Obviously, the triple $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ is an i.a.f.. The Φ-atoms will be equalities between Φ-terms, i.e. first-order atomic formulae of the kind $t_1 = t_2$. Word problem in $\Phi{=}\langle \mathcal{L}, T, \mathcal{S} \rangle$ is standard uniform word problem (as defined for the case of equational theories for instance in [6]), whereas constraint satisfiability problem is the problem of deciding satisfiability of a finite set of equations and inequations.

Example 3.7 (Universal first-order fragments)**.** The previous example disregards the relational symbols of the first-order signature $\mathcal{L}$. To take also them into consideration, it is sufficient to make some slight adjustment: besides first-order terms, also atomic formulae ($\top, \bot$ included), as well as propositional variables (namely variables having type Ω) will be terms of the fragment.[13] The semantic class $\mathcal{S}$ where the fragment is to be interpreted can be taken to be again the class of the models of some first-order theory. Then, for $\Phi{=}\langle \mathcal{L}, T, \mathcal{S} \rangle$ so defined, the constraint satisfiability problem becomes the problem of deciding the satisfiability of an arbitrary finite set of $\mathcal{L}$-literals in the models belonging to $\mathcal{S}$[14] (the complementary problem is equivalent to the problem of deciding validity of a universal first-order formula in $\mathcal{S}$).

We now define different kinds of i.a.f.'s starting from the set F of first-order formulae of a

[12] Remember that, when no semantic class $\mathcal{S}$ is mentioned in the definition of an i.a.f., it is intended that $\mathcal{S}$ consists of all possible interpretations for the language.

[13] Propositional variables are added to the set of terms in order for closure under codomain variables to be satisfied, see Definition 3.2.

[14] $\mathcal{L}$-atomic formulae A (resp. negated $\mathcal{L}$-atomic formulae $\neg A$) can be seen as the Φ-atoms $A = \top$ (resp. $A = \bot$). One should include also equations $A = B$ and inequations $A \neq B$ among $\mathcal{L}$-atomic formulae and/or propositional variables. However, for instance, $A = B$ is satisfiable iff $A \wedge B$ is satisfiable or $\neg A \wedge \neg B$ is satisfiable: this means that, by case splitting, we can anyway reduce satisfiability of Φ-constraints to satisfiability of conjunctions of $\mathcal{L}$-atomic and negated $\mathcal{L}$-atomic formulae.

first-order signature $\mathcal{L}$; for simplicity, let us suppose also that $\mathcal{L}$ is relational and one-sorted (call W its unique sort).

Example 3.8 (Full First-Order Language, plain version)**.** We take T to be the union of F with the sets of the individual variables and of the propositional variables. Of course, $\Phi=\langle\mathcal{L},T\rangle$ so defined is an algebraic fragment, whose types are W and Ω. By Church theorem, both word and constraint satisfiability problem are undecidable here (the two problems reduce to satisfiability of a first-order formula with equality); they may be decidable in case the fragment is interpreted into some specific semantic class $\mathcal{S}$.

In the next example, we build formulae (out of the symbols of our fixed first order relational one-sorted signature $\mathcal{L}$) by using at most N (free or bound) individual variables; however we are allowed to use also second order variables of arity at most K:

Example 3.9 (Full First-Order Language, NK-version)**.** Fix cardinals $K \leq N \leq \omega$ and consider, instead of F, the set F_{NK} of formulae φ that contains at most N (free or bound) individual variables and that are built up by applying boolean connectives and individual quantifiers to atomic formulae of the following two kinds:

– $P(x_{i_1}, \ldots, x_{i_n})$, where P is a relational constant and $x_{i_1}, \ldots, x_{i_n}$ are individual variables (since at most $x_1, \ldots, x_N$ can be used, we require that $i_1, \ldots, i_n \leq N$);

– $X(x_{i_1}, \ldots, x_{i_n})$, where $i_1, \ldots, i_n \leq N$, and X is a variable of type $W^n \to \Omega$ with $n \leq K$ (here W^n abbreviates $W \cdots W$, n-times).

The terms in the algebraic fragment $\Phi^{\mathcal{L}}_{NK} = \langle\mathcal{L}_{NK}, T^{\mathcal{L}}_{NK}\rangle$ are now the terms t such that $t \sim_{\beta\eta} \{x_1, \ldots, x_n \mid \varphi\}$, for some $n \leq K$ and for some $\varphi \in F_{NK}$, with $fvar_W(\varphi) \subseteq \{x_1, \ldots, x_n\}$.[15] Types in such $\Phi^{\mathcal{L}}_{NK}$ are now $W^n \to \Omega$ ($n \leq K$) and this fact makes a big difference with the previous example (the difference will be sensible when combined fragments enter into the picture). Constraint satisfiability problems still reduce to satisfiability problems for sentences: in fact, once second order variables are replaced by the names of the subsets assigned to them by some assignment α in a $\mathcal{L}$-structure, $\Phi^{\mathcal{L}}_{NK}$-atoms like $\{\underline{x} \mid \varphi\} = \{\underline{x} \mid \psi\}$ are equivalent to first-order sentences $\forall\underline{x}(\varphi \leftrightarrow \psi)$ and conversely any first-order sentence θ (with at most N bound individual variables) is equivalent to the $\Phi^{\mathcal{L}}_{NK}$-atom $\theta = \top$.

The cases $N = 1, 2$ are particularly important, because in these cases the satisfiability problem for sentences (and hence also constraint satisfiability problems in our fragments) becomes decidable [43],[52], [47], [53], [18].

[15]We need to use $\beta\eta$-equivalence here to show that the properties of Definition 3.2 (namely closure under composition and under domain/codomain variables) are satisfied.

We mention that the previous two examples admit very important weaker versions in which some of the first order operators are omitted. For instance, if universal quantifiers and negations are omitted, constraint satisfiability in the $\lambda_{\omega\omega}$-version becomes the problem of deciding whether a geometric sequent is entailed by a finitely axiomatized geometric theory (for this terminology, see [4] or some book in categorical logic, like [45]).

Further examples can be obtained by using the large information contained in the textbook [15] (see also [20]). We shall continue here by investigating fragments that arise from research in knowledge representation area, especially in connection to modal and description logics.

Example 3.10 (Modal/Description Logic Fragments, global version)**.** A *modal signature* is a set O_M, whose elements are called unary 'Diamond' modal operators (the case of n-ary modal operators does not create special difficulties and it is left to the reader). O_M-modal formulae are built up from a countable set of propositional variables $\mathrm{x}, \mathrm{y}, \mathrm{z}, \ldots$ by applying $\top, \bot, \neg, \wedge, \vee$ as well as the operators $\Diamond_k \in O_M$.

With every modal signature O_M we associate the first-order signature $\mathcal{L}_M$, containing a unique sort W and, for every $\Diamond_k \in O_M$, a relational constant R_k of type $WW \to \Omega$. Suppose we are given a bijective correspondence $\mathrm{x} \longmapsto X$ between propositional variables and second order variables of type $W \to \Omega$. Given an O_M-modal formula φ and a variable w of type W, the standard translation $ST(\varphi, w)$ is the $\mathcal{L}_M$-term of type Ω inductively defined as follows:

$$\begin{aligned}
ST(\top, w) &= \top \\
ST(\bot, w) &= \bot \\
ST(\mathrm{x}, w) &= X(w) \\
ST(\neg\psi, w) &= \neg ST(\psi, w) \\
ST(\psi_1 \vee \psi_2, w) &= ST(\psi_1, w) \vee ST(\psi_2, w) \\
ST(\psi_1 \wedge \psi_2, w) &= ST(\psi_1, w) \wedge ST(\psi_2, w) \\
ST(\Diamond\psi, w) &= \exists v(R(w, v) \wedge ST(\psi, v))
\end{aligned}$$

where v is a variable of type W (different from w). Let T_M be the set of those $\mathcal{L}_M$-terms t for which there exists a modal formula φ s.t. $t \sim_{\beta\eta} \{w \mid ST(\varphi, w)\}$. The pair $\langle \mathcal{L}_M, T_M \rangle$ is an algebraic fragment and it becomes an i.a.f. $\Phi_M = \langle \mathcal{L}_M, T_M, \mathcal{S}_M \rangle$ if we specify also a class $\mathcal{S}_M$ of $\mathcal{L}_M$-structures closed under isomorphisms (notice that $\mathcal{L}_M$-structures, usually called *Kripke frames* in modal logic, are just sets endowed with a binary relation R_k for every $\Diamond_k \in O_M$).

Φ_M-constraints can be equivalently represented in the form

$$\{w \mid ST(\psi, w)\} = \{w \mid \top\} \wedge \{w \mid ST(\varphi_1, w)\} \neq \{w \mid \bot\} \wedge \cdots \wedge \{w \mid ST(\varphi_n, w)\} \neq \{w \mid \bot\};$$

they are satisfied iff there exists a Kripke model[16] based on a frame in $\mathcal{S}_M$ in which ψ holds globally (namely in every state), whereas $\varphi_1, \dots, \varphi_n$ hold in some states $s_1, \dots, s_n$, respectively. If $\mathcal{S}_M$ is closed under disjoint unions, we can limit ourselves to the case $n = 1$: thus constraint satisfiability problem becomes, in the description logics terminology, just the relativized satisfiability problem for a given concept description wrt to a given T-Box (we call *T-Box* a Φ_M-atom like $\{w \mid ST(\psi, w)\} = \{w \mid \top\}$).[17]

Example 3.11 (Modal/Description Logic Fragments, local version)**.** If we want to capture A-Box reasoning too, we need to build a slightly different fragment. The type-theoretic signature $\mathcal{L}_{ML}$ of our fragment $\langle \mathcal{L}_{ML}, T_{ML} \rangle$ is again $\mathcal{L}_M$, but T_{ML} now contains: a) the set of terms which are $\beta\eta$-equivalent to terms of the kind $ST(\varphi, w)$ (these terms are called 'concept assertions'); b) the terms of the kind $R_k(v, w)$ (these terms are called 'role assertions'); c) the variables of type W, Ω and $W \to \Omega$.

The pair $\langle \mathcal{L}_{ML}, T_{ML} \rangle$ is an algebraic fragment and it becomes an interpreted algebraic fragment $\Phi_{ML} = \langle \mathcal{L}_{ML}, T_{ML}, \mathcal{S}_{ML} \rangle$ if we specify also a class $\mathcal{S}_{ML}$ of $\mathcal{L}_{ML}$-structures closed under isomorphisms. Constraints in this fragment can be represented as conjunctions of concept assertions and role assertions,[18] plus in addition: a) identities among individual names (i.e. among variables of type W); b) identities among atomic concepts (i.e. among second order variables of type $W \to \Omega$); c) propositional variables (i.e. variables of type Ω); d) negations of identities among atomic concepts; e) negations of propositional variables; f) negations of role assertions; g) negations of identities among individual names.

Since a)-b)-c)-d)-e) can be eliminated without loss of generality,[19] we can conclude that Φ_{LM}-constraints are just standard A-Boxes with, in addition, negations of role assertions and of identities among individual names.[20] Let us call *A-Boxes* these slightly more general constraints and let us reserve the name of *positive A-Boxes* to conjunctions of concept assertions and role assertions.

[16] A Kripke model is a Kripke frame together with an assignment of subsets for second order variables of type $W \to \Omega$.

[17] Usually, a T-Box is defined as a conjunction of 'generalized concept inclusions' that are required to hold globally: this can be reduced to the requirement for a single formula to hold globally, because all boolean connectives are at our disposal.

[18] These arise because $\top \equiv ST(\top, w)$ is a term of the fragment. One should consider also atoms of the kind $ST(\varphi, w) = ST(\psi, v)$, $ST(\varphi, w) = R(v_1, v_2)$, etc. (and their negations). However, we can eliminate them by Boolean case splitting, like in Example 3.7.

[19] All variable identities can be eliminated by replacements; negations of identities among atomic concepts can be replaced by concept assertions involving fresh variables. Propositional variables and their negations do not interact with the remaining part of the constraint and can be ignored.

[20] Traditional A-Boxes automatically include all negations of identities among distinct individual variables by the so-called 'unique name assumption'.

Example 3.12 (Modal/Description Logic Fragments, full version)**.** If we want to deal with satisfiability of an A-Box wrt a T-Box, it is sufficient to join the two previous fragments. More precisely, we can build the fragments $\Phi_{MF} = \langle \mathcal{L}_{MF}, T_{MF}, \mathcal{S}_{MF} \rangle$, where $\mathcal{L}_{MF} = \mathcal{L}_M$ and $T_{MF} = T_M \cup T_{ML}$. Types in this fragment are W, Ω and $W \to \Omega$; constraints are conjunctions of a T-Box and an A-Box.

Example 3.13 (Modal/Description Logic Fragments, non-normal case)**.** If we want to consider the case in which some of the operators in O_M are non-normal, we can use higher order constants $f_k : (W \to \Omega) \to (W \to \Omega)$ (instead of binary relations $R_k : WW \to \Omega$) and define a different translation. Such a translation $NT(\varphi, w)$ differs from $ST(\varphi, w)$ for the inductive step relative to modal operators which now reads as follows:

$$NT(\Diamond_k \psi, w) = f_k(\{w \mid NT(\psi, w)\})(w).$$

Now global, local and full algebraic fragments can be defined as in the normal case. If the easy extension to n-ary non normal cases is included and if we also interpret the resulting fragments, we get precisely the abstract description systems of [10].

Guarded and packed guarded fragments were introduced as generalizations of modal fragments [1], [33], [46]: in fact, they form classes of formulae which are remarkably large but still inherit relevant syntactic and semantic features of the more restricted modal formulae. In particular, guarded and packed guarded formulae are decidable for satisfiability (with the appropriate settings, decision procedures can be obtained also by running standard superposition provers [25]).

For simplicity, we give here the instructions on how to build only one version of the packed guarded fragment with equality (other versions can be built by following the methods we used above for the first-order and the modal cases). We notice that packed guarded fragments without equality are also important: to built them it is sufficient to erase any reference to the equality predicate in the relevant definitions.

Example 3.14 (Packed Guarded Fragments)**.** Let us consider a first-order one-sorted relational signature $\mathcal{L}_G$. A *guard* π is a $\mathcal{L}_G$-formula like $\bigwedge_{i=1}^{k} \pi_i$, where:

- π_i is obtained by applying existential quantifiers to atomic formulae $P_i(x_{i1}, \dots, x_{in_i})$ where the P_i are constants of type $W^{n_i} \to \Omega$ and $x_{i1}, \dots, x_{in_i}$ are variables of type W;
- for all $x_1, x_2 \in fvar(\pi)$, there exists an $i \in \{1, \dots, k\}$ such that $\{x_1, x_2\} \subseteq fvar(\pi_i)$.

We define the *packed guarded formulae* as follows:

- if $X : W \to \Omega$ and $x : W$ are variables, $X(x)$ is a packed guarded formula;

– if $P : W^n \to \Omega$ is a constant and $y_1 : W, \dots, y_n : W$ are variables, $P(y_1, \dots, y_n)$ is a packed guarded formula;

– if φ is a packed guarded formula, $\neg\varphi$ is a packed guarded formula;

– if φ_1 and φ_2 are packed guarded formulae, $\varphi_1 \wedge \varphi_2$ and $\varphi_1 \vee \varphi_2$ are packed guarded formulae;

– if φ is a packed guarded formula and π is a guard such that $fvar_W(\varphi) \subseteq fvar(\pi)$, then $\forall \underline{y}(\pi[\underline{x}, \underline{y}] \to \varphi[\underline{x}, \underline{y}])$ and $\exists \underline{y}(\pi[\underline{x}, \underline{y}] \wedge \varphi[\underline{x}, \underline{y}])$ are packed guarded formulae.[21]

Notice that we used second order variables of type $W \to \Omega$ only (and not of type $W^n \to \Omega$ for $n > 1$): the reason, besides the applications to combined decision problems we have in mind, is that we want constraint problems to be equivalent to sentences which are still packed guarded, see below. Packed guarded formulae not containing variables of type $W \to \Omega$ are called *elementary* (or first-order) packed guarded formulae.

If we let T_G be the set of $\mathcal{L}_G$-terms t such that t is $\beta\eta$-equivalent to a term of the kind $\{w \mid \varphi\}$ (where φ is a packed guarded formula such that $fvar_W(\varphi) \subseteq \{w\}$), then the pair $\langle \mathcal{L}_G, T_G \rangle$ is an algebraic fragment. The only type in this fragment is $W \to \Omega$ and constraint satisfiability problem in this fragment is equivalent to satisfiability of guarded sentences: this is because, in case φ_1, φ_2 are packed guarded formulae with $fvar_W(\varphi_i) \subseteq \{w\}$ (for $i = 1, 2$), then $\{w \mid \varphi_1\} = \{w \mid \varphi_2\}$ is equivalent to $\forall w(\varphi_1 \leftrightarrow \varphi_2)$ which is packed guarded (just use $w = w$ as a guard).

3.3 Reduced Fragments and Residues

If $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ is an i.a.f. and $\underline{x}$ is a finite set of Φ-variables, let us $\Phi(\underline{x})$ denote the Φ-clauses whose free variables are among the $\underline{x}$. If Γ is a set of such $\Phi(\underline{x})$-clauses and $C \equiv L_1 \vee \cdots \vee L_k$ is a $\Phi(\underline{x})$-clause, we say that C is a Φ-*consequence of* Γ (written $\Gamma \models_\Phi C$), iff the set of formulae $\Gamma \cup \{\neg L_1, \dots, \neg L_k\}$ is not Φ-satisfiable (i.e. iff such a set is not satisfiable in any $\mathcal{A} \in \mathcal{S}$).

The notion of consequence is too strong for certain applications; for instance, when we simply need to delete certain deductively useless data, a weaker notion of redundancy (based e.g. on subsumption) is preferable. We shall give abstract notions of redundancy and of consequences enumeration in this subsection, however these notions are less sophisticated than similar notions introduced within saturation-based theorem proving (see e.g. [12]). The reason why we need them here is that they will make our combined decision procedure more flexible, as explained below.

[21] If $\underline{y} = \{y_1, \dots, y_n\}$, then $\forall \underline{y}$ means $\forall y_1 \cdots \forall y_n$ and $\exists \underline{y}$ means $\exists y_1 \cdots \exists y_n$.

Our abstract axiomatization of a notion of redundancy is the following (recall that we conventionally included $\top$ and $\bot$ among Φ-atoms in any i.a.f. Φ):

Definition 3.15. *A* redundancy notion *for a fragment* Φ *is a* recursive *binary relation* Red_Φ *between a finite set of* Φ*-clauses* Γ *and a* Φ*-clause* C *satisfying the following properties:*

(i) $Red_\Phi(\Gamma, C)$ *implies* $\Gamma \models_\Phi C$ *(soundness);*

(ii) $Red_\Phi(\emptyset, \top)$ *and* $Red_\Phi(\{\bot\}, C)$ *both hold;*

(iii) $Red_\Phi(\Gamma, C)$ *and* $\Gamma \subseteq \Gamma'$ *imply* $Red_\Phi(\Gamma', C)$ *(monotonicity);*

(iv) $Red_\Phi(\Gamma, C)$ *and* $Red_\Phi(\Gamma \cup \{C\}, D)$ *imply* $Red_\Phi(\Gamma, D)$ *(transitivity);*

(v) *if* C *is subsumed by some* $C' \in \Gamma$,[22] *then* $Red_\Phi(\Gamma, C)$ *holds.*

Whenever a redundancy notion Red_Φ *is fixed, we say that* C *is* Φ*-redundant wrt* Γ *when* $Red_\Phi(\Gamma, C)$ *holds.*

For example, the *minimum* redundancy notion is obtained by stipulating that $Red_\Phi(\Gamma, C)$ holds precisely when ($\bot \in \Gamma$ or $C \equiv \top$ or $C \equiv \top \vee D$ or C is subsumed by some $C' \in \Gamma$). On the contrary, if the constraint solving problem for Φ is decidable, there is a maximum redundancy notion (called the *full* redundancy notion) given by the Φ-consequence relation. In fact, it is evident that a recursive procedure for Φ-constraint solving is a recursive procedure deciding $\Gamma \models_\Phi C$, for finite Γ.

Let $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ be an i.a.f. on the signature $\mathcal{L} = \langle \mathcal{T}, \Sigma, a \rangle$ and let $\mathcal{L}_0 = \langle \mathcal{T}_0, \Sigma_0, a_0 \rangle$ be a subsignature of $\mathcal{L}$. The i.a.f. *restricted to* $\mathcal{L}_0$ is the i.a.f. $\Phi_{|\mathcal{L}_0} = \langle \mathcal{L}_0, T_{|\mathcal{L}_0}, \mathcal{S}_{|\mathcal{L}_0} \rangle$ that is so defined:

– $T_{|\mathcal{L}_0}$ is the set of terms obtained by intersecting T with the set of $\mathcal{L}_0$-terms;

– $\mathcal{S}_{|\mathcal{L}_0}$ consists of the structures of the kind $\mathcal{A}_{|\mathcal{L}_0}$, varying $\mathcal{A} \in \mathcal{S}$.

An i.a.f. $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$ is said to be a $\mathcal{L}_0$-*subfragment* (or simply a subfragment, leaving the subsignature $\mathcal{L}_0 \subseteq \mathcal{L}$ as understood) of $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ iff $T_0 \subseteq T_{|\mathcal{L}_0}$ and $\mathcal{S}_0 \supseteq \mathcal{S}_{|\mathcal{L}_0}$. In this case, we may also say that Φ is an *expansion* (or an *extension*) of Φ_0.

Given a set Γ of $\Phi(\underline{x})$-clauses and a redundancy notion Red_{Φ_0} on a subfragment Φ_0 of Φ, we call Φ_0-*basis for* Γ a set Δ of $\Phi_0(\underline{x}_0)$-clauses such that (here $\underline{x}_0$ collects those variables among the $\underline{x}$ which happen to be Φ_0-variables):

(i) all clauses $D \in \Delta$ are positive and are such that $\Gamma \models_\Phi D$;[23]

[22] As usual, this means that every literal of C' is also in C.

[23] Recall that we conventionally included $\bot$ among Φ-atoms, so $\bot$ is considered as a positive clause.

(ii) every positive $\Phi_0(\underline{x}_0)$-clause C such that $\Gamma \models_\Phi C$ is Φ_0-redundant with respect to Δ.

Since we will be interested in exchange information concerning consequences over shared signatures, we need a notion of a residue, like in partial theory reasoning (see e.g. [14] for comprehensive information on the subject and the relevant pointers to the literature). Again, we prefer an abstract approach and treat residues as clauses which are recursively enumerated by a suitable device (the device may for instance be an enumerator of certain proofs of a calculus, but there is no need to think of it in this way):

Definition 3.16. *Suppose we are given a subfragment Φ_0 of a fragment Φ. A* positive residue Φ-enumerator for Φ_0 *(often shortened as Φ-p.r.e.) is a recursive function mapping a finite set $\underline{x}$ of Φ-variables, a finite set Γ of $\Phi(\underline{x})$-clauses and a natural number i to a Φ_0-clause $Res^{\underline{x}}_\Phi(\Gamma, i)$ (to be written simply as $Res_\Phi(\Gamma, i)$) in such a way that:*

- *$Res_\Phi(\Gamma, i)$ is a positive clause;*
- *$fvar(Res_\Phi(\Gamma, i)) \subseteq \underline{x}$;*
- *$\Gamma \models_\Phi Res_\Phi(\Gamma, i)$ (soundness).*

Any Φ_0-clause of the kind $Res_\Phi(\Gamma, i)$ (for some $i \geq 0$) will be called a Φ_0-residue *of Γ.*

Having also a redundancy notion for Φ_0 at our disposal, we can axiomatize the notion of an 'optimized' (i.e. of a non-redundant) Φ-p.r.e. for Φ_0. The version of the Nelson-Oppen combination procedure we give in Subsection 4.2 has non-redundant p.r.e.'s as main ingredients and it is designed to be 'self-adaptive' for termination in the relevant cases when termination follows from our results. These are basically the noetherian and the locally finite cases mentioned in Subsection 3.4, where p.r.e.'s which are non redundant with respect to the full redundancy notion usually exist and enjoy the termination property below.

Definition 3.17. *A Φ-p.r.e. Res_Φ for Φ_0 is said to be* non-redundant *(wrt a redundancy notion Red_{Φ_0}) iff it satisfies also the following properties for every $\underline{x}$, for every finite set Γ of $\Phi(\underline{x})$-clauses and for every $i \geq 0$ (we write $\Gamma_{|\Phi_0}$ for the set of clauses in Γ which are Φ_0-clauses):*

(i) *if $Res_\Phi(\Gamma, i)$ is Φ_0-redundant with respect to $\Gamma_{|\Phi_0} \cup \{Res_\Phi(\Gamma, j) \mid j < i\}$, then $Res_\Phi(\Gamma, i)$ is either $\bot$ or $\top$;*

(ii) *if $\bot$ is Φ_0-redundant with respect to $\Gamma_{|\Phi_0} \cup \{Res_\Phi(\Gamma, j) \mid j < i\}$, then $Res_\Phi(\Gamma, i)$ is equal to $\bot$;*

(iii) *if $Res_\Phi(\Gamma, i)$ is equal to $\top$, then $\Gamma_{|\Phi_0} \cup \{Res_\Phi(\Gamma, j) \mid j < i\}$ is a Φ_0-basis for Γ.*

Definition 3.18. *A non-redundant Φ-p.r.e. for Φ_0 is said to be* complete *iff for every $\underline{x}$, for every finite set Γ of $\Phi(\underline{x})$-clauses and for every positive $\Phi_0(\underline{x})$-clause C, we have that $\Gamma \models_\Phi C$ implies that C is Φ_0-redundant wrt $\Gamma_{|\Phi_0} \cup \{Res_\Phi(\Gamma, j) \mid j \leq i\}$ for some i.*

A non-redundant Φ-p.r.e. Res_Φ is said to be terminating *iff for for every $\underline{x}$, for every finite set Γ of $\Phi(\underline{x})$-clauses there is an i such that $Res_\Phi(\Gamma, i)$ is equal to $\perp$ or to $\top$.*

Let us make a few comments on Definition 3.17: first, only non redundant residues can be produced at each step (condition (i)), if possible. If this is not possible, this means that all the relevant information has been accumulated (a Φ_0-basis has been reached). In this case, if the inconsistency $\perp$ is discovered (in the sense that it is perceived as redundant), then the residue enumeration in practice stops, because it becomes constantly equal to $\perp$ (condition (ii)). The tautology $\top$ has the special role of marking the opposite outcome: it is the residue that is returned precisely when Γ is consistent and a Φ_0-basis has been produced, meaning that all relevant semantic consequences of Γ have been discovered (conditions (ii)-(iii)).

If the redundancy notion we use is trivial (i.e. it is the minimum one), then only very mild corrections are needed for a Φ-p.r.e. for Φ_0 to become non-redundant: apart from minor ad hoc modifications,[24] we only need to make it constantly equal to $\perp$, as soon as $\perp$ becomes redundant in the enumeration. This observation shows that, in practice, any Φ-p.r.e. for Φ_0 can be made non-redundant and can consequently be used as input of our combined decision procedure.

The role of $\top$ as a residue is precious in case for some special reasons (typically exemplified in computational algebra, see Subsection 3.5 below), we have an effective procedure which is able to recognize whether a given set of positive Φ_0-clauses forms a Φ_0-basis for Γ with respect to the full redundancy notion: if these *full* Φ_0*-bases* for Γ can be effectively recognized and if also Φ_0-consequence is decidable, we can always turn a complete Φ-p.r.e. for Φ_0 into a non-redundant one with respect to the full redundancy notion. The advantage of this optimization is that the combined decision procedure of Subsection 4.2, after getting $\top$ or $\perp$ as residues, automatically recognizes that the residue exchange is over and halts.

[24] These modifications are possible provided that there are countably many closed Φ_0-atoms equivalent to $\top$ but syntactically different from it: if there are such infinitely many closed Φ_0-atoms which are 'copies' of $\top$, then we can replace $Res_\Phi(\Gamma, i)$ by one of them in case $Res_\Phi(\Gamma, i)$ is redundant with respect to $\Gamma_{|\Phi_0} \cup \{Res_\Phi(\Gamma, j) \mid j < i\}$. By using this trick, conditions (i) and (iii) of Definition 3.17 can be forced, if the underlying redundancy notion for Φ_0 is the *minimum* one. The hypothesis that Φ_0 is endowed with such infinitely many 'copies' of $\top$ is not really restrictive and can be always obtained by slight modifications of Φ_0.

3.4 Noetherian, Locally Finite and Convex Fragments

The above mentioned optimization for p.r.e.'s usually apply to the cases in which the 'small' fragment Φ_0 is noetherian: this important notion is borrowed from Algebra. Noetherianity conditions known from Algebra [44] say that there are no infinite ascending chains of congruences. In finitely presented algebras, congruences are represented as sets of equations among terms, hence noetherianity can be expressed there by saying that there are no infinite ascending chains of sets of atoms, modulo logical consequence. If we translate this into our general setting, we get the following definition.

An i.a.f. Φ_0 is called *noetherian* if and only if for every finite set of variables $\underline{x}$, every infinite ascending chain

$$\Theta_1 \subseteq \Theta_2 \subseteq \cdots \subseteq \Theta_n \subseteq \cdots$$

of sets of $\Phi_0(\underline{x})$-atoms is eventually constant for Φ_0-consequence (meaning that there is an n such that for all m and $A \in \Theta_m$, we have $\Theta_n \models_{\Phi_0} A$).

An i.a.f. Φ_0 is said to be *effectively locally finite* iff

(i) the set of Φ_0-types is recursive and constraint satisfiability problem for Φ_0 is decidable;

(ii) for every finite set of Φ_0-variables $\underline{x}$, there are finitely many computable $\Phi_0(\underline{x})$-terms $t_1, \ldots, t_n$ such that for every further $\Phi_0(\underline{x})$-term u one of the literals $t_1 \neq u, \ldots, t_n \neq u$ is not Φ_0-satisfiable (that is, in the class of the structures in which Φ_0 is interpreted, every $\Phi_0(\underline{x})$-term is equal, as an interpreted function, to one of the t_i).

The terms $t_1, \ldots, t_n$ in (ii) are called the $\underline{x}$-representative terms of Φ_0.

Effective local finiteness is often used in order to make Nelson-Oppen procedures terminating [28], [11], [4]:[25] we shall see however that noetherianity (which is clearly a weaker condition) is already sufficient for that, once it is accompanied by a suitable effectiveness condition.

Proposition 3.19. *In a noetherian fragment Φ_0 every infinite ascending chain of sets of positive $\Phi_0(\underline{x})$-clauses is eventually constant for Φ_0-consequence.*

Proof. Suppose not; in this case there are infinitely many positive $\Phi_0(\underline{x})$-clauses $C_1, C_2, \ldots$, such that for all i, the clause C_i is not a Φ_0-consequence of $\{C_k \mid k < i\}$.[26]

Let us build a chain of trees $T_0 \subseteq T_1 \subseteq T_2 \subseteq \cdots$, whose nodes are labeled by $\Phi_0(\underline{x})$-atoms as follows. T_0 consists of the root only, which is labeled $\top$. Suppose T_{i-1} is already built and

[25] Notice that the above definition of local finiteness becomes slightly redundant in the first order universal case considered in these papers.

[26] This is an equivalent formulation of the negation of the statement of the proposition.

consider the clause $C_i \equiv B_1 \vee \cdots \vee B_m$. To build T_i, do the following for every leaf K of T_{i-1} (let the branch leading to K be labeled by $A_1, \dots, A_k$): append new sons to K labeled $B_1, \dots, B_m$, respectively, if C_i is not a Φ_0-consequence of $\{A_1, \dots, A_k\}$ (if it is, do nothing for the leaf K).

Consider now the union tree $T := \bigcup T_i$: since, whenever a node labeled A_{k+1} is added, A_{k+1} is not a Φ_0-consequence of the formulae labeling the predecessor nodes, by the noetherianity of Φ_0, all branches are then finite and by König lemma the whole tree is itself finite. This means that for some index j, the examination of clauses C_i (for $i > j$) did not yield any modification of the already built tree. Now, C_{i+1} is not a Φ_0-consequence of $\{C_1, \dots, C_i\}$: this means that there is a structure in the class of the structures in which Φ_0 is interpreted, in which under some assignment α, all atoms of C_{i+1} are false and one atom in each of the $C_1, \dots, C_i$ is true. This contradicts the fact that the tree T_i has not been modified in step $i+1$. $\square$

Suppose that Φ_0 is noetherian and that Φ is an expansion of it: by the above proposition, it is immediate to see that every finite set of $\Phi(\underline{x})$-clauses Γ has a finite full Φ_0-basis (i.e. there is a finite Φ_0-basis for Γ with respect to the full redundancy notion). The following noetherianity requirement for a p.r.e. is intended to be nothing but an effectiveness requirement for the computation of finite full Φ_0-bases.

A Φ-p.r.e. Res_Φ for a noetherian fragment Φ_0 is said to be *noetherian* iff it is non redundant with respect to the full redundancy notion for Φ_0. An immediate consequence of Proposition 3.19 is that:

Proposition 3.20. *A noetherian Φ-p.r.e. Res_Φ for Φ_0 is terminating and also complete.*

The following proposition is also easy, but let us fix it for future reference:

Proposition 3.21. *If Φ_0 is effectively locally finite and Φ is any extension of it having decidable constraint satisfiability problem, then there always exists a noetherian Φ-p.r.e. for Φ_0.*

Proof. Once a $\Phi(\underline{x})$-constraint Γ is given, first check Γ for consistency: if it is inconsistent, the residue enumeration just returns $\perp$. If it is consistent, test the finitely many $\Phi_0(\underline{x}_0)$-positive clauses built up from the representative $\underline{x}_0$-terms of Φ_0 for being a Φ-consequence of Γ (here $\underline{x}_0$ are those variables among the $\underline{x}$'s which are Φ_0-variables). To build the desired Φ-p.r.e., it is then sufficient to list (up to Φ_0-redundancy, which can be effectively checked) the clauses whose test is positive and to give $\top$ as a final output. $\square$

We shall see that, when dealing with noetherian p.r.e.'s over a noetherian shared fragment, the combination procedure of Subsection 4.2 becomes automatically terminating.

Noetherianity is the essential ingredients for the termination of Nelson-Oppen combination procedures; on the other hand, for efficiency, convexity is the crucial property, as it makes the combination procedure deterministic [50]. Following an analogous notion introduced in [55], we say that an i.a.f. Φ is Φ_0-*convex* (here Φ_0 is a subfragment of Φ) iff every finite set Γ of Φ-literals having as a Φ-consequence the disjunction of $n > 1$ Φ_0-atoms, actually has as a Φ-consequence one of them.[27] Similarly, a Φ-p.r.e. for Φ_0 is Φ_0-*convex* iff, for each finite set of Φ-literals Γ, $Res_\Phi(\Gamma, i)$ is always a Φ_0-atom (recall that by our conventions, this includes the case in which it is $\top$ or $\bot$). Any complete non-redundant Φ-p.r.e. for Φ_0 can be turned into a Φ_0-convex complete non-redundant Φ-p.r.e. for Φ_0, in case Φ is Φ_0-convex. Thus the combination procedure of Subsection 4.2 is designed in such a way that it becomes automatically *deterministic* if the component fragments are both convex with respect to the shared fragment.

3.5 Further Examples

In this subsection we collect some examples from Algebra which concretely illustrate the notions of a noetherian fragment and of a noetherian Φ-p.r.e. (the subsection can be skipped by readers mainly interested in examples of a different nature).[28]

Example 3.22 (K-algebras)**.** Given a field K, let us consider the one-sorted language $\mathcal{L}_{Kalg}$, whose signature contains the constants $0, 1$ of type V (V is the unique sort of $\mathcal{L}_{Kalg}$), the two binary function symbols $+, \circ$ of type $VV \to V$, the unary function symbol $-$ of type $V \to V$ and a K-indexed family of unary function symbols g_k of type $V \to V$. We consider the i.a.f. $\Phi_{Kalg} = \langle \mathcal{L}_{Kalg}, T_{Kalg}, \mathcal{S}_{Kalg} \rangle$ where T_{Kalg} is the set of first order terms in the above signature (we shall use infix notation for $+$ and write kv, $v_1 v_2$ for $g_k(v)$, $\circ(v_1, v_2)$, respectively). Furthermore, the class $\mathcal{S}_{Kalg}$ consists of the structures which happen to be models for the theory of (commutative, for simplicity) K-algebras: these are structures having both a commutative ring with unit and a K-vector space structure (the two structures are related by the equations $k(v_1 v_2) = (kv_1)v_2 = v_1(kv_2)$). It is clear that Φ_{Kalg} is an algebraic fragment. The Φ_{Kalg}-atoms have a normalized representation as $p = 0$, where p is a polynomial; the theory of K-algebras is equational, hence convex, so that the constraint satisfiability problem is just the problem of deciding whether an equation $p = 0$ is a logical consequence of a finite number of equations $\{p_1 = 0, \dots, p_n = 0\}$. Since the polynomial ring $K[x_1, \dots, x_n]$ is the free K-algebra on n-generators, this problem is equivalent to the membership of the polynomial p to the ideal $\langle p_1, \dots, p_n \rangle$ generated by the polynomials $p_1, \dots, p_n$. The Buchberger's algorithm solves the problem by computing the Gröbner basis associated to the ideal $\langle p_1, \dots, p_n \rangle$ (Gröbner basis

[27] When we say that a fragment Φ is *convex* tout court, we mean that it is Φ-convex. The fragments $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ analyzed in Example 3.6 are convex in case $\mathcal{S}$ is the class of the models of a first-order Horn theory.

[28] However, Example 3.24 below may be of some interest for people working in software verification.

can morally be considered as confluent and terminating rewriting systems for the conditional word problem in K-algebras).

Example 3.23 (K-vector spaces)**.** As a subfragment of the previous fragment, let us consider the fragment $\Phi_K = \langle \mathcal{L}_K, T_K, \mathcal{S}_K \rangle$, where we forgot in the signature the ring multiplication $\circ$ and the ring unit 1; the structures in $\mathcal{S}_K$ are now the K-vector spaces and the terms in T_K (namely the first order terms in $\mathcal{L}_K$) can consequently be represented as linear homogeneous polynomials with K-coefficients. Φ_K is noetherian because there does not exist an infinite properly ascending chain of subspaces of a finitely generated K-vector space,[29] but it is not locally finite because linear polynomials in n unknowns are infinite (if coefficients are infinite). In order to obtain a noetherian Φ_{Kalg}-p.r.e. for Φ_K through Buchberger algorithm, we need membership of a linear polynomial to a finitely generated ideal to be decided only by linear reduction rules (in this case, our Φ_{Kalg}-p.r.e. for Φ_K *may simply extract the linear polynomials from a Gröbner basis*). This can be obtained through a further requirement on the admissibility of the term ordering (in fact, since we need that linear terms can be rewritten to linear ones only, not every admissible order is suitable). As an example of an order matching such a requirement, one can take for instance the following one [6]: say that $x_1^{m_1} \cdots x_n^{m_n} \succ x_1^{l_1} \cdots x_n^{l_n}$ holds iff (1) $\sum_{i=1}^n k_i > \sum_{i=1}^n l_i$ or (2) $\sum_{i=1}^n k_i = \sum_{i=1}^n l_i$ and $(k_1, \cdots, k_n) >_{lex}^n (l_1, \cdots, l_n)$ (where $>_{lex}^n$ is the lexicographic extension of the natural numbers ordering). We can now briefly outline how to extract from the Buchberger's algorithm a *noetherian* Φ_{Kalg}-p.r.e. for the subfragment Φ_K. First notice that, according to the above requirement on the ordering for polynomials, only linear rewrite rules can be used to reduce a linear term. Thus in order to compute a full Φ_K-basis for the Φ_{Kalg}-constraint $\Gamma = \{p_1 = 0, \ldots, p_n = 0, r_1 \neq 0, \ldots, r_m \neq 0\}$ it is sufficient to: (a) run the Buchberger's algorithm procedure on $\{p_1, \ldots, p_n\}$ (let Q be the corresponding Gröbner basis); (b) normalize the r_j wrt $\to_Q$ to check consistency; (c) if consistency holds, return the equations $q = 0$, where q in the basis is a linear polynomial (if consistency does not hold, simply return $\perp$).

Example 3.24. For general algebraic reasons [44], the observations of the previous example concerning noetherianity and convexity of the i.a.f. $\Phi_K = \langle \mathcal{L}_K, T_K, \mathcal{S}_K \rangle$ applies also in the analogous case of the theory of modules over a noetherian ring K. This implies that the theory of abelian groups is a noetherian fragment and, since integer or rational arithmetic (namely the theory of the integers or of the rationals under addition) is an extension of the latter, it is noetherian too (however noetherianity is lost if we add the ordering to the language).

We give a further example (to be combined in Subsection 5.2 with the fragment of Example 3.22):

[29] For a similar reason, concerning chains of ideals in $K[x_1, \ldots, x_n]$ this time, Φ_{Kalg} is noetherian too.

Example 3.25 (K-vector spaces endowed with an endomorphism)**.** This is an expansion of the fragment in Example 3.23. We augment the signature $\mathcal{L}_K$ with a unary function symbol f and, in order to interpret the fragment, we take K-vector spaces endowed with an endomorphism (call this fragment $\Phi_{Kend} = \langle \mathcal{L}_{Kend}, T_{Kend}, \mathcal{S}_{Kend} \rangle$ and the structures in $\mathcal{S}_{Kend}$ f-K-vector spaces). Terms in this fragment formally represent vectors in finitely generated free $\mathcal{S}_{Kend}$-algebras and hence normalize to the form $k_1 f^{m_1}(x_{i_1}) + \cdots + k_n f^{m_n}(x_{i_n})$ ($k_j \in K$, $j \in \{1, \dots, n\}$). Given the convexity of the fragment, constraint satisfiability amounts to decide if an equation $p = 0$ is a logical consequence of given equations $\{p_1 = 0, \dots, p_n = 0\}$. This is equivalent to check (in free f-K-vector spaces) if the vector p belongs to the subspace $\langle p_1, \dots, p_n \rangle_f$ obtained from the closure under the endomorphism f of the subspace $\langle p_1, \dots, p_n \rangle$ generated by the vectors $p_1, \dots, p_n$. In order to show the decidability of the problem, we can use standard completion methods (imitating Buchberger algorithm): we show here what to do, leaving the details to the reader.

(i) Let us call 'vector components' the terms of the form $f^m(x_i)$; vector components must be given a *terminating total order* (satisfying suitable admissibility requirements). An example of such an ordering is the following: define $f^{n_1}(x_i) \prec f^{n_2}(x_j)$ iff $\langle n_1, i \rangle <_{lex} \langle n_2, j \rangle$ (where $<_{lex}$ is the lexicographic order on $\mathbb{N} \times \mathbb{N}$). We will call head component the greatest vector component occurring in a term p (denoted by $H(p)$), head coefficient the coefficient of the head component (denoted by $HC(p)$) and remainder the term $R(p) = p - HC(p) \cdot H(p)$.

(ii) Given a finite set P of terms with head coefficient 1, the *reduction relation* $\to_P$ is introduced as follows: $p_1 \to_P p_2$ holds iff for some $p \in P$

   (a) p_1 contains a component m whose coefficient is $k \neq 0$;

   (b) there exists n such that $m = f^n(H(p))$;

   (c) $p_2 = p_1 - k \cdot f^n(p)$.

(iii) Critical pairs are identified with *S-vectors*, in the sense that $\to_P$ is confluent iff all S-vectors $S(p_l, p_r)$ (for $p_l, p_r \in P$) normalize to 0. Here if $p_l = H(p_l) + R(p_l)$ and $p_r = H(p_r) + R(p_r)$, the term $S(p_l, p_r)$ is defined as follows:

   – if there exists n such that $H(p_i) = f^n(H(p_j))$ ($i, j \in \{l, r\}$, $i \neq j$), then $S(p_l, p_r) = p_i - f^n(p_j)$;

   – $S(p_l, p_r) = 0$ otherwise.

(iv) A terminating *completion procedure* turns an arbitrary $\to_P$ into some confluent and equivalent $\to_Q$: the procedure simply adds a normal form of the S-vectors that do not normalize to 0 to the current set of terms. For efficiency, the procedure may also perform backward simplification steps.

We finally mention how to extract from the above completion procedure a *noetherian* Φ_{Kend}-p.r.e. for the subfragment Φ_K: to this aim, notice that according to the above ordering for vector components, only f-free rewrite rule can be used to reduce an f-free term. So it is possible to compute a full Φ_K-basis for the Φ_{Kend}-constraint $\Gamma = \{p_1 = 0, \dots, p_n = 0, r_1 \neq 0, \dots, r_m \neq 0\}$ by running the completion procedure on $\{p_1, \dots, p_n\}$ and by following the steps (b), (c) of Example 3.23.

4 Combined Fragments

We give now the formal definition for the operation of combining fragments.

Definition 4.1. *Let $\Phi_1 = \langle \mathcal{L}_1, T_1, \mathcal{S}_1 \rangle$ and $\Phi_2 = \langle \mathcal{L}_2, T_2, \mathcal{S}_2 \rangle$ be i.a.f.'s on the languages $\mathcal{L}_1$ and $\mathcal{L}_2$ respectively; we define the* shared fragment of Φ_1, Φ_2 *as the i.a.f. $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$, where*

- $\mathcal{L}_0 := \mathcal{L}_1 \cap \mathcal{L}_2$;
- $T_0 := T_{1|\mathcal{L}_0} \cap T_{2|\mathcal{L}_0}$;
- $\mathcal{S}_0 := \mathcal{S}_{1|\mathcal{L}_0} \cup \mathcal{S}_{2|\mathcal{L}_0}$.

Thus the Φ_0-terms are the $\mathcal{L}_0$-terms that are both Φ_1-terms and Φ_2-terms, whereas the Φ_0-structures are the $\mathcal{L}_0$-structures which are reducts either of a Φ_1- or of a Φ_2-structure. According to the above definition, Φ_0 is a subfragment of both Φ_1 and Φ_2.

Definition 4.2. *The* combined fragment *of the i.a.f.'s Φ_1 and Φ_2 is the i.a.f.*

$$\Phi_1 \oplus \Phi_2 = \langle \mathcal{L}_1 \cup \mathcal{L}_2, T_1 \oplus T_2, \mathcal{S}_1 \oplus \mathcal{S}_2 \rangle$$

in the language $\mathcal{L}_1 \cup \mathcal{L}_2$ such that:

- *$T_1 \oplus T_2$ is the smallest set of $\mathcal{L}_1 \cup \mathcal{L}_2$-terms which includes $T_1 \cup T_2$, is closed under composition and contains domain and codomain variables;*
- *$\mathcal{S}_1 \oplus \mathcal{S}_2 = \{\mathcal{A} \mid \mathcal{A}$ is a $\mathcal{L}_1 \cup \mathcal{L}_2$-structure s.t. $\mathcal{A}_{|\mathcal{L}_1} \in \mathcal{S}_1$ and $\mathcal{A}_{|\mathcal{L}_2} \in \mathcal{S}_2\}$.*

$T_1 \oplus T_2$ is defined in such a way that conditions (i)-(ii)-(iii) from Definition 3.2 are matched;[30] of course, since $\Phi_1 \oplus \Phi_2$-types turn out to be just the types which are either Φ_1- or Φ_2-types, closure under domain and codomain variables comes for free.

[30]In Subsection 4.1 we shall prove that $T_1 \oplus T_2$ is recursive (given that T_1 and T_2 are recursive).

4.1 The Purification Steps

We say that a $\Phi_1 \oplus \Phi_2$-term is *pure* iff it is a Φ_i-term ($i = 1$ or $i = 2$) and that a $\Phi_1 \oplus \Phi_2$-constraint Γ is *pure* iff it for each literal $L \in \Gamma$ there is $i = 1$ or $i = 2$ such that L is a Φ_i-literal. Constraints in combined fragments can be purified, as we shall see. Before giving the related procedure, we first have a better look to terms in a combined fragment $\Phi_1 \oplus \Phi_2 = \langle \mathcal{L}_1 \cup \mathcal{L}_2, T_1 \oplus T_2, \mathcal{S}_1 \oplus \mathcal{S}_2 \rangle$. For a $\mathcal{L}_1 \cup \mathcal{L}_2$-term t and for a natural number n, the relation $\delta(t, n)$ (written as $\delta(t) \leq n$) holds whenever one of the following non mutually exclusive conditions apply:

- $n \geq 0$ and t is a shared variable (i.e. a Φ_0-variable);
- $n \geq 1$ and $t \in T_1 \cup T_2$;
- $n \geq 2$ and there are $r, s > 0$, there are terms $u[x_1, \ldots, x_k], t_1, \ldots, t_k$ such that $n = r + s$, $\delta(u) \leq r$, $\delta(t_1) \leq s, \ldots, \delta(t_k) \leq s$ and t is equal to $u[t_1/x_1, \ldots, t_k/x_k]$.

Notice that if $\delta(t) \leq n$ holds and if $n \leq m$, then $\delta(t) \leq m$ holds too. The *degree* $\delta(t)$ of a $\mathcal{L}_1 \cup \mathcal{L}_2$-term t is the minimum d such that $\delta(t) \leq d$ holds (provided such a d exists, otherwise the degree of t is said to be infinite). It turns out that terms having degree 0 are just shared variables and terms having degree 1 are pure Φ_i-terms which are not shared variables.

Lemma 4.3. *A term $t \in \mathcal{L}_1 \cup \mathcal{L}_2$ belongs to $T_1 \oplus T_2$ iff it has a finite degree.*

Proof. As a preliminary observation (to be used below), notice that terms t satisfying $\delta(t) \leq n$ are closed under substitutions mapping variables into variables (this is easily checked by induction on $n \geq 0$).

For $n \geq 1$, the set of terms t satisfying $\delta(t) \leq n$ contains domain and codomain variables (essentially because T_1, T_2 contain their domain and codomain variables).

Let us show that terms having finite degree are closed under composition: take terms $u[x_1, \ldots, x_k]$ and $t_1, \ldots, t_k$ (all having finite degree) and suppose that types are compatible for substitution. We must show that $u[t_1/x_1, \ldots, t_k/x_k]$ has finite degree: this is obvious if u is a variable and, if all t_i are variables, we can use the above mentioned fact that terms having finite degree are closed under substitutions mapping variables into variables. Otherwise $\delta(u[x_1, \ldots, x_k]) = s_1 > 0$ and $\delta(t_1) = m_1, \ldots, \delta(t_k) = m_k$. For $s_2 := \max(m_1, \ldots, m_k) > 0$, we have that $\delta(u[t_1/x_1, \ldots, t_k/x_k]) \leq s_1 + s_2$ has finite degree too.

We proved that terms of finite degree satisfy conditions (i)-(ii)-(iii) of Definition 3.2. Viceversa, if $\delta(t) \leq n$, then it is immediate to see that t belongs to any set of $\mathcal{L}_1 \cup \mathcal{L}_2$-terms containing $T_1 \cup T_2$ and satisfying such conditions. $\square$

From Lemma 4.3 it follows that one can effectively determine whether a given term $t \in \mathcal{L}_1 \cup \mathcal{L}_2$ belongs or not to the combined fragment: it is sufficient to this aim to check whether it is a pure Φ_i-term and, in the negative case, to split it as $u[t_1, \dots, t_k]$ and to recursively check whether $u, t_1, \dots, t_k$ are in the combined fragment. This is well defined (by an induction on the size of t), because we can limit ourselves to the case in which u and at least one of the $t_1, \dots, t_k$ are not a variable:[31] in fact, for t to have degree precisely some $n > 1$, we must have $\delta(u) \leq r$, $\delta(t_1) \leq s, \dots, \delta(t_k) \leq s$ (for $r, s > 0$ such that $n = r + s$) and this implies that u and at least one of the $t_1, \dots, t_n$ are not a variable.[32]

The membership problem $t \in T_1 \oplus T_2$ however might be computationally hard: since we basically have to guess a subtree of the position tree of the term t, *the procedure we sketched is in NP*. For instance, if we combine $T_1 = \{f_1(f_0^{2n+1}(x)) \mid n \geq 0\}$ with $T_2 = \{f_0^{2n+1}(f_2(x)) \mid n \geq 0\}$,[33] then it is evident that in order to get a good splitting of $f_1(f_0^4(f_2(x)))$ one might need to backtrack a first inappropriate attempt like

$$f_1(f_0^2(y)) \quad \text{and} \quad y \mapsto f_0^2(f_2(x)).$$

Notice however that these complications in complexity (with respect to the plain Nelson-Oppen case) are due to our level of generality and that they disappear in customary situations where don't know non-determinism can be avoided by looking for 'alien' subterms, see [8] for a thorough discussion of the problem in standard first-order cases.

Let Γ be now any $\Phi_1 \oplus \Phi_2$-constraint: we shall provide finite sets Γ_1, Γ_2 of Φ_1- and Φ_2-literals, respectively, such that Γ is $\Phi_1 \oplus \Phi_2$-satisfiable iff $\Gamma_1 \cup \Gamma_2$ is $\Phi_1 \oplus \Phi_2$-satisfiable.

The purification process is obtained by iterated applications of the following:

Purification Rule

$$\frac{\Gamma', A[t, \underline{x}]}{\Gamma',\ A[y, \underline{x}],\ y = t} \tag{1}$$

where (we use notations like $\Gamma', A[t, \underline{x}]$ for the constraint $\Gamma' \cup \{A[t, \underline{x}]\}$)

– t is a non-variable term (let τ be its type);

[31] Notice that we do not require our terms to be in $\rhd_{\beta\eta}$-normal form (otherwise said, substitution is not supposed to be followed by normalization).

[32] Within the proof of Lemma 4.3, we already observed that terms having a certain finite degree are closed under variable-for-variables substitutions: hence, if all t_i are variables, from $\delta(u) \leq r$, it follows that $\delta(t) \equiv \delta(u[t_1, \dots, t_k]) \leq r < n$, contradiction.

[33] To complete the settings for this example, we may assume that $a(f_1) = S_0 \to S_1$, $a(f_2) = S_2 \to S_0$, $a(f_0) = S_0 \to S_0$ (f_0 is the unique shared symbol). Suitable variables should also be added to T_1, T_2 to formally fulfill the conditions of Definition 3.2.

– y is a variable of type τ occurring in $A[y, \underline{x}]$ but not occurring in $\Gamma', A[t, \underline{x}]$;[34]

– the literal $A[y, \underline{x}]$ is not an equation between variables;

– Γ', $A[y, \underline{x}]$, $y = t$ is a $\Phi_1 \oplus \Phi_2$-constraint (this means that it still consists of equations and inequations among $\Phi_1 \oplus \Phi_2$-terms).

The meaning of the Purification Rule is that *we are allowed to simultaneously abstract out in a constraint one or more occurrences of a non-variable subterm* t, provided we still produce a $\Phi_1 \oplus \Phi_2$-constraint (for termination, we must also take care of not introducing variable equations).[35]

Proposition 4.4. *An application of Purification Rule produces an equi-satisfiable constraint.*

Proof. The constraint $\Gamma', A[t, \underline{x}]$ is satisfied in a $\mathcal{L}_1 \cup \mathcal{L}_2$-structure $\mathcal{A} \in \mathcal{S}_1 \oplus \mathcal{S}_2$ under the assignment α iff the constraint Γ', $A[y, \underline{x}]$, $y = t$ produced by the rule is satisfied in $\mathcal{A}$ under the assignment obtained by incrementing α with $y \mapsto \mathcal{I}^{\alpha}_{\mathcal{A}}(t)$. □

The *purification process* takes as input an arbitrary $\Phi_1 \oplus \Phi_2$-constraint Γ and applies it the Purification Rule as far as possible. The Purification Rule can be applied in a don't care non deterministic way (however recall that in order to apply the rule one must before take care of the fact that the constraint produced by it still consists of $\Phi_1 \oplus \Phi_2$-literals, hence don't know non-determinism may arise inside a single application of the rule).

Proposition 4.5. *The purification process terminates and returns a set* $\Gamma_1 \cup \Gamma_2$*, where* Γ_i *is a set of* Φ_i*-literals.*

Proof. The terminating property is proved as follows. First notice that, after an application of the Purification Rule, the number N of the non variable subterm positions of the current constraint cannot increase. New equations are added by the rule, but these are only equations between a variable and a non-variable term occurring in the constraint, so that the overall number of equations that can be added during the purification process does not exceed N (notice that, after the rule has produced Γ', $A[y, \underline{x}], y = t$, the new position in which the

[34]Recall that, from our conventions in Subsection 2.3, the notation $A[y, \underline{x}]$ means that $fvar(A) \subseteq \{y, \underline{x}\}$ and $A[t, \underline{x}]$ means the formula obtained by applying to A the substitution $y \mapsto t$.

[35] The present formulation of the Purification Rule is preferable to the formulation of [30]. The latter is also slightly inaccurate, although it can be easily corrected by using a non empty *subset* of the set of the positions in the current constraint of a given nonvariable subterm. One position only is not enough, because the simultaneous abstraction of more than one occurrence of a subterm is indeed needed in the general approach of this paper (abstracting occurrences once a time may produce constraints which are not $\Phi_1 \oplus \Phi_2$-constraints, whereas simultaneous abstraction may happen to succeed).

subterm t is now is not available for another purification step, since purification steps cannot produce variables equations).

Let us now show that if the Purification Rule does not apply to Γ, then Γ splits into two pure Φ_i-constraints. First notice that any term t in a literal $t = v$ or $t \neq v$ of Γ has degree at most 1 (i.e. it is either in T_1 or in T_2): otherwise we have $t \equiv u[t_1/x_1, \dots, t_k/x_k]$, with $u[x_1, \dots, x_k], t_1, \dots, t_k$ all having lower degree than t. We then show that the Purification Rule applies, because u and at least one of the t_i are not a variable:[36] suppose for instance that t_1 is not a variable and that the constraint Γ is $\Gamma', u[t_1, \dots, t_k] = v$. The Purification Rule can now produce the constraint

$$\Gamma',\, u[x_1, t_2, \dots, t_k] = v,\, x_1 = t_1 \tag{2}$$

(x_1 can be renamed, if needed):[37] in fact fragments are closed under domain/codomain variables, hence the variable we need is at our disposal, so that (2) is a $\Phi_1 \oplus \Phi_2$-constraint (notice that $u[x_1, t_2, \dots, t_k]$ has a degree, hence it is a $\Phi_1 \oplus \Phi_2$-term).

Having established that terms in Γ are all pure, we wonder whether there are impure (in)equations. This is also impossible, because the Purification Rule can replace e.g. $t_1 = t_2$ by $t_1 = x \wedge x = t_2$ in case $t_1 \in T_1, t_2 \in T_2$ are non-variable terms (since fragments are closed under codomain variables, if $t_1 : \tau_1 \in T_1$, $t_2 : \tau_2 \in T_2$ and $\tau := \tau_1 = \tau_2$, then the type τ is shared and $t_i = x$ is a Φ_i-atom for every variable $x : \tau$). $\square$

Actually, one can prove that the Purification Rule (if exhaustively applied) can bring the current constraint into a pure and *flat* form (i.e. in a form in which negative literals just contain variables and positive literals do not contain equations among two non-variable terms).

4.2 The Combination Procedure

In this subsection, we develop a procedure which is designed to solve constraint satisfiability problems in combined fragments: the procedure is sound and we shall investigate afterwards sufficient conditions for it to be terminating and complete. Let us fix (once and for all) relevant notation for the involved data.

Assumptions/Notational Conventions.

We suppose that we are given two i.a.f.'s $\Phi_1 = \langle \mathcal{L}_1, T_1, \mathcal{S}_1 \rangle$ *and* $\Phi_2 = \langle \mathcal{L}_2, T_2, \mathcal{S}_2 \rangle$*, with shared fragment* $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$*. We suppose also that a redundancy notion* Red_{Φ_0} *for* Φ_0 *and*

[36] See the argument around footnote 32.

[37] This might be needed in order to fulfill the Purification Rule requirement that x_1 does not occur in Γ.

two non-redundant Φ_i-p.r.e.'s for Φ_0 (call them Res_{Φ_1}, Res_{Φ_2}) are available.[38] *We also fix a purified $\Phi_1 \oplus \Phi_2$-constraint $\Gamma_1 \cup \Gamma_2$ to be tested for $\Phi_1 \oplus \Phi_2$-consistency; we can freely suppose that Γ_1 and Γ_2 contain the same subset Γ_0 of Φ_0-literals (i.e. that $\Gamma_0 := \Gamma_{1|\Phi_0} = \Gamma_{2|\Phi_0}$).*[39] *We indicate by $\underline{x}_i$ the free variables occurring in Γ_i ($i = 1, 2$); $\underline{x}_0$ are those variables among $\underline{x}_1 \cup \underline{x}_2$ which happen to be Φ_0-variables (again we can freely suppose that $\underline{x}_0 = \underline{x}_1 \cap \underline{x}_2$).*[40]

In order to describe the procedure we also need a selection function in the sense of the following definition:

Definition 4.6. *A* selection function $\textsc{Choose}(\Lambda)$ *is a recursive function accepting as input a set Λ of $\Phi_0(\underline{x}_0)$-atoms and returning a positive $\Phi_0(\underline{x}_0)$-clause C such that:*

(i) *C is a Φ_i-consequence of $\Gamma_i \cup \Lambda$, for $i = 1$ or $i = 2$;*

(ii) *if $\bot$ is Φ_0-redundant wrt $\Gamma_0 \cup \Lambda$, then C is $\bot$;*

(iii) *if C is Φ_0-redundant wrt $\Gamma_0 \cup \Lambda$, then C is $\top$ or $\bot$.*

The recursive function $\textsc{Choose}(\Lambda)$ will be subject also to a fairness requirement that will be explained below.

The Procedure FComb.

Our combined procedure generates a tree whose internal nodes are labeled by sets of $\Phi_0(\underline{x}_0)$-atoms; leaves are labeled by "unsatisfiable" or by "saturated". The root of the tree is labeled by the empty set and if a node is labeled by the set Λ, then the successors are:

– *a single leaf labeled "unsatisfiable", if $\textsc{Choose}(\Lambda)$ is equal to $\bot$;*

– *or a single leaf labeled "saturated", if $\textsc{Choose}(\Lambda)$ is equal to $\top$;*

– *or nodes labeled by $\Lambda \cup \{A_1\}, \dots, \Lambda \cup \{A_k\}$, if $\textsc{Choose}(\Lambda)$ is $A_1 \vee \cdots \vee A_k$.*

The branches which are infinite or end with the *"saturated"* message are called *open*, whereas the branches ending with the *"unsatisfiable"* message are called *closed*. The procedures stops (and the generation of the above tree is interrupted) iff all branches are closed or if there is an open finite branch (of course termination is not guaranteed in the general case).[41]

[38] Of course, Res_{Φ_1} and Res_{Φ_2} are assumed to be both non-redundant with respect to Red_{Φ_0}.

[39] Otherwise, Φ_0-literals can be added to Γ_1 and Γ_2 till this holds.

[40] Otherwise, equations like $x = x$ can be added to the Γ_i.

[41] During our combination procedure Φ_0-residues are exchanged, till a saturation state is reached or till an inconsistency is detected. One may worry about the fact that information concerning $\mathcal{L}_0$-terms which are not

Algorithm 1 The combination procedure

procedure FComb(Λ)
    $C \leftarrow$ Choose(Λ)
    if $C = \bot$ **then**
        return *"unsatisfiable"*
    else if $C = \top$ **then**
        return *"saturated"*
    end if
    for all $A \in C$ **do**
        if FComb($\Lambda \cup \{A\}$) = *"saturated"* **then**
            return *"saturated"*
        end if
    end for
    return *"unsatisfiable"*
end procedure

Fair Selection Functions.

The function Choose(Λ) is *fair* iff the following happens for every open branch $\Lambda_0 \subseteq \Lambda_1 \subseteq \cdots$: if $C \equiv Res_{\Phi_i}(\Gamma_i \cup \Lambda_k, l)$ for some $i = 1, 2$ and for some $k, l \geq 0$, then C is Φ_0-redundant with respect to $\Gamma_0 \cup \Lambda_n$ for some n (roughly, *residues wrt Φ_i of an open branch are redundant with respect to the atoms in the branch*).

We first show how to build fair selection functions (under the current assumptions/notational conventions):

Proposition 4.7. *There always exists a fair selection function.*

Proof. For a finite set Λ and for a list Θ of $\Phi_0(\underline{x}_0)$-clauses, let us define the auxiliary procedure LMin(Λ, Θ). We have that

- LMin($\Lambda, [C]$) $= C$;
- LMin($\Lambda, [C|\Theta]$) $=$ LMin(Λ, Θ), if $Red_{\Phi_0}(\Lambda \cup \{D\}, C)$ holds for some $D \in \Theta$;

Φ_0-terms is not exchanged (there might *in principle* be such terms, according to Definition 4.1). However, we just pointed out that, without additional conditions, the procedure is not complete and, when sufficient conditions for completeness are introduced (see Proposition 5.5), these will be strong enough to guarantee that the exchanged information is sufficient: more specifically, the set of Φ_0-terms will be big enough for Φ_0-*equivalence* of structures to be converted into $\mathcal{L}_0$-*isomorphism*, through Φ_i-equivalence preserving semantic operations.

– $\textsc{LMin}(\Lambda, [C|\Theta]) = C$, otherwise

(roughly, the procedure takes the leftmost Red_{Φ_0}-maximal element of the list Θ, using the set Λ as a parameter).

Fix now a surjective recursive function

$$\delta = (\delta_1, \delta_2, \delta_3) : \mathbb{N} \longrightarrow \{1, 2\} \times \mathbb{N} \times \mathbb{N}$$

such that $n \geq \delta_2(n)$ holds for every n (this function can be easily built by using a recursive encoding of pairs, see e.g. [49]).

For $\Lambda = \{A_1, \dots, A_n\}$, define now $\textsc{Choose}(\Lambda)$ to be

$$\textsc{LMin}(\Gamma_0 \cup \Lambda, [Res_{\Phi_{\delta_1(n)}}(\Gamma_{\delta_1(n)} \cup \{A_1, \dots, A_{\delta_2(n)}\}, \delta_3(n)), Res_{\Phi_1}(\Gamma_1 \cup \Lambda, 0), Res_{\Phi_2}(\Gamma_2 \cup \Lambda, 0)]). \tag{3}$$

The intuitive explanation of this definition is as follows: for $i = 1, 2$ and $j \leq n$, residues $Res_{\Phi_i}(\Gamma_i \cup \{A_1, \dots, A_j\}, k)$ can be disposed in two matrices having n infinite rows ($Res_{\Phi_i}(\Gamma_i \cup \{A_1, \dots, A_j\}, k)$ is the k-th entry of the j-th row in the i-th matrix). Now our selection function explores the rows of these two matrices by a diagonal path, but before making the final choice it checks whether in the first entries of the two last rows there is anything more informative.

Clearly $\textsc{Choose}(\Lambda)$ is a Φ_i-consequence of $\Gamma_i \cup \Lambda$, for $i = 1$ or $i = 2$, because of the soundness condition of Definition 3.16; moreover if $\textsc{Choose}(\Lambda)$ is redundant wrt $\Gamma_0 \cup \Lambda$, then (by Definition 3.15 (iii)) it must be equal to $Res_{\Phi_2}(\Gamma_2 \cup \Lambda, 0)$, hence it is $\top$ or $\bot$, according to Definition 3.17 (i) (and it is $\bot$, if $\bot$ is redundant wrt $\Gamma_0 \cup \Lambda$, by Definition 3.17 (ii)).

To show fairness, pick a consistent branch labeled by the increasing sets of Φ_0-atoms $\Lambda_0 \subseteq \Lambda_1 \subseteq \cdots$ and suppose that $C \equiv Res_{\Phi_i}(\Gamma_i \cup \Lambda_k, l)$ for some $i = 1, 2$ and for some $k, l \geq 0$. Let us distinguish the case in which the consistent branch is finite and the case in which it is infinite.

If it is finite, it ends with a saturation message, which means that for some n, we have $\textsc{Choose}(\Lambda_n) \equiv \top$. From (3) and Definition 3.15 (ii)-(iii)-(iv), we must have that $Res_{\Phi_1}(\Gamma_1 \cup \Lambda_n, 0)$ and $Res_{\Phi_2}(\Gamma_2 \cup \Lambda_n, 0)$ are both equal to $\top$. To show this, notice that: (a) a residue equal to $\top$ selected by the function $\textsc{Choose}(\Lambda_n)$ according to (3) cannot be either $Res_{\Phi_{\delta_1(n)}}(\Gamma_{\delta_1(n)} \cup \{A_1, \dots, A_{\delta_2(n)}\}, \delta_3(n))$ or $Res_{\Phi_1}(\Gamma_1 \cup \Lambda_n, 0)$, because $\top$ is always redundant; (b) hence it must be $Res_{\Phi_2}(\Gamma_2 \cup \Lambda_n, 0)$, which implies however (by the way the procedure $\textsc{LMin}$ is defined) that $Res_{\Phi_1}(\Gamma_1 \cup \Lambda_n, 0)$ is redundant wrt $\Gamma_0 \cup \Lambda_n \cup \{Res_{\Phi_2}(\Gamma_2 \cup \Lambda_n, 0)\}$ (which is equal to $\Gamma_0 \cup \Lambda_n \cup \{\top\}$) and hence wrt $\Gamma_0 \cup \Lambda_n$ by transitivity. The latter implies that $Res_{\Phi_1}(\Gamma_1 \cup \Lambda_n, 0)$ is also equal to $\top$ by Definition 3.17 (i).[42]

[42] It cannot be equal to $\bot$, because it is redundant wrt $\Gamma_0 \cup \Lambda_n$: in that case, $Res_{\Phi_2}(\Gamma_2 \cup \Lambda_n, 0)$ would be $\bot$ too by Definition 3.17(ii).

Since $Res_{\Phi_1}(\Gamma_1 \cup \Lambda_n, 0)$ and $Res_{\Phi_2}(\Gamma_2 \cup \Lambda_n, 0)$ are both equal to $\top$, by Definition 3.17 (iii), we conclude that $\Gamma_0 \cup \Lambda_n$ is a Φ_0-basis for both $\Gamma_1 \cup \Lambda_n$ and $\Gamma_2 \cup \Lambda_n$, which means that the Φ_i-residue C is redundant with respect to $\Gamma_0 \cup \Lambda_n$: in fact, since $C \equiv Res_{\Phi_i}(\Gamma_i \cup \Lambda_k, l)$, by Definition 3.16, C is a Φ_i-consequence of $\Gamma_i \cup \Lambda_k$ (for $\Lambda_k \subseteq \Lambda_n$) and hence also of $\Gamma_i \cup \Lambda_n$, thus the definition of a Φ_0-basis applies.

If the branch is infinite, for some n, we have $\delta_1(n) = i, \delta_2(n) = k, \delta_3(n) = l$. Hence, either C has been selected, or some better choice (from the redundancy point of view) has been made according to (3). Since this better choice D cannot be $\top$ or $\bot$ because the branch is infinite, some atom of D (or of C, if C has been directly selected) is in Λ_{n+1}: this means that C is redundant with respect to $\Gamma_0 \cup \Lambda_{n+1}$ because of Definition 3.15(iii)-(iv)-(v). □

We underline that the fair selection function given in (3) above can be optimized in specific situations, where extra information on the input residue enumerators is available; however, the existence of a uniform schema for defining a fair selection function is an interesting property of our combination procedure.

4.3 Soundness

One possible exit of our procedure is when it generates a finite tree whose leaves are all labeled *"unsatisfiable"*: this is precisely the case in which the whole procedure returns *"unsatisfiable"*.

Proposition 4.8 (Soundness)**.** *If the procedure* FCOMB *returns "unsatisfiable", then the purified constraint* $\Gamma_1 \cup \Gamma_2$ *is* $\Phi_1 \oplus \Phi_2$*-unsatisfiable.*

Proof. We consider the tree generated by the execution of the procedure described in section 4.2. The thesis consists of proving that, if such a tree is closed, then the purified constraint $\Gamma_1 \cup \Gamma_2$ is $\Phi_1 \oplus \Phi_2$-unsatisfiable. The proof applies an inductive argument on the tree.

Consider a node labeled with Λ which is the parent of a leaf labeled with *"unsatisfiable"*: by construction CHOOSE(Λ) is equal to $\bot$. By Definition 4.6, there exists an i such that $\Gamma_i \cup \Lambda$ is Φ_i-unsatisfiable. Recalling the definition of combined fragment, it follows that $\Gamma_i \cup \Lambda$ is $\Phi_1 \oplus \Phi_2$-unsatisfiable, thus $\Gamma_1 \cup \Gamma_2 \cup \Lambda$ is unsatisfiable too.

Consider now a tree whose leaves are labeled with *"unsatisfiable"* and whose root is labeled by Λ. Suppose now, by inductive hypothesis, that each child of the root (labeled by $\Lambda \cup \{A_j\}$) is such that $\Gamma_1 \cup \Gamma_2 \cup \Lambda \cup \{A_j\}$ is $\Phi_1 \oplus \Phi_2$-unsatisfiable ($j \in \{1, \dots, k\}$). $\Gamma_1 \cup \Gamma_2 \cup \Lambda \cup \{A_j\}$ is $\Phi_1 \oplus \Phi_2$-unsatisfiable for each j iff $\Gamma_1 \cup \Gamma_2 \cup \Lambda \cup \{A_1 \vee \dots \vee A_k\}$ is $\Phi_1 \oplus \Phi_2$-unsatisfiable: this means that our inductive hypothesis entails the $\Phi_1 \oplus \Phi_2$-unsatisfiability of $\Gamma_1 \cup \Gamma_2 \cup \Lambda \cup \{A_1 \vee \dots \vee A_k\}$. By construction, our internal nodes are labeled by $\Lambda \cup \{A_1\}, \dots, \Lambda \cup \{A_k\}$ iff CHOOSE(Λ) is $A_1 \vee \dots \vee A_k$; hence, by Definition 4.6, there exists an $i \in \{1, 2\}$ such that $\Gamma_i \cup \Lambda \models_{\Phi_i} A_1 \vee \dots \vee A_k$, thus being $A_1 \vee \dots \vee A_k$ a $\Phi_1 \oplus \Phi_2$-consequence of $\Gamma_1 \cup \Gamma_2 \cup \Lambda$.

This means that for each $\Phi_1 \oplus \Phi_2$-structure in which $\Gamma_1 \cup \Gamma_2 \cup \Lambda$ is true wrt an assigment, $A_1 \vee \cdots \vee A_k$ is true wrt to the same assignment too. Moreover, the $\Phi_1 \oplus \Phi_2$-unsatisfiability of $\Gamma_1 \cup \Gamma_2 \cup \Lambda \cup \{A_1 \vee \cdots \vee A_k\}$ means that there does not exist a $\Phi_1 \oplus \Phi_2$-structure in which $\Gamma_1 \cup \Gamma_2 \cup \Lambda$ and $\{A_1 \vee \cdots \vee A_k\}$ are both true wrt at least an assignment. It follows that $\Gamma_1 \cup \Gamma_2 \cup \Lambda$ is $\Phi_1 \oplus \Phi_2$-unsatisfiable itself.

The thesis follows by the consideration that, when we run the procedure for the pure constraint $\Gamma_1 \cup \Gamma_2$, the root of the tree is labeled by the empty set by construction. $\square$

4.4 Termination

Next, we identify a relevant termination case:

Proposition 4.9 (Termination). *If Φ_0 is noetherian and Red_{Φ_0} is the full redundancy notion, then the procedure* FCOMB *terminates on the purified constraint* $\Gamma_1 \cup \Gamma_2$.

Proof. Let us consider the tree T generated by the execution of the procedure FCOMB as described in Section 4.2. Recalling that T is finite iff FCOMB($\emptyset$) terminates, we now suppose that FCOMB($\emptyset$) does not terminate. In this way T, which is a finitely branching tree by construction, is not finite and it has an infinite branch by König lemma.

This means that there is a infinite chain of sets of $\Phi_0(\underline{x}_0)$-atoms $\Lambda_1 \subset \Lambda_2 \subset \cdots \subset \Lambda_n \subset \cdots$, where Λ_i is the label of a node that belongs to that infinite path, $\Lambda_{i+1} = \Lambda_i \cup \{A_i\}$ and $Red_{\Phi_0}(\Gamma_0 \cup \Lambda_i, A_i)$ does not hold by Definitions 3.15(v) and 4.6. Since Red_{Φ_0} is the full redundancy notion, we obtained an infinite sequence $A_1, A_2, A_3, \ldots$ such that $\Gamma_0 \cup \{A_j \mid j < i\} \not\models_{\Phi_0} A_i$ for every i. This contradicts our hypothesis on the noetherianity of Φ_0. $\square$

4.5 Towards completeness

Completeness of the procedure FCOMB cannot be achieved easily, heavy conditions are needed. In this section, we nevertheless identify what is the 'semantic meaning' of a run of the procedure that either does not terminate or terminates with a saturation message.

Since our investigations are taking a completeness-oriented route, it is quite obvious that we must consider from now on only the case in which the input Φ_i-p.r.e.'s are *complete* (see Definition 3.18). In addition we need a compactness-like assumption. We say that an i.a.f. Φ is Φ_0-*compact* (where Φ_0 is a subfragment of Φ) iff, given a Φ-constraint Γ and a generalized Φ_0-constraint Γ_0, we have that $\Gamma \cup \Gamma_0$ is Φ-satisfiable if and only if for all finite $\Delta_0 \subseteq \Gamma_0$, we have that $\Gamma \cup \Delta_0$ is Φ-satisfiable.

Proposition 4.10. *Any extension Φ of a locally finite fragment Φ_0 is Φ_0-compact.*

Proof. This is due to the existence of only finitely many $\underline{x}$-representative terms in Φ_0: for this reason, a generalized $\Phi_0(\underline{x})$-constraint Γ_0 is equivalent to the constraint in which all terms have been replaced by their representatives. $\square$

The above Proposition means that, if we assume effective local finiteness in order to guarantee termination, Φ_0-compactness is guaranteed too. Notice that only special kinds of generalized Φ-constraints are involved in the definition of Φ_0-compactness, namely those that contain finitely many proper Φ-literals; thus, Φ_0-compactness is a rather weak condition (that's why it may hold for any extension whatsoever of a given fragment, as shown by Proposition 4.10). Finally, it goes without saying that, by the compactness theorem for first order logic, Φ_0-compactness is guaranteed whenever Φ is a first-order fragment.[43]

Proposition 4.11. *Suppose that Φ_1, Φ_2 are both Φ_0-compact, that the function* CHOOSE*(Λ) is fair wrt two complete Φ_i-p.r.e.'s and that the procedure* FCOMB *does not return "unsatisfiable" on the purified constraint $\Gamma_1 \cup \Gamma_2$. Then there are $\mathcal{L}_i$-structures $\mathcal{M}_i \in \mathcal{S}_i$ and $\mathcal{L}_i$-assignments α_i ($i = 1, 2$) such that:*

(i) $\mathcal{M}_1 \models_{\alpha_1} \Gamma_1$ *and* $\mathcal{M}_2 \models_{\alpha_2} \Gamma_2$*;*

(ii) *for every $\Phi_0(\underline{x}_0)$-atom A, we have that $\mathcal{M}_1 \models_{\alpha_1} A$ iff $\mathcal{M}_2 \models_{\alpha_2} A$.*

Proof. A set of positive $\Phi_0(\underline{x}_0)$-clauses $\Gamma_0^\star$ is *saturated* if and only if it is closed under the two rules:

$$\begin{array}{lcl} \Gamma_1 \cup \Gamma_0^\star \models_{\Phi_1} C & \Rightarrow & C \in \Gamma_0^\star \\ \Gamma_2 \cup \Gamma_0^\star \models_{\Phi_2} C & \Rightarrow & C \in \Gamma_0^\star, \end{array}$$

for every positive $\Phi_0(\underline{x}_0)$-clause C.

Let us suppose that Res_{Φ_1} and Res_{Φ_2} are complete p.r.e.'s. Consider now the tree T generated by an execution as described in Section 4.2, in case FCOMB does not return *"unsatisfiable"* on the purified constraint $\Gamma_1 \cup \Gamma_2$. If T is finite, then it has a branch whose leaf is labeled by *"saturated"*, otherwise it has an infinite branch (we recall that, by construction, T is a finitely branching tree and, by König lemma, a finitely branching tree which is infinite has an infinite branch).

Let us consider that (finite or infinite) open branch labeled $\Lambda_0 \subseteq \Lambda_1 \subseteq \cdots$ and let us take $\Lambda := \bigcup_j \Lambda_j$; we define $\Gamma_0^\star := \{C \mid C$ is a positive $\Phi_0(\underline{x}_0)$-clause s.t. $\Gamma_1 \cup \Lambda \models_{\Phi_1} C\}$ (we remark that $\Lambda \subseteq \Gamma_0^\star$). $\Gamma_0^\star$ is saturated and, for $i = 1, 2$, the generalized constraints $\Gamma_i \cup \Gamma_0^\star$

[43] An i.a.f. $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ is said to be *first-order* iff $\mathcal{L}$ is a first-order signature, all Φ-atoms are equivalent to first-order formulae and $\mathcal{S}$ is an *elementary class* (i.e. it is the class of the models of a first-order theory).

are Φ_i-satisfiable, as shown by Lemma 4.12. Thus, Lemma 4.13 applies and there are two $\mathcal{L}_i$-structures $\mathcal{M}_i \in \mathcal{S}_i$ satisfying $\Gamma_i \cup \Gamma_0^\star$ under assignments α_i, such that $\mathcal{M}_1, \alpha_1$ and $\mathcal{M}_2, \alpha_2$ satisfy the same $\Phi_0(\underline{x}_0)$-atoms. $\square$

Lemma 4.12. *The set $\Gamma_0^\star$ defined above is saturated, $\Gamma_1 \cup \Gamma_0^\star$ is Φ_1-satisfiable and $\Gamma_2 \cup \Gamma_0^\star$ is Φ_2-satisfiable.*

Proof. To prove that $\Gamma_0^\star$ is saturated, we need to show that

$$\Gamma_2 \cup \Gamma_0^\star \models_{\Phi_2} C \qquad \Rightarrow \qquad C \in \Gamma_0^\star$$

where C is a positive $\Phi_0(\underline{x}_0)$-clause. We prove that $\Gamma_1 \cup \Lambda \models_{\Phi_1} C$ implies $\Gamma_2 \cup \Lambda \models_{\Phi_2} C$ (and conversely, but the proof of the converse is the same).

$\Gamma_1 \cup \Lambda \models_{\Phi_1} C$ iff there exists n such that $\Gamma_1 \cup \Lambda_n \models_{\Phi_1} C$ (by Φ_0-compactness of Φ_1) iff there exists k such that[44] $Red_{\Phi_0}(\Gamma_0 \cup \Lambda_n \cup \{C_0, \dots, C_k\}, C)$ holds, where $C_j \equiv Res_{\Phi_1}(\Gamma_1 \cup \Lambda_n, j)$ (by the completeness of the Φ_1-p.r.e.). By the fairness requirement on the CHOOSE function, there exist m_j's such that $Red_{\Phi_0}(\Gamma_0 \cup \Lambda_{m_j}, C_j)$ holds ($j \in \{1, \dots, k\}$), hence by monotonicity of redundancy there exists $m \geq n, m \geq m_j$ such that $Red_{\Phi_0}(\Gamma_0 \cup \Lambda_m, C_j)$ holds for each $j \in \{1, \dots, k\}$; by transitivity of redundancy we have $Red_{\Phi_0}(\Gamma_0 \cup \Lambda_m, C)$ and consequently also $\Gamma_0 \cup \Lambda_m \models_{\Phi_0} C$. Thus $\Gamma_2 \cup \Lambda_m \models_{\Phi_2} C$ and finally $\Gamma_2 \cup \Lambda \models_{\Phi_2} C$.

We showed that $\Gamma_1 \cup \Lambda \models_{\Phi_1} C$ holds iff $\Gamma_2 \cup \Lambda \models_{\Phi_2} C$ holds; it follows that:

$$\begin{aligned} \Gamma_0^\star &= \{C \text{ is a positive } \Phi_0(\underline{x}_0)\text{-clause} \mid \Gamma_1 \cup \Lambda \models_{\Phi_1} C\} = \\ &= \{C \text{ is a positive } \Phi_0(\underline{x}_0)\text{-clause} \mid \Gamma_2 \cup \Lambda \models_{\Phi_2} C\}, \end{aligned}$$

that is $\Gamma_2 \cup \Gamma_0^\star \models_{\Phi_2} C \Rightarrow C \in \Gamma_0^\star$.

We finally prove that $\Gamma_i \cup \Gamma_0^\star$ is Φ_i-satisfiable for $i = 1, 2$. To this aim notice that $\Gamma_i \cup \Lambda$ is Φ_i-satisfiable iff $\Gamma_i \cup \Delta$ is Φ_i-satisfiable for each $\Delta \subseteq \Lambda$, Δ finite (by Φ_0-compactness of Φ_i). For each such Δ, there exists an index n such that $\Delta \subseteq \Lambda_n$. $\Gamma_i \cup \Lambda_n$ is Φ_i-satisfiable for each n (we have $\Gamma_i \cup \Lambda_n \not\models_{\Phi_i} \bot$ by completeness of Res_{Φ_i}, by definition 3.17 (ii) and by the fairness requirement of the selection function CHOOSE),[45] thus every $\Gamma_i \cup \Delta$ is Φ_i-satisfiable and consequently so is $\Gamma_i \cup \Lambda$. Since we noticed above that $\Gamma_0^\star$ consists of the clauses C such that $\Gamma_i \cup \Lambda \models_{\Phi_1} C$, it follows that $\Gamma_i \cup \Gamma_0^\star$ is Φ_i-satisfiable. $\square$

[44] Recall that $\Gamma_0 = \Gamma_{1|\Phi_0} = \Gamma_{2|\Phi_0}$ according to the 'Notational Conventions' at the beginning of Subsection 4.2.

[45] In more detail, if $\Gamma_i \cup \Lambda_n \models_{\Phi_i} \bot$ for some n, then $\bot$ appears as a residue of $\Gamma_i \cup \Lambda_n$ (by completeness of Res_{Φ_i} and by Definition 3.17 (ii)). By the fairness of the selection function it is then redundant with respect to some $\Gamma_0 \cup \Lambda_m$, which implies that CHOOSE(Λ_m) is $\bot$, by Definition 4.6, contrary to fact that the branch is not closed.

Lemma 4.13. *Suppose that we are given a saturated set of positive $\Phi_0(\underline{x}_0)$-clauses $\Gamma_0^\star$, such that $\Gamma_1 \cup \Gamma_0^\star$ is Φ_1-satisfiable and $\Gamma_2 \cup \Gamma_0^\star$ is Φ_2-satisfiable. Then there are structures $\mathcal{M}_1 \in \mathcal{S}_1$, $\mathcal{M}_2 \in \mathcal{S}_2$ and two assignments α_1, α_2 such that $\mathcal{M}_1 \models_{\alpha_1} \Gamma_1 \cup \Gamma_0^\star$ and $\mathcal{M}_2 \models_{\alpha_2} \Gamma_2 \cup \Gamma_0^\star$. Moreover, for every $\Phi_0(\underline{x}_0)$-atom A, $\mathcal{M}_1 \models_{\alpha_1} A$ holds if and only if $\mathcal{M}_2 \models_{\alpha_2} A$ holds.*

Proof. A set Δ of $\Phi_0(\underline{x}_0)$-literals is *exhaustive* iff for each $\Phi_0(\underline{x}_0)$-atom A, $A \in \Delta$ or $\neg A \in \Delta$.

Let us consider any terminating strict total order on $\Phi_0(\underline{x}_0)$-atoms (it exists by the well ordering principle) and let us extend it to a terminating strict total order on multisets.[46] We use such an ordering to define increasing subsets Δ_C, varying C among positive $\Phi_0(\underline{x}_0)$-clauses in $\Gamma_0^\star$ (positive clauses are identified here with multisets of atoms).

We say that the $\Phi_0(\underline{x}_0)$-clause $C \equiv A \vee A_1 \vee \cdots \vee A_n$ from $\Gamma_0^\star$ is *productive* (and produces the $\Phi_0(\underline{x}_0)$-atom A) iff $\{A\} > \{A_1, \ldots, A_n\}$ and $A_1, \ldots, A_n \not\in \Delta^+_{<C}$ where $\Delta^+_{<C} := \bigcup_{D<C, D\in\Gamma_0^\star} \Delta^+_D$. If C is productive and produces A, then $\Delta^+_C := \Delta^+_{<C} \cup \{A\}$, otherwise $\Delta^+_C := \Delta^+_{<C}$.

Let us define $\Delta^+ := \bigcup_{C\in\Gamma_0^\star} \Delta^+_C$ and $\Delta := \Delta^+ \cup \{\neg A \mid A \text{ is a } \Phi_0(\underline{x}_0)\text{-atom and } A \not\in \Delta^+\}$. By construction, $\Delta \models_{\Phi_0} \Gamma_0^\star$ (because Δ contains a $\Phi_0(\underline{x}_0)$-atom for every $\Phi_0(\underline{x}_0)$-clause in $\Gamma_0^\star$).

We need to show that $\Gamma_1 \cup \Delta$ is Φ_1-satisfiable and $\Gamma_2 \cup \Delta$ is Φ_2-satisfiable. First of all, we *claim* that if a $\Phi_0(\underline{x}_0)$-clause $C \equiv A \vee A_1 \vee \cdots \vee A_n$ is productive and $\{A\} > \{A_1, \ldots, A_n\}$, then $A_1, \ldots, A_n \not\in \Delta^+$. To show this, recall that, by definition, $A_i \in \Delta^+$ ($i \in \{1, \ldots, n\}$) iff A_i belongs to a productive $\Phi_0(\underline{x}_0)$-clause C_i and A_i is the maximum atom in it, thus $C_i < C$ (by multisets ordering): however none of the A_i can be in $\Delta^+_{<C}$, because C is productive, thus justifying our claim.

We suppose now that $\Gamma_1 \cup \Delta$ is Φ_1-unsatisfiable. By Φ_0-compactness of the i.a.f. Φ_1, there are $\Phi_0(\underline{x}_0)$-atom $B_1, \ldots, B_m \not\in \Delta^+$ and productive $\Phi_0(\underline{x}_0)$-clauses

$$\begin{array}{rcl} C_1 & \equiv & A_1 \vee A_{11} \vee \cdots \vee A_{1k_1} \\ & & \cdots \\ C_n & \equiv & A_n \vee A_{n1} \vee \cdots \vee A_{nk_n} \end{array}$$

(with maximum $\Phi_0(\underline{x}_0)$-atoms $A_1, \ldots, A_n$ respectively) s.t. $\Gamma_1 \cup \{A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m\}$ is Φ_1-unsatisfiable. It follows that

$$\Gamma_1 \cup \{C_1, \ldots, C_n\} \cup \{\neg A_{11}, \cdots, \neg A_{nk_n}, \neg B_1, \ldots, \neg B_m\}$$

is also unsatisfiable. As $C_1, \ldots, C_n$ are positive $\Phi_0(\underline{x}_0)$-clauses in $\Gamma_0^\star$ and $\Gamma_0^\star$ is saturated, the

[46] We are using basic information on multiset orderings that can be found in textbooks like [6].

positive $\Phi_0(\underline{x}_0)$-clause

$$D \equiv \bigvee_{i,j} A_{ij} \vee B_1 \vee \cdots \vee B_m$$

is also in $\Gamma_0^\star$. By construction, some of its $\Phi_0(\underline{x}_0)$-atoms belongs to Δ^+. $A_{11}, \ldots, A_{nk_n}$ cannot be there because the $C_1, \ldots, C_n$ are productive (see the above claim), thus at least one of the B_j's is in Δ^+: contradiction.[47] The case of Γ_2 is analogous.

We finally show that, given two structures $\mathcal{M}_1 \in \mathcal{S}_1$, $\mathcal{M}_2 \in \mathcal{S}_2$ and two assignments α_1, α_2 such that $\mathcal{M}_1 \models_{\alpha_1} \Gamma_1 \cup \Delta$ and $\mathcal{M}_2 \models_{\alpha_2} \Gamma_2 \cup \Delta$, we have that $\mathcal{M}_1, \alpha_1$ and $\mathcal{M}_2, \alpha_2$ satisfy the same $\Phi_0(\underline{x}_0)$-atoms. This is clear, because Δ an exhaustive set of $\Phi_0(\underline{x}_0)$-literals. $\square$

5 Isomorphism Theorems and Completeness

Proposition 4.11 explains what is the main problem for completeness: we would like an open branch to produce Φ_i-structures ($i = 1, 2$) whose $\mathcal{L}_0$-reducts are isomorphic and we are only given Φ_i-structures whose $\mathcal{L}_0$-reducts are $\Phi_0(\underline{x}_0)$-equivalent (in the sense that they satisfy the same $\Phi_0(\underline{x}_0)$-atoms). Hence we need a powerful semantic device that is able to *transform* $\Phi_0(\underline{x}_0)$*-equivalence into* $\mathcal{L}_0$*-isomorphism*: this device will be called an isomorphism theorem. The precise formulation of what we mean by an isomorphism theorem needs some preparation. First of all, for the notion of an isomorphism theorem to be useful for us, it should apply to fragments extended with free constants.

Given an i.a.f. $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$, we denote by $\Phi(\underline{c}) = \langle \mathcal{L}(\underline{c}), T(\underline{c}), \mathcal{S}(\underline{c}) \rangle$ the following i.a.f.: (i) $\mathcal{L}(\underline{c}) := \mathcal{L} \cup \{\underline{c}\}$ is obtained by adding to $\mathcal{L}$ finitely many new constants $\underline{c}$ (the types of these new constants must be types of Φ); (ii) $T(\underline{c})$ contains the terms of the kind $t[\underline{c}/\underline{x}, \underline{y}]$ for $t[\underline{x}, \underline{y}] \in T$; (iii) $\mathcal{S}(\underline{c})$ contains precisely the $\mathcal{L}(\underline{c})$-structures whose $\mathcal{L}$-reduct is in $\mathcal{S}$. Fragments of the kind $\Phi(\underline{c})$ are called *finite expansions* of Φ.

Let $\Phi(\underline{c})$ be a finite expansion of $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ and let $\mathcal{A}, \mathcal{B}$ be $\mathcal{L}(\underline{c})$-structures. We say that $\mathcal{A}$ is $\Phi(\underline{c})$*-equivalent* to $\mathcal{B}$ (written $\mathcal{A} \equiv_{\Phi(\underline{c})} \mathcal{B}$) iff for every closed $\Phi(\underline{c})$-atom A we have that $\mathcal{A} \models A$ iff $\mathcal{B} \models A$. By contrast, we say that $\mathcal{A}$ is $\Phi(\underline{c})$*-isomorphic* to $\mathcal{B}$ (written $\mathcal{A} \simeq_{\Phi(\underline{c})} \mathcal{B}$) iff there is an $\mathcal{L}(\underline{c})$-isomorphism from $\mathcal{A}$ onto $\mathcal{B}$.

We can now specify what we mean by a structural operation on an i.a.f. $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$. We will be very liberal here and define *structural operation on* Φ_0 any family of correspondences $O = \{O^{\underline{c}_0}\}$ associating with any finite set of free constants $\underline{c}_0$ and with any $\mathcal{A} \in \mathcal{S}_0(\underline{c}_0)$ some $O^{\underline{c}_0}(\mathcal{A}) \in \mathcal{S}_0(\underline{c}_0)$ such that $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} O^{\underline{c}_0}(\mathcal{A})$. If no confusion arises, we omit the indication of $\underline{c}_0$ in the notation $O^{\underline{c}_0}(\mathcal{A})$ and write it simply as $O(\mathcal{A})$.[48]

[47] Notice that we cannot have $k_1 = \cdots = k_n = m = 0$, because $\Gamma_1 \cup \{C_1, \ldots, C_n\} \subseteq \Gamma_1 \cup \Gamma^\star$ and the latter is consistent by hypothesis.

[48] The notion of a structural operation we propose here is sufficient to state and prove our results. However

A collection $\mathcal{O}$ of structural operations on Φ_0 *admits a* Φ_0*-isomorphism theorem* if and only if, for every $\underline{c}_0$, for every $\mathcal{A}, \mathcal{B} \in \mathcal{S}_0(\underline{c}_0)$, if $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} \mathcal{B}$ then there exist $O_1, O_2 \in \mathcal{O}$ such that $O_1(\mathcal{A}) \simeq_{\Phi_0(\underline{c}_0)} O_2(\mathcal{B})$.

Example 5.1 (Ultrapowers)**.** Ultrapowers [17] are basic constructions in the model theory of first-order logic. An ultrapower $\prod_{\mathcal{U}}$ (technically, an ultrafilter $\mathcal{U}$ over a set of indices is needed to describe the operation) transforms a first-order structure $\mathcal{A}$ into a first-order structure $\prod_{\mathcal{U}} \mathcal{A}$ which is elementarily equivalent to it (meaning that $\mathcal{A}$ and $\prod_{\mathcal{U}} \mathcal{A}$ satisfy the same first-order sentences). Hence if we take a fragment $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$, where $\mathcal{S}_0$ is an elementary class and $\langle \mathcal{L}_0, T_0 \rangle$ is an algebraic fragment of the kind analyzed in Example 3.8, then $\prod_{\mathcal{U}}$ is a structural operation on Φ_0. The following deep result in classical model theory (known as the *Keisler-Shelah isomorphism theorem* [17]) gives here a Φ_0-isomorphism theorem in our sense:[49]

Theorem 5.2 (Keisler-Shelah Isomorphism Theorem)**.** *Let $\mathcal{L}$ be a first-order signature and let $\mathcal{A}, \mathcal{B}$ be $\mathcal{L}$-structures. Then $\mathcal{A}$ is elementarily equivalent to $\mathcal{B}$ iff there is an ultrafilter $\mathcal{U}$ such that the ultrapowers $\prod_{\mathcal{U}} \mathcal{A}$ and $\prod_{\mathcal{U}} \mathcal{B}$ are $\mathcal{L}$-isomorphic.*

We shall mainly be interested into operations that can be extended to a preassigned expanded fragment. Here is the related definition. Let an i.a.f. $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ extending $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$ be given; a structural operation O on Φ_0 is Φ*-extensible* if and only if for every $\underline{c}$ and every $\mathcal{A} \in \mathcal{S}(\underline{c})$ there exist $\mathcal{B} \in \mathcal{S}(\underline{c})$ such that

$$\mathcal{B}_{|\mathcal{L}_0(\underline{c}_0)} \simeq_{\Phi_0(\underline{c}_0)} O(\mathcal{A}_{|\mathcal{L}_0(\underline{c}_0)}) \quad \text{and} \quad \mathcal{B} \equiv_{\Phi(\underline{c})} \mathcal{A},$$

(where $\underline{c}_0$ denotes the set of those constants in $\underline{c}$ whose type is a Φ_0-type).

Example 5.3. Taking the reduct of a first-order structure to a smaller signature commutes with ultrapowers, hence if $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ is an extension of $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$, both $\langle \mathcal{L}, T \rangle, \langle \mathcal{L}_0, T_0 \rangle$ are fragments of the kind analyzed in Example 3.8 and $\mathcal{S}, \mathcal{S}_0$ are elementary classes, then the Φ_0-structural operation $\prod_{\mathcal{U}}$ is Φ-extensible (the structure $\mathcal{B}$ required in the definition of Φ-extensibility is again $\prod_{\mathcal{U}} \mathcal{A}$, where the ultrapower is now taken at the level of $\mathcal{L}$-structures).

all the structural operations we shall use in the paper enjoy additional properties: for instance, since they are induced by suitable endofunctors on the category of sets, they are functorial. For similar reasons, properties like $(O^{\underline{c}_0}(\mathcal{A}))_{|\mathcal{L}_0(\underline{c}'_0)} = O^{\underline{c}'_0}(\mathcal{A}_{|\mathcal{L}_0(\underline{c}'_0)})$ (for $\underline{c}'_0 \subseteq \underline{c}_0$) are always true in our examples.

[49] If $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$ is from Example 3.6-3.7 and quantifier elimination holds in $\mathcal{S}_0$, then the $\prod_{\mathcal{U}}$'s are also structural operations on Φ_0 admitting a Φ_0-isomorphism theorem (this observation will be implicitly used in the proof of Theorem 5.7 below.)

Sometimes an isomorphism theorem does not hold precisely for a fragment $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$, but for an inessential variation (called specialization) of it. A *specialization* of Φ_0 is an i.a.f. $\Phi_0^\star$ which has the same language and the same terms as Φ_0, but whose class of $\mathcal{L}_0$-structures is a smaller class $\mathcal{S}_0^\star \subseteq \mathcal{S}_0$ satisfying the following condition: for every $\underline{c}_0$ and for every $\mathcal{A} \in \mathcal{S}_0(\underline{c}_0)$, there exists $\mathcal{A}^\star \in \mathcal{S}_0^\star(\underline{c}_0)$ such that $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} \mathcal{A}^\star$.

Given an i.a.f. $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ extending Φ_0, we say that Φ is *compatible* with respect to a given specialization $\Phi_0^\star = \langle \mathcal{L}_0, T_0, \mathcal{S}_0^\star \rangle$ of Φ_0 iff $\Phi^\star = \langle \mathcal{L}, T, \mathcal{S}^\star \rangle$ is a specialization of Φ, where $\mathcal{S}^\star$ contains exactly those $\mathcal{L}$-structures from $\mathcal{S}$ whose $\mathcal{L}_0$-reduct belongs to $\mathcal{S}_0^\star$.

This Φ_0-compatibility notion is intended to recapture, in our general setting, T_0-compatibility as introduced in [28]. The latter generalizes, in its turn, the standard stable infiniteness requirement of Nelson-Oppen procedure:

Example 5.4 (Stably Infinite First-Order Theories)**.** Let $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ be an i.a.f. of the kind considered in Example 3.6 or in Example 3.7: we say that Φ is *stably infinite* iff every satisfiable Φ-constraint is satisfiable in some infinite $\mathcal{L}$-structure $\mathcal{A} \in \mathcal{S}$. To see that this is a Φ_0-compatibility requirement, consider the i.a.f. $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$ so specified: (i) $\mathcal{L}_0$ is the empty one-sorted signature; (ii) T_0 contains only the individual variables; (iii) $\mathcal{S}_0$ is the totality of $\mathcal{L}_0$-structures (i.e. the totality of sets). A specialization $\Phi_0^\star$ of Φ_0 is obtained by considering the class $\mathcal{S}_0^\star$ formed by the infinite sets. By an easy compactness argument (compactness theorem holds here because Φ is a first-order fragment), it is easily seen that Φ is stably infinite iff it is compatible with respect to the specialization $\Phi_0^\star$ of Φ_0.

During the paper, we shall see other examples of extensible structural operations and of isomorphisms theorems (we shall divide indeed our applications into three groups, depending on the particular isomorphism theorem which is involved in them).[50]

5.1 The Main Combination Result

We are now ready to formulate a sufficient condition for our combined procedure to be complete:

Proposition 5.5. *Suppose that Φ_1, Φ_2 are both Φ_0-compact and Φ_0-compatible with respect to a specialization $\Phi_0^\star$ of Φ_0; suppose also that there is a collection $\mathcal{O}$ of structural operations on $\Phi_0^\star$ which are all $\Phi_1^\star$ and $\Phi_2^\star$-extensible and admit a $\Phi_0^\star$-isomorphism theorem. In this case, if the function* Choose*(Λ) is fair wrt two complete Φ_i-p.r.e.'s and the procedure* FComb

[50] Still, there might be further examples which are not considered in this paper and which deserve further investigation: model theory of modal logic [32] seems to be interesting in this sense (we thank V. Goranko for dropping our attention to this opportunity).

does not return "unsatisfiable" on the purified constraint $\Gamma_1 \cup \Gamma_2$*, then such a constraint is* $\Phi_1 \oplus \Phi_2$*-satisfiable.*

Proof. By Proposition 4.11, there are two structures $\mathcal{M}_1$, $\mathcal{M}_2$ and two assignments α_1, α_2 such that: (i) $\mathcal{M}_1 \in \mathcal{S}_1$, $\mathcal{M}_2 \in \mathcal{S}_2$; (ii) $\mathcal{M}_1 \models_{\alpha_1} \Gamma_1$ and $\mathcal{M}_2 \models_{\alpha_2} \Gamma_2$; (iii) $\mathcal{M}_1, \alpha_1$ and $\mathcal{M}_2, \alpha_2$ satisfy the same $\Phi_0(\underline{x}_0)$-atoms. If we put variables into bijective correspondence with free constants, we may identify the pairs $(\mathcal{M}_i, \alpha_i)$ with structures in $\mathcal{S}_i(\underline{c}_i)$, for finite sets of free constants $\underline{c}_i$. Thus we can say that there are structures $\mathcal{N}_1 \in \mathcal{S}(\underline{c}_1), \mathcal{N}_2 \in \mathcal{S}(\underline{c}_2)$ satisfying $\Gamma_1[\underline{c}_1], \Gamma_2[\underline{c}_2]$, respectively, such that $\mathcal{N}_{1|\mathcal{L}_0(\underline{c}_0)} \equiv_{\Phi_0(\underline{c}_0)} \mathcal{N}_{2|\mathcal{L}_0(\underline{c}_0)}$ (where $\underline{c}_0 = \underline{c}_1 \cap \underline{c}_2$ are the free constants whose types are Φ_0-types).[51]

Now we will show that there is a $\mathcal{L}_1(\underline{c}_1) \cup \mathcal{L}_2(\underline{c}_2)$-structure $\mathcal{M}$ such that $\mathcal{M}_{|\mathcal{L}_i} \in \mathcal{S}_i$ and $\mathcal{M} \models \Gamma_i[\underline{c}_i]$ $(i = 1, 2)$. By Φ_0-compatibility with respect to $\Phi_0^\star$, we may assume that $\mathcal{N}_{1|\mathcal{L}_0(\underline{c}_0)}$ and $\mathcal{N}_{2|\mathcal{L}_0(\underline{c}_0)}$ are in a class $\mathcal{S}_0^\star$, over which the collection of structural operations $\mathcal{O}$ admits an isomorphism theorem.

Thus there are two structural operations $O_1, O_2 \in \mathcal{O}$ such that $O_1(\mathcal{N}_{1|\mathcal{L}_0(\underline{c}_0)}) \simeq_{\Phi_0(\underline{c}_0)} O_2(\mathcal{N}_{2|\mathcal{L}_0(\underline{c}_0)})$. Since O_1, O_2 are $\Phi_1^\star$- and $\Phi_2^\star$-extensible, there exist two structures $\mathcal{B}_1 \in \mathcal{S}_1^\star(\underline{c}_1)$ and $\mathcal{B}_2 \in \mathcal{S}_2^\star(\underline{c}_2)$ such that $\mathcal{B}_{1|\mathcal{L}_0(\underline{c}_0)} \simeq_{\Phi_0(\underline{c}_0)} O_1(\mathcal{N}_{1|\mathcal{L}_0(\underline{c}_0)})$ and $\mathcal{B}_{2|\mathcal{L}_0(\underline{c}_0)} \simeq_{\Phi_0(\underline{c}_0)} O_2(\mathcal{N}_{2|\mathcal{L}_0(\underline{c}_0)})$, $\mathcal{B}_1 \equiv_{\Phi_1(\underline{c}_1)} \mathcal{N}_1$ and $\mathcal{B}_2 \equiv_{\Phi_2(\underline{c}_2)} \mathcal{N}_2$. Thus, $\mathcal{B}_1$ satisfies $\Gamma_1[\underline{c}_1]$ and $\mathcal{B}_2$ satisfies $\Gamma_2[\underline{c}_2]$. Moreover, $\mathcal{B}_{1|\mathcal{L}_0(\underline{c}_0)} \simeq_{\Phi_0(\underline{c}_0)} \mathcal{B}_{2|\mathcal{L}_0(\underline{c}_0)}$; we can now easily build the desired $\mathcal{M}$ in two steps.

In the first step, we define $\mathcal{B}'_2$ such that $\mathcal{B}_{1|\mathcal{L}_0(\underline{c}_0)} = \mathcal{B}'_{2|\mathcal{L}_0(\underline{c}_0)}$ and $\mathcal{B}_2 \simeq_{\Phi_2(\underline{c}_2)} \mathcal{B}'_2$ (notice that $\mathcal{B}'_2 \in \mathcal{S}_2(\underline{c}_2)$ by the closure under isomorphisms of $\mathcal{S}_2$, see Definition 3.3). Let ι be the isomorphism $\mathcal{B}_{1|\mathcal{L}_0(\underline{c}_0)} \longrightarrow \mathcal{B}_{2|\mathcal{L}_0(\underline{c}_0)}$; to define $\mathcal{B}'_2$, we interpret $\mathcal{L}_0$-sorts as in $\mathcal{B}_1$ and $\mathcal{L}_2 \setminus \mathcal{L}_0$-sorts as in $\mathcal{B}_2$. Put now $\iota'_S := \iota_S$ for $S \in \mathcal{L}_0$ and let ι'_S be the identity for $\mathcal{L}_2 \setminus \mathcal{L}_0$-sorts; by taking standard inductive extension to all $\mathcal{L}_2$-types, we get a family of bijections $\iota' = \{\iota'_\tau : [\![\, \tau \,]\!]_{\mathcal{B}'_2} \longrightarrow [\![\, \tau \,]\!]_{\mathcal{B}_2}\}$ (indexed by the $\mathcal{L}_2$-types) that can be used in order to complete the definition of $\mathcal{B}'_2$ (in the sense that we define the $\mathcal{B}'_2$-interpretation of every constant $d : \tau$ of $\mathcal{L}_2(\underline{c}_2)$ as $(\iota'_\tau)^{-1}(\mathcal{I}_{\mathcal{B}_2}(d))$). It is easily seen that the $\mathcal{L}_2(\underline{c}_2)$-structure $\mathcal{B}'_2$ matches the desired requirements.

Since the $\mathcal{L}_0(\underline{c}_0)$-reducts of $\mathcal{B}_1$ and $\mathcal{B}'_2$ are now just the same structure, it is easy to define (through a trivial join of both sorts and constants interpretations) a $\mathcal{L}_1(\underline{c}_1) \cup \mathcal{L}_2(\underline{c}_2)$-structure $\mathcal{M}$ such that $\mathcal{M}_{|\mathcal{L}_1(\underline{c}_1)} = \mathcal{B}_1$ and $\mathcal{M}_{|\mathcal{L}_2(\underline{c}_2)} = \mathcal{B}'_2$. Thus, the $\mathcal{L}_1 \cup \mathcal{L}_2$-reduct of $\mathcal{M}$ belongs to $\mathcal{S}_1 \oplus \mathcal{S}_2$ and satisfies $\Gamma_1[\underline{x}_1] \cup \Gamma_2[\underline{x}_2]$. $\square$

The facts we established so far can be collected into our main decidability transfer theorem:

Theorem 5.6. *Suppose that:*

[51] See the parallel convention for $\underline{x}_0$ at the beginning of Subsection 4.2. Notice that because of the Definition 4.1 concerning the shared fragment Φ_0, we have that both $\mathcal{N}_{1|\mathcal{L}_0(\underline{c}_0)}$ and $\mathcal{N}_{2|\mathcal{L}_0(\underline{c}_0)}$ belong to $\mathcal{S}_0(\underline{c}_0)$.

(1) *the interpreted algebraic fragments* Φ_1, Φ_2 *have decidable constraint satisfiability problems;*

(2) *the shared fragment* Φ_0 *is effectively locally finite (or more generally,* Φ_1, Φ_2 *are both* Φ_0*-compact,* Φ_0 *is noetherian and there exist noetherian positive residue* Φ_1*- and* Φ_2*-enumerators for* Φ_0*);*

(3) Φ_1 *and* Φ_2 *are both* Φ_0*-compatible with respect to a specialization* $\Phi_0^\star$ *of* Φ_0*;*

(4) *there is a collection* $\mathcal{O}$ *of structural operations on* $\Phi_0^\star$ *which are all* $\Phi_1^\star$*- and* $\Phi_2^\star$*-extensible and admit a* $\Phi_0^\star$*-isomorphism theorem.*

Then the procedure FComb *(together with the preprocessing Purification Rule) decides constraint satisfiability in the combined fragment* $\Phi_1 \oplus \Phi_2$.

Proof. From Propositions 4.4, 4.5, 4.7,4.8, 3.21, 4.9, 4.10, 3.20, and 5.5. □

Remark. In case the shared fragment Φ_0 is locally finite, a combination procedure can be obtained also simply by guessing a maximal set Θ_0 of $\Phi_0(\underline{x}_0)$-literals and by testing the Φ_i-satisfiability of $\Theta_0 \cup \Gamma_i$. This non-deterministic version of the procedure does not require the machinery developed in Section 3.3 (but it does not apply to noetherian cases and does not yield automatic optimizations in Φ_0-convexity cases).

Remark. Theorem 5.6 cannot be used to transfer decidability of word problems to our combined fragments: the reason is that, in case the procedure FComb is initialized with only a single negative literal, constraints containing positive literals are nevertheless generated during the execution (and also by the Purification Rule). However, since negative literals are never run-time generated, Theorem 5.6 can be used to transfer decidability of *conditional word problems*, namely of satisfiability problems for constraints containing just one negative literal.[52]

In the next three subsections we shall investigate families of concrete applications of Theorem 5.6, based on suitable isomorphism theorems.

5.2 Applications: Decidability Transfer through Ultrapowers

We shall use the isomorphism Theorem 5.2 to get the transfer decidability results of [28] as a special case of Theorem 5.6. For simplicity, we show how to do it for one-sorted functional

[52] For instance, usual relativized satisfiability problems in modal logic are the same as conditional word problems in the fragments $\Phi_M = \langle \mathcal{L}_M, T_M, \mathcal{S}_M \rangle$ from Example 3.10; here constraint satisfiability problems can be reduced to conditional word problems in case of closure under disjoint unions of the Kripke frames in Φ_M (this is a special case of the straightforward general fact that, *in convex fragments, conditional word problems and constraint satisfiability problems are inter-reducible*).

first-order signatures (and leave to the reader the easy extension to first order signatures containing also relational symbols like in Example 3.7).

Let $\Phi_1 = \langle \mathcal{L}_1, T_1, \mathcal{S}_1 \rangle$ and $\Phi_2 = \langle \mathcal{L}_2, T_2, \mathcal{S}_2 \rangle$ be equational first-order i.a.f.'s (i.e. i.a.f.'s of the kind considered in the Example 3.6) and let $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$ be their shared fragment. The hypothesis for the decidability transfer result of [28] are (equivalent to) the following:

(C1) there is a universal theory T_0 in the shared signature $\mathcal{L}_0$ such that every $\mathcal{A} \in \mathcal{S}_0$ is a model of T_0;

(C2) T_0 admits a model-completion $T_0^\star$;

(C3) for $i = 1, 2$, for every tuple $\underline{c}$ of free constants and for every $\mathcal{A} \in \mathcal{S}_i(\underline{c})$, there is some $\mathcal{A}' \in \mathcal{S}_i(\underline{c})$, s.t. $\mathcal{A}(\underline{c}) \equiv_{\Phi_i(\underline{c})} \mathcal{A}'(\underline{c})$ and $\mathcal{A}'_{|\mathcal{L}_0}$ is a model of $T_0^\star$;[53]

(C4) Φ_0 is effectively locally finite.

Condition (C2) means that: (a) $T_0 \subseteq T_0^\star$; (b) every model of T_0 embeds into a model of $T_0^\star$; (c) the following holds for every finite set of free $\mathcal{L}_0$-constants $\underline{c}_0$ and for every pair of $\mathcal{L}_0(\underline{c}_0)$-structures $\mathcal{A}, \mathcal{B}$ which are models of T_0^*: if $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} \mathcal{B}$, then $\mathcal{A}$ and $\mathcal{B}$ are $\mathcal{L}_0(\underline{c}_0)$-elementarily equivalent (see any textbook on model theory like [17] for more information).[54]

Theorem 5.7 ([28]). *Suppose that Φ_1 and Φ_2 are first-order equational i.a.f.'s satisfying conditions* (C1)-(C4) *above; if constraint satisfiability problems are decidable in Φ_1 and Φ_2, then they are decidable in $\Phi_1 \oplus \Phi_2$ too.*

Proof. We check the conditions of Theorem 5.6. We take $\mathcal{S}_0^\star$ to be the class of $\mathcal{L}_0$-structures $\mathcal{B}$ such that $\mathcal{B} \in \mathcal{S}_0$ and $\mathcal{B}$ is a model of the model-completion $T_0^\star$ of T_0 (see (C2) above). To check that $\Phi_0^\star = \langle \mathcal{L}_0, T_0, \mathcal{S}_0^\star \rangle$ is a specialization of Φ_0 argue as follows. If $\mathcal{A} \in \mathcal{S}_0(\underline{c}_0)$, then by Definition 4.1, there is $\mathcal{M} \in \mathcal{S}_i(\underline{c}_0)$ $(i = 1, 2)$ such that $\mathcal{M}_{|\mathcal{L}_0(\underline{c}_0)} = \mathcal{A}$. By (C3) above there is some $\mathcal{M}' \in \mathcal{S}_i$ which is a model of $T_0^\star$ and such that $\mathcal{M}' \equiv_{\Phi_i(\underline{c}_0)} \mathcal{M}$ holds. It follows that the $\mathcal{L}(\underline{c}_0)$-reduct of $\mathcal{M}'$ is in $\mathcal{S}_0^\star$ and it is $\Phi_0(\underline{c}_0)$-equivalent to $\mathcal{A}$. The same argument proves also the $\Phi_0^\star$-compatibility of Φ_i.

If we have two decision procedures for the constraint satisfiability problem over Φ_1 and Φ_2, then 5.6 (1) holds; (C4) guarantees 5.6 (2) and 5.6 (3) has been just proved. To check the remaining condition 5.6 (4), we use ultrapowers and the isomorphism Theorem 5.2.

[53] Instead of (C3), in [28] it is asked the (apparently stronger but in fact equivalent) condition: (C3') every $\mathcal{A} \in \mathcal{S}_i$ embeds into some $\mathcal{A}' \in \mathcal{S}_i$ which is a model of $T_0^\star$.

[54] Usually condition (c) is formulated by saying that: (c') the union of $T_0^\star$ and of the Robinson diagram of a model of T_0 is a complete first order theory. It can be shown that (c) is equivalent to (c'), but for our purpose of deriving the results from [28], it is sufficient to observe that (c') implies (c): this is clear since $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} \mathcal{B}$ means precisely that $\mathcal{A}$ and $\mathcal{B}$ are both models of the Robinson diagram of the substructure generated by the $\underline{c}_0$ (the latter is a model of T_0 because T_0 is universal).

In fact, for every $\underline{c}_0$, if $\mathcal{A}, \mathcal{B} \in \mathcal{S}_0^\star(\underline{c}_0)$ are such that $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} \mathcal{B}$, then $\mathcal{A}$ and $\mathcal{B}$ are $\mathcal{L}_0(\underline{c}_0)$-elementarily equivalent because they are models of $T_0^\star$. Hence we have $\prod_{\mathcal{U}} \mathcal{A} \simeq_{\Phi_0(c_0)} \prod_{\mathcal{U}} \mathcal{B}$, for a suitable ultrafilter $\mathcal{U}$.

Ultrapowers as structural operations are $\Phi_i^\star$-extensible, meaning that for every $\mathcal{U}$, for every $\underline{c}$ and for every $\mathcal{A} \in \mathcal{S}_i^\star(\underline{c})$, there exists $\mathcal{B} \in \mathcal{S}_i^\star(\underline{c})$ such that $\mathcal{B}_{|\mathcal{L}_0(\underline{c})} \simeq_{\Phi_0(\underline{c})} \prod_{\mathcal{U}}(\mathcal{A}_{|\mathcal{L}_0(\underline{c})})$ and $\mathcal{B} \equiv_{\Phi_i(\underline{c})} \mathcal{A}$.[55] To see this, take $\mathcal{B} := \prod_{\mathcal{U}} \mathcal{A}$; since the class of structures which are models of an elementary theory is closed under ultrapowers, $\mathcal{B} \in \mathcal{S}_i^\star(\underline{c})$. Furthermore we have

$$\mathcal{B}_{|\mathcal{L}_0(\underline{c})} = (\prod_{\mathcal{U}} \mathcal{A})_{|\mathcal{L}_0(\underline{c})} \simeq_{\Phi_0(\underline{c})} \prod_{\mathcal{U}}(\mathcal{A}_{|\mathcal{L}_0(\underline{c})}),$$

as desired. □

If we take as T_0 the empty theory (in the one-sorted first-order empty language with equality), then $T_0^\star$ is the theory of an infinite set and condition (C3) is equivalent to stable infiniteness (by a simple argument based on compactness); thus, Theorem 5.7 *reduces to the standard Nelson-Oppen result* [48], [50], [56] *concerning stably infinite theories over disjoint signatures.*

We recall from [28] that among relevant examples of theories to which Theorem 5.7 is easily seen to apply, we have Boolean algebras with operators (namely the theories axiomatizing algebraic semantics of modal logic): thus, decidability of conditional word problem transfers from two theories axiomatizing varieties of Boolean algebras with operators to their union (provided only Boolean operators are shared). This result, proved in [59] by specific techniques, is the algebraic version of the *fusion transfer of decidability of global consequence relation in modal logic.*

We remark that condition (C4) can be weakened to

(C4′) Φ_0 is noetherian and there exist noetherian positive residue Φ_1- and Φ_2-enumerators for Φ_0

(as suggested by Theorem 5.6 (2)) and we give an example of an application of Theorem 5.7 under this weaker condition.

Example 5.8 (A Combination of noetherian fragments)**.** We consider the combined fragment $\Phi_1 \oplus \Phi_2$ where Φ_1 is the fragment Φ_{Kalg} of the Example 3.22 and Φ_2 is the fragment Φ_{Kend} of Example 3.25 (here, however, we require K-algebras to be non degenerate, i.e. to satisfy the condition $0 \neq 1$). From Definition 4.2, it follows that the class $\mathcal{S}_1 \oplus \mathcal{S}_2$ consists of the models of the theory of the non degenerate K-algebras endowed with a linear endomorphism (i.e.

[55] Since, for the sake of simplicity, we limited our analysis to the one-sorted case, the set of types of Φ_0 is the same as the set of types of Φ_i, so the $\underline{c}_0$ in the definition of an extensible operation are equal to the $\underline{c}$ here.

endowed with a function f preserving sum and scalar multiplication). The set $\mathcal{S}_0$ of structures of the shared fragment Φ_0 consists of the models of the theory T_0 of K-vector spaces. The theory T_0 is universal and admits as a model completion the theory $T_0^\star = T_0 \cup \{\exists x\, (x \neq 0)\}$, if K is an infinite field, and the theory $T_0 \cup \{\exists x_1 \cdots \exists x_n \bigwedge_{i \neq j} x_i \neq x_j\}_{n \in \mathbb{N}}$, otherwise: in both cases, the models of $T_0^\star$ are just infinite K-vector spaces. Thus conditions (C1) and (C2) are satisfied. Since every non degenerate K-algebra (resp. every f-K-vector space) can be embedded into an infinite K-algebra (resp. into an infinite f-K-vector space), condition (C3) holds too. Condition (C4′) is also satisfied, as pointed out in Subsection 3.5 when discussing Examples 3.22, 3.23 and 3.25. Hence the combination procedure FComb decides conditional word problems for the theory of (non degenerate) K-algebras endowed with a linear endomorphism.

As another application of Theorem 5.6 based on Keisler-Shelah isomorphism theorem, we show *how to include a first order equational theory within description logic A-Boxes.*

An *A-Box fragment* is an i.a.f. of the kind $\Phi_{ML} = \langle \mathcal{L}_{ML}, T_{ML}, \mathcal{S}_{ML} \rangle$, where $\langle \mathcal{L}_{ML}, T_{ML} \rangle$ is defined (out of a modal signature O_M) as in Example 3.11 and $\mathcal{S}_{ML}$ is a class of $\mathcal{L}_{ML}$-structures closed under isomorphisms and disjoint I-copies. The latter operation is defined as follows:

Definition 5.9 (Disjoint I-copy). *Consider a first order one-sorted relational signature $\mathcal{L}$ and a (non empty) index set I. The operation $\sum_I$, defined on $\mathcal{L}$-structures and called* disjoint I-copy, *associates with an $\mathcal{L}$-structure $\mathcal{M} = \langle [\![-]\!]_\mathcal{M}, \mathcal{I}_\mathcal{M} \rangle$ the $\mathcal{L}$-structure $\sum_I \mathcal{M}$ such that $[\![W]\!]_{\sum_I \mathcal{M}}$ is the disjoint union of I-copies of $[\![W]\!]_\mathcal{M}$ (here W is the unique sort of $\mathcal{L}$). The interpretation of relational predicates is defined as follows*[56]

$$\sum_I \mathcal{M} \models P(\langle d_1, i_1 \rangle, \dots, \langle d_n, i_n \rangle) \quad \Longleftrightarrow \quad i_1 = i_2 = \cdots = i_n \text{ and } \mathcal{M} \models P(d_1, \dots, d_n) \tag{4}$$

for every n-ary predicate P.

Disjoint I-copy is a special case of a more general disjoint union operation: the latter is defined again by (4) and applies to any I-indexed family of structures (which may not coincide with each other). Our specific interest for disjoint I-copies is motivated by the following Lemma, concerning satisfiability of packed guarded formulae:[57]

Lemma 5.10. *Consider a first order one-sorted relational signature $\mathcal{L}$, the $\mathcal{L}$-structure $\mathcal{M}$ and its disjoint I-copy $\sum_I \mathcal{M}$. The following statements hold:*

[56] Elements of the disjoint union of I-copies of a set S are represented as pairs $\langle s, i \rangle$ (meaning that $\langle s, i \rangle$ is the i-th copy of $s \in S$).

[57] See Example 3.14 for the related definition.

(i) *for every elementary packed guarded formula* $\varphi[x_1, \dots, x_n]$ *(*$n \geq 0$*), for every* $d_1, \dots, d_n$ *in the support of* $\mathcal{M}$ *and for every index* $i \in I$*, we have that*

$$\sum_I \mathcal{M} \models \varphi[\langle d_1, i\rangle, \dots, \langle d_n, i\rangle] \quad \Longleftrightarrow \quad \mathcal{M} \models \varphi[d_1, \dots, d_n];$$

(ii) *a packed guarded elementary sentence is satisfiable in* $\mathcal{M}$ *iff it is satisfiable in* $\sum_I \mathcal{M}$*.*

Proof. We check the first claim by induction on φ (the second claim follows immediately for the case $n = 0$). If φ is atomic, just apply (4), and the case of Boolean connectives is immediate. Suppose now that φ is the packed guarded existential quantification

$$\exists y_1 \cdots \exists y_m(\pi[x_{i_1}, \dots, x_{i_k}, y_1, \dots, y_m] \wedge \psi[x_{i_1}, \dots, x_{i_k}, y_1, \dots, y_m])$$

where $x_{i_1}, \dots, x_{i_k}$ are the variables among $x_1, \dots, x_n$ that really occur free in $\varphi[x_1, \dots, x_n]$ (notice that they must all occur free in the guard π, as well as the $y_1, \dots, y_m$). That $\mathcal{M} \models \varphi[d_1, \dots, d_n]$ implies $\sum_I \mathcal{M} \models \varphi[\langle d_1, i\rangle, \dots, \langle d_n, i\rangle]$ is trivial; for the converse suppose that

$$\begin{aligned}\sum_I \mathcal{M} \models\ & \pi[\langle d_{i_1}, i\rangle, \dots, \langle d_{i_k}, i\rangle, \langle e_1, j_1\rangle, \dots, \langle e_m, j_m\rangle] \wedge \\ & \wedge \psi[\langle d_{i_1}, i\rangle, \dots, \langle d_{i_k}, i\rangle, \langle e_1, j_1\rangle, \dots, \langle e_m, j_m\rangle]\end{aligned}$$

for some $\langle e_1, j_1\rangle, \dots, \langle e_m, j_m\rangle$. By (4) and the definition of a guard, all indices $j_1, \dots, j_m$ must be equal to some j (and, if $k \neq 0$, j must be i). Thus $\mathcal{M} \models \pi[d_{i_1}, \dots, d_{i_k}, e_1, \dots, e_m] \wedge \psi[d_{i_1}, \dots, d_{i_k}, e_1, \dots, e_m]$ holds by induction hypothesis and by (4). $\square$

Let O_M be a modal signature, as defined in Example 3.10: notice that formulae like $ST(\varphi, w)$ and $\forall w ST(\varphi, w)$ are packed guarded, hence if we replace in them the second order variables of type $W \to \Omega$ by free constants for subsets of W (which are first-order relational symbols), Lemma 5.10 (i)-(ii) applies to the formulae so obtained.

We now want to combine an equational first-order i.a.f. $\Phi = \langle \mathcal{L}, T, \mathcal{S}\rangle$ from Example 3.6 and an A-Box fragment $\Phi_{ML} = \langle \mathcal{L}_{ML}, T_{ML}, \mathcal{S}_{ML}\rangle$ (we suppose that the signatures $\mathcal{L}$ and $\mathcal{L}_{ML}$ are disjoint). Assume in addition that $\mathcal{S}_{ML}$ is an elementary class (i.e it is the class of the models of a first-order $\mathcal{L}_{ML}$-theory) and that Φ is stably infinite.

Since all our data are first-order, the argument of the proof of Theorem 5.7 works, provided conditions (C1)-(C4) hold. We take as T_0 the empty theory and as $T_0^\star$ the theory of an infinite set, so that we only have to check condition (C3) for both Φ and Φ_{ML}. For the former, the condition holds trivially (the situation is the same as in the standard disjoint Nelson-Oppen case mentioned above). For the latter, for every $\mathcal{L}_{ML}(\underline{c})$-structure $\mathcal{A}(\underline{c})$ and for an infinite I, by Lemma 5.10 (i) we can expand $\sum_I \mathcal{A}$ to a $\mathcal{L}_{ML}(\underline{c})$-structure in such a way that

$\mathcal{A}(\underline{c}) \equiv_{\Phi_{ML}(\underline{c})} \sum_I \mathcal{A}(\underline{c})$ holds[58]: this proves condition (C3) (obviously, $\sum_I \mathcal{A}(\underline{c})$ is infinite, because I is infinite). We so proved the following result:

Theorem 5.11. *Suppose that we are given an equational first-order i.a.f. $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ and an A-Box fragment $\Phi_{ML} = \langle \mathcal{L}_{ML}, T_{ML}, \mathcal{S}_{ML} \rangle$; suppose also that the signatures $\mathcal{L}$ and $\mathcal{L}_{ML}$ are disjoint, that Φ is stably infinite and that $\mathcal{S}_{ML}$ is an elementary class. Then decidability of constraint satisfiability problems transfers from Φ and Φ_{ML} to $\Phi \oplus \Phi_{ML}$.*

The fragment $\Phi \oplus \Phi_{ML}$ of Theorem 5.11 is quite peculiar, because its combined terms all arise from a single composition step (they all have degree 2, in the terminology of Lemma 4.3).

5.3 Applications: Decidability Transfer through Disjoint Copies

Disjoint copies are the key tool for transfer decidability results in modal fragments too. Let O_M be a modal signature, as defined in Example 3.10. A *modal i.a.f. over O_M* is a fragment of the kind $\Phi_M = \langle \mathcal{L}_M, T_M, \mathcal{S}_M \rangle$, where $\mathcal{L}_M$ and T_M are as defined in Example 3.10, whereas $\mathcal{S}_M$ is a class of $\mathcal{L}_M$-structures *closed under isomorphisms and disjoint I-copies.* Next Proposition translates into our settings the main ingredient of the decidability transfer proof for relativized satisfiability in [10].

In the following, we indicate by O_{M_0} the empty modal signature.

Proposition 5.12. *Let Φ_M be a modal i.a.f. over the modal signature O_M and consider a modal subfragment Φ_{M_0} of it, based on the empty modal signature; the structural operations $\{\sum_I\}_I$ over Φ_{M_0} are Φ_M-extensible and form a collection admitting a Φ_{M_0}-isomorphism theorem.*

Proof. Recall from Example 3.10 that $W \to \Omega$ is the only type of the i.a.f. Φ_M, hence the relevant free constants $\underline{c}$ in expanded languages are second-order constants for subsets of W (this means, in particular, that their interpretation can be extended to disjoint I-copies like any other relational first-order predicate symbol as shown in Definition 5.9).

That taking disjoint I-copies $\sum_I$ is a structural Φ_M-extensible operation is clear: to define $\mathcal{N} \in \mathcal{S}_M(\underline{c})$ which is $\Phi_M(\underline{c})$-equivalent to some $\mathcal{M} \in \mathcal{S}_M(\underline{c})$ and whose $\Phi_{M_0}(\underline{c})$-reduct is $\mathcal{L}_{M_0}(\underline{c})$-isomorphic to $\sum_I(\mathcal{M}_{|\mathcal{L}_{M_0}(\underline{c})})$ it is sufficient to take the $\mathcal{L}_M(\underline{c})$-structure $\sum_I \mathcal{M}$ as $\mathcal{N}$ and apply Lemma 5.10 (ii) (recall that constraints in Φ_M are equivalent to conjunctions of formulae of the kind $\forall w ST(\varphi, w)$ and their negations). That taking disjoint I-copies is a structural operation (i.e. that a $\mathcal{L}_{M_0}(\underline{c})$-structure and its disjoint I-copy are $\Phi_{M_0}(\underline{c})$-equivalent) is clear by the same reasons.

[58] To this aim, interpret a free individual constant $c \in \underline{c}$ in $\sum_I \mathcal{A}$ as $\langle \mathcal{I}_{\mathcal{A}(\underline{c})}(c), i \rangle$, where i is some arbitrarily chosen element of I (to be the same for all the free individual constants c that belong to $\underline{c}$).

To show that a Φ_{M_0}-isomorphism theorem holds, suppose that we are given free constants $\underline{c}_0 := \{P_1, \dots, P_n\}$ and two structures $\mathcal{M}_1$ and $\mathcal{M}_2$ in $\mathcal{S}_{M_0}(\underline{c}_0)$ such that $\mathcal{M}_1 \equiv_{\Phi_{M_0}(\underline{c}_0)} \mathcal{M}_2$; we show that $\sum_I \mathcal{M}_1 \simeq_{\Phi_{M_0}(\underline{c}_0)} \sum_I \mathcal{M}_2$ holds for some I.

Consider every boolean combination of the form $\varepsilon(w) = Q_1(w) \wedge \cdots \wedge Q_n(w)$ where $Q_j \equiv P_j$ or $Q_j \equiv \neg P_j$ (thus the number of such formulae is 2^n). For a given $\mathcal{L}_{M_0}(\underline{c}_0)$-structure $\mathcal{N}$, let $\varepsilon(\mathcal{N}) := \{a \in [\![W]\!]_{\mathcal{N}} \mid \mathcal{N} \models \varepsilon(a)\}$ and let us associate with $\mathcal{N}$ the 2^n cardinal invariants $a_\varepsilon(\mathcal{N}) := \sharp\varepsilon(\mathcal{N})$. Now two $\mathcal{L}_{M_0}(\underline{c}_0)$-structures $\mathcal{N}_1$ and $\mathcal{N}_2$ having the same invariants are isomorphic, because we can glue bijections $\varepsilon(\mathcal{N}_1) \longrightarrow \varepsilon(\mathcal{N}_2)$ to a $\mathcal{L}_{M_0}(\underline{c}_0)$-isomorphism $\mathcal{N}_1 \simeq \mathcal{N}_2$.

Finally, we note that $\mathcal{M}_1 \equiv_{\Phi_{M_0}(\underline{c}_0)} \mathcal{M}_2$ means that $\mathcal{M}_1 \models A$ holds iff $\mathcal{M}_2 \models A$ holds for every closed $\Phi_{M_0}(\underline{c}_0)$-atom A. In particular, $\mathcal{M}_1 \models \{w \mid \varepsilon(w)\} = \{w \mid \bot\}$ iff $\mathcal{M}_2 \models \{w \mid \varepsilon(w)\} = \{w \mid \bot\}$: thus $\varepsilon(\mathcal{M}_1) = \emptyset$ iff $\varepsilon(\mathcal{M}_2) = \emptyset$ holds for all ε. Let now consider a set I whose cardinality m is such that $m \geq \varepsilon(\mathcal{M}_i)$ for all ε and for $i \in \{1, 2\}$: we show that $\sum_I \mathcal{M}_1 \simeq_{\Phi_{M_0}(\underline{c}_0)} \sum_I \mathcal{M}_2$ proving that the two structures have the same invariants. In fact the cardinal identities $a_\varepsilon(\sum_I \mathcal{M}_1) = m \cdot a_\varepsilon(\mathcal{M}_1) = m = m \cdot a_\varepsilon(\mathcal{M}_2) = a_\varepsilon(\sum_I \mathcal{M}_2)$ hold for all ε. $\square$

If O_{M_1} and O_{M_2} are modal signatures, let us $O_{M_1 \oplus M_2}$ indicate their disjoint union ($O_{M_1 \oplus M_2}$ is called the fusion of the modal signatures O_{M_1} and O_{M_2}). Given a modal i.a.f. Φ_{M_1} over O_{M_1} and a modal i.a.f. Φ_{M_2} over O_{M_2}, let us define their *fusion* as the modal i.a.f.

$$\Phi_{M_1 \oplus M_2} = \langle \mathcal{L}_{M_1 \oplus M_2}, T_{M_1 \oplus M_2}, \mathcal{S}_{M_1} \oplus \mathcal{S}_{M_2} \rangle.$$

Notice that for two modal i.a.f.'s $\Phi_{M_1} = \langle \mathcal{L}_{M_1}, T_{M_1}, \mathcal{S}_{M_1} \rangle$ and $\Phi_{M_2} = \langle \mathcal{L}_{M_1}, T_{M_1}, \mathcal{S}_{M_1} \rangle$ over disjoint modal signatures, the shared fragment $\Phi_{M_0} = \langle \mathcal{L}_{M_0}, T_{M_0}, \mathcal{S}_{M_0} \rangle$ is locally finite, because it is a modal i.a.f. over the empty modal signature (for any finite set of Φ_{M_0}-variables $\underline{x}_0$, the representative $\Phi_{M_0}(\underline{x}_0)$-terms are those of the kind $\{w \mid \psi(w)\}$, where ψ is a boolean combination of the second order variables $\underline{x}_0$).

Now if Φ_{M_1} and Φ_{M_2} have decidable constraint satisfiability problems, then so does the combined i.a.f. $\Phi_{M_1} \oplus \Phi_{M_2}$: in fact, the hypotheses of Theorem 5.6 are satisfied by the previous Proposition.[59] To infer the transfer decidability result to the fusion modal i.a.f., we need to clarify the relationship between $\Phi_{M_1 \oplus M_2}$ and $\Phi_{M_1} \oplus \Phi_{M_2}$.

Given two i.a.f.'s $\Phi_1 = \langle \mathcal{L}_1, T_1, \mathcal{S}_1 \rangle$ and $\Phi_2 = \langle \mathcal{L}_2, T_2, \mathcal{S}_2 \rangle$, we say that they are $\beta\eta$-*equivalent* (written $\Phi_1 \sim_{\beta\eta} \Phi_2$) iff $\mathcal{L}_1 = \mathcal{L}_2$, $\mathcal{S}_1 = \mathcal{S}_2$ and moreover for every $t_1 \in T_1$ one can effectively compute some $t_2 \in T_2$ such that $t_1 \sim_{\beta\eta} t_2$, and vice versa. Clearly, $\beta\eta$-equivalent i.a.f.'s can be considered to be just the same.

[59] We obviously take $\Phi_0^\star$ to be Φ_{M_0} in 5.6 (3).

Lemma 5.13. *If $\Phi_{M_1 \oplus M_2}$ and $\Phi_{M_1} \oplus \Phi_{M_2}$ are as above, we have that $\Phi_{M_1 \oplus M_2} \sim_{\beta\eta} \Phi_{M_1} \oplus \Phi_{M_2}$.*

Proof. Since $T_{M_1} \oplus T_{M_2}$ is defined to be the minimum set of terms closed under substitutions and containing T_{M_1} and T_{M_2} and since $T_{M_1 \oplus M_2}$ enjoys these properties, clearly any $t \in T_{M_1} \oplus T_{M_2}$ belongs to $T_{M_1 \oplus M_2}$.

Conversely, let us take $t \in T_{M_1 \oplus M_2}$; then $t \sim_{\beta\eta} \{w \mid ST(\varphi, w)\}$ for some $O_{M_1 \oplus M_2}$-modal formula φ.[60] By induction on φ, we define $u \in T_{M_1} \oplus T_{M_2}$ such that $u \sim_{\beta\eta} \{w \mid ST(\varphi, w)\}$ (then $t \sim_{\beta\eta} u$ follows by transitivity). If φ is a propositional variable we can take u to be $\{w \mid ST(\varphi, w)\}$. If φ is $\psi_1 \wedge \psi_2$, by induction there are $u_1, u_2 \in T_{M_1} \oplus T_{M_2}$ such that $u_i \sim_{\beta\eta} \{w \mid ST(\psi_i, w)\}$ for $i = 1, 2$. Then $\{w \mid ST(\varphi, w)\} = \{w \mid ST(\psi_1, w) \wedge ST(\psi_2, w)\} \sim_{\beta\eta} \{w \mid \{w \mid ST(\psi_1, w)\}(w) \wedge \{w \mid ST(\psi_2, w)\}(w)\} \sim_{\beta\eta} \{w \mid u_1(w) \wedge u_2(w)\}$. The latter is obtained by replacing in the term $\{w \mid ST(\mathrm{x}_1 \wedge \mathrm{x}_2, w)\} = \{w \mid X_1(w) \wedge X_2(w)\}$ the terms $u_1, u_2 \in T_{M_1} \oplus T_{M_2}$ for the second order variables X_1, X_2, respectively, hence it is a term that belongs to $T_{M_1} \oplus T_{M_2}$ too, because the latter is closed under substitution. The cases of $\vee, \neg, \Diamond_k$ are analogous. $\square$

We have so proved the following well-known decidability transfer result (see e.g. [10] and the literature quoted therein):

Theorem 5.14 (Decidability transfer for modal i.a.f.'s)**.** *If two modal i.a.f.'s Φ_{M_1} and Φ_{M_2} have decidable constraint satisfiability problems, so does their fusion $\Phi_{M_1 \oplus M_2}$.*

Fragments of the kind examined in Example 3.11 are not interesting for being combined with each other, because the absence of the type $W \to \Omega$ makes such combinations trivial. On the contrary, full modal fragments from Example 3.12 are quite interesting in this respect (we recall that they reproduce both A-Box and T-Box reasoning from the point of view of description logics). In fact very slight modifications are sufficient to get a result analogous to Theorem 5.14: we just sketch how to do it.

Let O_M be a modal signature; a *full modal i.a.f. over O_M* is a fragment of the kind $\Phi_{FM} = \langle \mathcal{L}_{FM}, T_{FM}, \mathcal{S}_{FM} \rangle$, where $\mathcal{L}_{FM}$ and T_{FM} are as defined in Example 3.12, whereas $\mathcal{S}_{FM}$ is again a class of $\mathcal{L}_M$-structures closed under isomorphisms and disjoint I-copies.

There is a little complication arising now: since W is a type of an i.a.f. like Φ_{FM}, when we expand languages with free constants, we now get (besides constants of type $W \to \Omega$) also individual constants of type W. The interpretation of these constants is not defined in disjoint I-copies, because taking disjoint I-copies is an operation defined only for first-order relational signatures. We proceed as follows: we take index sets I which are pointed, namely

[60] Notice that $\{w \mid ST(\varphi, w)\}$ - hence also φ - can be effectively computed because it is in long-$\beta\eta$-normal form and so it is the long-$\beta\eta$-normal form of t.

some $i_0 \in I$ is specified. Then, we define the interpretation of an individual constant c of type W in $\sum_I \mathcal{M}$ as $\langle \mathcal{I}_{\mathcal{M}}(c), i_0 \rangle$.

The definition of fusion for full modal i.a.f.'s is the obvious one and it leads to the following:

Theorem 5.15 (Decidability transfer for full modal i.a.f.'s [10])**.** *If two full modal i.a.f.'s have decidable constraint satisfiability problems, so does their fusion.*[61]

Proof. We sketch the little modifications required to prove Proposition 5.12 in the present context (Lemma 5.13 does not need any essential change).

Let Φ_{FM} be a full modal i.a.f. over O_M and let Φ_{FM_0} be a subfragment of it on the empty modal signature. According to the considerations in Examples 3.10-3.11, when considering languages expanded with free constants $\underline{c}$, closed $\Phi_{FM}(\underline{c})$-atoms are now of the kind $c_1 = c_2$, $R_k(c_1, c_2)$, $ST(\psi, c/w)$, and $\forall w ST(\psi, w)$ (where second order variables in ψ have been replaced by first-order unary predicate constants). From the pointed definition of disjoint I-copy given above and Lemma 5.10, it is then clear that the $\mathcal{L}_{FM}(\underline{c})$-structures $\mathcal{M}$ and $\sum_I \mathcal{M}$ still are $\Phi_{FM}(\underline{c})$-equivalent and this is all what matters in order to check that pointed disjoint I-copies are Φ_{FM}-extensible structural operations over Φ_{FM_0}.

For the Φ_{FM_0}-isomorphism theorem, we just need to add to the invariants of a $\mathcal{L}_{FM_0}(\underline{c})$-structure $\mathcal{N}$ considered in the proof of Proposition 5.12 also the indication about the truth/falsity in $\mathcal{N}$ of the ground atoms of the kind $\varepsilon(c)$ and $c_1 = c_2$, varying c, c_1, c_2 among the individual constants in $\underline{c}$. □

The decidability transfer theorem for the non-normal case of Example 3.13 (i.e. for the full strength of abstract description systems in the sense of [10]) requires a simple adaptation of Definition 5.9 and of Lemma 5.10.

We now try to extend our decidability transfer results to cover also combinations of packed guarded or of two-variable fragments. However, to get positive results, we need to keep shared signatures under control (otherwise undecidability phenomena arise). In addition, we still want to exploit the isomorphism theorem of Proposition 5.12 and for that we need the shared signature to be empty and second order variables appearing as terms in the fragments to be monadic only. The kind of combination that arise in this way is a form of fusion, that we shall call monadic fusion. We begin by identifying a class of fragments to which our techniques apply.

[61] The statement of Theorem 5.15 seems not to allow the decidability transfer of only *positive* A-Box satisfiability with respect to T-Box axioms; however this limited decidability transfer follows immediately once one realizes that the combined algorithm FComb never adds negative information to current constraints (so if non positive A-Boxes are not present from the very beginning, there won't be any call for a decision procedure involving them). See also the second Remark after Theorem 5.6 for the same observation.

Let us call $\Phi_\emptyset = \langle \mathcal{L}_\emptyset, T_\emptyset, \mathcal{S}_\emptyset \rangle$ the following i.a.f.: (i) $\mathcal{L}_\emptyset$ is the empty one-sorted first-order signature (that is, $\mathcal{L}_\emptyset$ does not contain any proper symbol, except for its unique sort which is called D); (ii) $T_\emptyset$ consists of the terms which are $\beta\eta$-equivalent to terms belonging to $T_{11}^{\mathcal{L}_\emptyset}$;[62] (iii) $\mathcal{S}_\emptyset$ contains all $\mathcal{L}_\emptyset$-structures.

Definition 5.16. *A* monadically suitable *i.a.f.* $\Phi = \langle \mathcal{L}, T, \mathcal{S} \rangle$ *is an i.a.f. such that:*

(i) $\mathcal{L}$ *is a relational one-sorted first-order signature;*

(ii) $T_{11}^{\mathcal{L}_\emptyset} \subseteq T \subseteq T_{\omega 1}^{\mathcal{L}}$;[63]

(iii) *the* $\Phi_\emptyset$*-structural operation of taking disjoint* I*-copies is* Φ*-extensible.*

We remark that, despite the fact that the definition of a monadically suitable fragment needs the present paper settings to be formulated, there is some anticipation of it in the literature on monodic fragments.[64] We give a couple of interesting examples of monadically suitable decidable fragments:

Example 5.17. *Packed guarded fragments* are i.a.f.'s of the kind $\Phi_G = \langle \mathcal{L}_G, T_G, \mathcal{S}_G \rangle$, where T_G is as defined in Example 3.14, whereas $\mathcal{S}_G$ is a class of $\mathcal{L}_G$-structures closed under isomorphisms and disjoint I-copies. To see that these are monadically suitable fragments, recall Lemma 5.10: by this Lemma, it is easy to see that for every free constants $\underline{c}$ of type $D \to \Omega$, for every $\mathcal{A} \in \mathcal{S}_G(\underline{c})$ and for every non empty set of indices I, we have that $\mathcal{A} \equiv_{\Phi_G(\underline{c})} \sum_I \mathcal{A}$. Thus taking disjoint I-copies is trivially Φ_G-extensible.

Before giving the second family of examples of monadically suitable fragments, we introduce an alternative construction for proving extensibility of the operation of taking disjoint I-copies. This construction is nicely behaved only for fragments without identity and is called I-conglomeration:[65]

Definition 5.18 (I-conglomeration)**.** *Consider a first order one-sorted relational signature* $\mathcal{L}$ *and a (non empty) index set* I*. The operation* $\sum^I$*, defined on* $\mathcal{L}$*-structures and called* I-conglomeration*, associates with an* $\mathcal{L}$*-structure* $\mathcal{M} = \langle [\![-]\!]_\mathcal{M}, \mathcal{I}_\mathcal{M} \rangle$ *the* $\mathcal{L}$*-structure* $\sum^I \mathcal{M}$ *such that* $[\![D]\!]_{\sum^I \mathcal{M}}$ *is the disjoint union of* I*-copies of* $[\![D]\!]_\mathcal{M}$ *(here* D *is the unique sort of* $\mathcal{L}$*). The interpretation of relational constants is defined in such a way that we have*

$$\sum^I \mathcal{M} \models P(\langle d_1, i_1 \rangle, \ldots, \langle d_n, i_n \rangle) \quad \Longleftrightarrow \quad \mathcal{M} \models P(d_1, \ldots, d_n)$$

[62] See Example 3.9 for this notation and for other similar notation used below.

[63] The inclusion $T \subseteq T_{\omega 1}^{\mathcal{L}}$ should be intended up to $\beta\eta$-equivalence (namely, for every $t \in T$ there is $t' \in T_{\omega 1}^{\mathcal{L}}$ such that $t \sim_{\beta\eta} t'$).

[64] See for instance statements like that of Theorem 11.21 in [24].

[65] Notice that, contrary to disjoint union, I-conglomeration cannot be defined for families of I-indexed structures different from each other.

for every n-ary relational predicate P different from equality.

Notice that I-conglomerations and disjoint I-copies *coincide* for relational first order signatures having only unary predicates. The preservation Lemma 5.10 can be reformulated as follows:

Lemma 5.19. *Consider a first order one-sorted relational signature $\mathcal{L}$ and the $\mathcal{L}$-structures $\mathcal{M}$ and $\sum^I \mathcal{M}$. The following statements hold:*

(i) *for every first-order formula $\varphi[x_1, \dots, x_n]$* not containing the equality predicate*, for every $d_1, \dots, d_n$ in the support of $\mathcal{M}$ and for every indexes $i_1, \dots, i_n \in I$, we have that*

$$\sum^{I} \mathcal{M} \models \varphi[\langle d_1, i_1\rangle, \dots, \langle d_n, i_n\rangle] \iff \mathcal{M} \models \varphi[d_1, \dots, d_n];$$

(ii) *a first order formula not containing the equality predicate is satisfiable in $\mathcal{M}$ iff it is satisfiable in $\sum^I \mathcal{M}$.*

Example 5.20. For a first-order relational one-sorted signature $\mathcal{L}_{2V}$, a *two variables i.a.f.* over $\mathcal{L}_{2V}$ is a fragment of the kind $\Phi_{2V} = \langle \mathcal{L}_{2V}, T_{2V}, \mathcal{S}_{2V}\rangle$, where: (i) T_{2V} contains the terms *without identity* which are $\beta\eta$-equivalent to terms belonging to the set $T_{NK}^{\mathcal{L}_{2V}}$ of Example 3.9 for $K = 1$ and $N = 2$; (ii) $\mathcal{S}_{2V}$ is a class of $\mathcal{L}_{2V}$-structures closed under isomorphisms and I-conglomerations.[66]

For two monadically suitable i.a.f.'s Φ_1 and Φ_2 operating on disjoint signatures, let us call the combined fragment $\Phi_1 \oplus \Phi_2$ the *monadic fusion* of Φ_1 and Φ_2. For monadic fusions we have the following:

Theorem 5.21 (Decidability transfer for monadically suitable i.a.f.'s)**.** *If two monadically suitable i.a.f.'s Φ_1, Φ_2 operating on disjoint signatures have decidable constraint satisfiability problems, so does their monadic fusion.*

Proof. Using Definition 5.16, we can say the following about the shared fragment $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0\rangle$: (i) $\mathcal{L}_0$ is the empty signature $\mathcal{L}_\emptyset$; (ii) T_0 contains $T_{11}^{\mathcal{L}_\emptyset}$ and hence it includes the terms T_{M_0} of Example 3.10 relative to the empty modal signature O_{M_0}; (iii) for every tuple of free constants $\underline{c}_0$, the closed $\Phi_0(\underline{c}_0)$-terms $t[\underline{c}_0]$, modulo $\beta\eta$-equivalence, are a subset of the terms of the kind $\{x \mid \varphi[x]\}$, where $\varphi[x]$ is a monadic formula of first order language, possibly with equality (that is, to build $\varphi[x]$, at most equality and the free constants $\underline{c}_0$ of type $D \to \Omega$ can be used); (iv) the structures in $\mathcal{S}_0(\underline{c}_0)$ are closed under disjoint I-copies and are $\Phi_0(\underline{c}_0)$-equivalent to their disjoint I-copies.

[66] These closure properties are guaranteed if $\mathcal{L}_{2V}$ is axiomatized by first-order formulae without identity (see [36] for recent interesting material, both in the decidable and in the undecidable case).

To justify (iv), argue as follows: if $\mathcal{A} \in \mathcal{S}_0(\underline{c}_0)$, then by Definition 4.1 it is the $\mathcal{L}_\emptyset(\underline{c}_0)$-reduct of a $\Phi_i(\underline{c}_0)$-structure $\mathcal{B}$ $(i = 1, 2)$; since taking disjoint I-copies of $\mathcal{L}_\emptyset(\underline{c}_0)$-structures is $\Phi_i(\underline{c}_0)$-extensible by Definition 5.16(iii), we have that for every index set I, there is a $\Phi_i(\underline{c}_0)$-structure $\mathcal{B}'$ having $\sum_I \mathcal{A}$ as a $\mathcal{L}_\emptyset(\underline{c}_0)$-reduct and such that $\mathcal{B}$ is $\Phi_i(\underline{c}_0)$-equivalent to $\mathcal{B}'$. Taking $\mathcal{L}_\emptyset(\underline{c}_0)$-reducts, it follows that $\mathcal{A} \equiv_{\Phi_0(\underline{c}_0)} \sum_I \mathcal{A}$.

Using (ii) and (iv) above, we can repeat word-by-word the proof of Proposition 5.12 and in order to apply Theorem 5.6 we only have to show that Φ_0 is effectively locally finite. Despite the fact that there are infinitely many non equivalent monadic first order sentences with equality, we shall show by using (iii)-(iv) that there are only finitely many closed $\Phi_0(\underline{c}_0)$-terms $t[\underline{c}_0]$ which are differently interpreted in at least one structure from $\mathcal{S}_0(\underline{c}_0)$ (here $\underline{c}_0 := \{P_1, \dots, P_n\}$ are free constants, which must be of type $D \to \Omega$, because this is the only type of Φ_0). Recall that $t[\underline{c}_0] \sim_{\beta\eta} \{x \mid \varphi[x]\}$, where $\varphi[x]$ is as in (iii) above.

By closure under disjoint I-copies and $\Phi_0(\underline{c}_0)$-equivalence to disjoint I-copies (see (iv)), we can limit ourselves to the consideration of at most 2^{2^n}-structures from $\mathcal{S}_0(\underline{c}_0)$: each of these structures is uniquely determined by the fact that the cardinal invariants[67] a_ε are either 0 or m in it (here m is an infinite, big enough, cardinal).

Each of these at most 2^{2^n} structures $\mathcal{A}_S$ is identified by a set S of formulae of the kind $\varepsilon(x)$, in the sense that we have $\mathcal{A}_S \models \psi_S$, where ψ_S is the one-variable monadic sentence $\bigwedge_{\varepsilon \in S} \exists x\, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x\, \varepsilon(x)$ (notice also that for $S \neq S'$, we have $\mathcal{A}_S \not\models \psi_{S'}$). Since, for every $\varepsilon \in S$, the set $\varepsilon(\mathcal{A}_S)$ is infinite, it is easily seen that quantifier elimination holds in $\mathcal{A}_S$:[68] this means in particular that, for a given monadic first-order formula with equality $\varphi[x]$, one can effectively compute $\theta_S[x]$ such that $\mathcal{A}_S \models \forall x(\varphi[x] \leftrightarrow \theta_S[x])$ holds and such that $\theta_S[x]$ does not contain quantifiers (that is, $\theta_S[x]$ is a boolean combination of the atomic formulae $P_j(x)$). Thus, in all the structures that belongs to $\mathcal{S}_0(\underline{c}_0)$, the $\Phi_0(\underline{c}_0)$-atom

$$\{x \mid \varphi[x]\} \;=\; \{x \mid \bigvee_S (\psi_S \wedge \theta_S[x])\}$$

is true, yielding the claim (because there are only finitely many possibilities for $\theta_S[x]$). $\square$

Theorem 5.21 offers various combination possibilities, however notice that: (a) the conditions for a fragment to be monadically suitable are rather strong (for instance, the two

[67] Here and below, we freely use notation from Proposition 5.12.

[68] Recall that in order to eliminate quantifiers, it is sufficient to eliminate them from primitive formulas, i.e. from formulas of the kind $\exists y\, \chi(y, \underline{x})$, where $\chi(y, \underline{x})$ is a conjunction of literals. In our case, these literals can only be $y = x_i, y \neq x_i, P_j(y), \neg P_j(y)$ (of course, literals in which y does not occur are not relevant). Since equations $y = x_i$ causes the quantifier $\exists y$ to be removed by replacement, we can assume that our χ is equivalent to a conjunction of negative literals $y \neq x_i$ and of a Boolean combination of atomic formulas of the kind $P_j(y)$. The set defined by this Boolean combination in $\mathcal{A}_S$ is either infinite or empty, so within $\mathcal{A}_S$, the formula $\exists y\, \chi(y, \underline{x})$ is equivalent either to $\perp$ or to $\top$.

variable fragment with identity is not monadically suitable); (b) the notion of monadic fusion is a restricted form of combination, because only unary second order variables are available for replacement when forming formulae of the combined fragment (thus, for instance, the monadic fusion of two two-variables fragments does not contain sentences like $\exists x \exists y(R_1(x, y) \wedge R_2(x, y))$, if R_1, R_2 do not belong to the same component signature).

5.4 Applications: Decidability Transfer for Monodic Fragments

Fragments in first-order *modal* predicate logic become undecidable quite soon: for instance, classical decidability results for the monadic or the two-variables cases do not extend to modal languages [38], [23], [37] (see also [16] for an essential account of these and related results). However there still are interesting modal predicate fragments which are decidable: one-variable fragments (corresponding to products with S5) are usually decidable [54], [24], as well as many monodic fragments. We recall that a *monodic* formula is a modal first order formula in which modal operators are applied only to subformulae containing at most one free variable; monodic fragments are intended to be classes of monodic formulae (of course the entire class of monodic formulae is too big to be decidable, being an extension of classical first order language).

Monodic fragments whose extensional (i.e. non modal) components is decidable seem to be decidable too [58],[24]: we shall give this fact a formulation in terms of a decidability transfer result for monodic fragments which are obtained as combinations of a suitable extensional fragment and of a one-variable first-order modal fragment. Since we prefer, for simplicity, not to introduce a specific formal notion of a modal fragment, we shall proceed through standard translations and rely on our usual notion of an i.a.f..

Constant Domains and Standard Translation.

Modal predicate formulae are built up from first order atomic formulae of a given first-order one-sorted relational signature $\mathcal{L}$ and from formulae of the kind $X(x)$ (where X is a unary second order variable), by using boolean connectives, individual quantifiers and a diamond operator $\Diamond$.[69]

There are actually different standard translations for first-order modal languages [16], we shall concentrate here on the translation corresponding to *constant domain* semantics. The latter is defined as follows. The signature $\mathcal{L}^W$ has, in addition to the unique sort D of $\mathcal{L}$, a new sort W; relational constants of type $D^n \to \Omega$ have corresponding relational constants in $\mathcal{L}^W$ of type $D^n W \to \Omega$. We use equal names for corresponding constants: this means for

[69] All the results in this subsection extend to the case of multimodal languages and to the case of n-ary modalities like Since, Until, etc.

instance that if P has type $D^2 \to \Omega$ in $\mathcal{L}$, the same P has type $D^2W \to \Omega$ in $\mathcal{L}^W$. We shall make the same conventions for second order variables: hence a second order $\mathcal{L}$-variable X of type $D \to \Omega$ has a corresponding second order variable X of type $DW \to \Omega$ in $\mathcal{L}^W$.

Notice that a $\mathcal{L}^W$-structure $\mathcal{A}$ is nothing but a $[\![W]\!]_{\mathcal{A}}$-indexed class of $\mathcal{L}$-structures,all having the same domain $[\![D]\!]_{\mathcal{A}}$: we indicate by $\mathcal{A}_w$ the structure corresponding to $w \in [\![W]\!]_{\mathcal{A}}$ and call it the *index structure over* w.[70]

The signature $\mathcal{L}^{WR}$ is obtained from $\mathcal{L}^W$ by adding it also a binary 'accessibility' relation R of type $WW \to \Omega$. This is the signature we need for defining the standard translation.

For a modal predicate $\mathcal{L}$-formula $\varphi[x_1^D, \dots, x_n^D]$ and for a variable $w : W$, we define the (non modal) $\mathcal{L}^{WR}$-formula $ST(\varphi, w)$ as follows:

$$\begin{aligned}
ST(\top, w) &= \top \\
ST(\bot, w) &= \bot \\
ST(X(x_i), w) &= X(x_i, w) \\
ST(P(x_{i_1}, \dots, x_{i_m}), w) &= P(x_{i_1}, \dots, x_{i_m}, w) \\
ST(\neg\psi, w) &= \neg ST(\psi, w) \\
ST(\psi_1 \vee \psi_2, w) &= ST(\psi_1, w) \vee ST(\psi_2, w) \\
ST(\psi_1 \wedge \psi_2, w) &= ST(\psi_1, w) \wedge ST(\psi_2, w) \\
ST(\Diamond\psi, w) &= \exists v^W (R(w, v) \wedge ST(\psi, v)) \\
ST(\exists x^D \psi, w) &= \exists x^D ST(\psi, w).
\end{aligned}$$

Monodic Fusions for Fragments

Next two definitions identify the ingredients of our combined problems.

Definition 5.22. *Let $\mathcal{F}_{1M}$ be a class of Kripke frames*[71] *closed under disjoint unions and isomorphisms. We call* one-variable modal fragment *induced by $\mathcal{F}_{1M}$ the interpreted algebraic fragment $\Phi_{1M} = \langle \mathcal{L}_{1M}, T_{1M}, \mathcal{S}_{1M} \rangle$, where:*

(i) $\mathcal{L}_{1M} := \mathcal{L}_{\emptyset}^{WR}$*, where $\mathcal{L}_{\emptyset}$ is the empty one-sorted first-order signature;*[72]

(ii) T_{1M} *contains the terms which are $\beta\eta$-equivalent to terms of the kind $\{w^W, x^D \mid ST(\varphi, w)\}$, where φ is a modal predicate formula having x as the only (free or bound) variable;*

[70] More precisely: $\mathcal{A}_w$ is the $\mathcal{L}$-structure having $[\![D]\!]_{\mathcal{A}}$ as a support and moreover, for $P : D^n \to \Omega$, we have that $\mathcal{I}_{\mathcal{A}_w}(P)$ contains the tuples $\underline{a} \in [\![D]\!]^n_{\mathcal{A}}$ such that $\mathcal{A} \models P(\underline{a}, w)$ holds.

[71] Recall that a Kripke frame is a nonempty set W endowed with a binary relation R.

[72] This means that $\mathcal{L}_{\emptyset}$ contains just the sort D and no other proper symbol.

(iii) $\mathcal{S}_{1M}$ *is the class of the* $\mathcal{L}_{1M}$*-structures* $\mathcal{A}$ *such that* $[\![D]\!]_{\mathcal{A}}$ *is not empty and such that the Kripke frame* $([\![W]\!]_{\mathcal{A}}, \mathcal{I}_{\mathcal{A}}(R))$ *belongs to* $\mathcal{F}_{1M}$.

Definition 5.23. *For a monadically suitable i.a.f.* $\Phi_e = \langle \mathcal{L}_e, T_e, \mathcal{S}_e \rangle$ *(recall Definition 5.16), we define the i.a.f.* $\Phi_e^W = \langle \mathcal{L}_e^W, T_e^W, \mathcal{S}_e^W \rangle$*, as follows:*

(i) T_e^W *contains the terms* $\beta\eta$*-equivalent to terms of the kind* $\{w^W, x^D \mid ST(\varphi, w)\}$*, for* $\{x^D \mid \varphi\} \in T_e$*;*

(ii) $\mathcal{S}_e^W$ *contains the* $\mathcal{L}_e^W$*-structures* $\mathcal{A}$ *whose index structures* $\mathcal{A}_w$ *are all in* $\mathcal{S}_e$*.*

We first make sure that decidability is not lost when passing from Φ_e to Φ_e^W:

Lemma 5.24. *If constraint satisfiability in a monadically suitable fragment* Φ_e *is decidable, so is constraint satisfiability in* Φ_e^W*.*

Proof. Consider a Φ_e^W-constraint:

$$\begin{aligned} &\textstyle\bigwedge_{i=1}^{n}(\{w^W, x^D \mid ST(\varphi_i, w)\} = \{w^W, x^D \mid ST(\varphi_i', w)\}) \wedge \\ &\wedge \textstyle\bigwedge_{j=1}^{m}(\{w^W, x^D \mid ST(\psi_j, w)\} \neq \{w^W, x^D \mid ST(\psi_j', w)\}); \end{aligned}$$

we claim that is is satisfiable iff the Φ_e-constraints

$$\bigwedge_{i=1}^{n}(\{x \mid \varphi_i\} = \{x \mid \varphi_i'\}) \wedge \{x \mid \psi_j\} \neq \{x \mid \psi_j'\} \tag{5}$$

are all satisfiable. One side is clear (just look at index structures); for the other side, consider Φ_e-structures $\mathcal{A}_j$ $(j = 1, \dots, m)$ satisfying (5); applying to their $\mathcal{L}_\emptyset$-reducts a large disjoint I-copy, the supports of these structures get the same cardinality, which means that such supports can be renamed to make them just the same. Since extensibility of taking disjoint I-copies is part of the definition of a monadically suitable fragment, this gives the desired model $\mathcal{A}$ for the original constraint: to this aim, it is sufficient to take $[\![W]\!]_{\mathcal{A}} := \{1, \dots, m\}$ and to let $\mathcal{A}$ be the $\mathcal{L}^W$-structure having the $\mathcal{A}_j$ so modified as index structures. □

Fix a one variable modal fragment Φ_{1M} and a first-order monadically suitable fragment Φ_e; we call *monodic fusion* of Φ_e and Φ_{1M} the combined fragment $\Phi_e^W \oplus \Phi_{1M}$.

Thus one may for instance combine packed guarded or two-variables fragments with one-variables modal fragments to get monodic fusions corresponding to the relevant cases analyzed in [58],[24].[73] In fact (modulo taking standard translation), in combined fragments like $\Phi_e^W \oplus$

[73] We recall that the two-variable fragment is monadically suitable only if we take out identity; consequently decidability of the monodic modal two-variable fragments with identity is not covered by our results (and as a matter of fact, it is not true, see [24] for the relevant pointers to the literature).

Φ_{1M} we can begin with formulae $\varphi[x]$ of Φ_e, apply to them a modal operator, then use the formulae so obtained to replace second order variables in other formulae from Φ_e, etc. Fragments of the kind $\Phi_e^W \oplus \Phi_{1M}$ formalize the intuitive notion of a monodic modal fragment whose extensional component is Φ_e. Since Φ_{1M} is also interpreted, constraint satisfiability in $\Phi_e^W \oplus \Phi_{1M}$ is restricted to a desired specific class of modal frames/flows of time.[74] We shall prove the following general transfer result for monodic fusions:

Theorem 5.25. *If the one variable modal i.a.f. Φ_{1M} and the monadically suitable i.a.f. Φ_e have decidable constraint satisfiability problems, then their monodic fusion $\Phi_e^W \oplus \Phi_{1M}$ also has decidable constraint satisfiability problems.*

If we try to use directly Theorem 5.6 to prove this result, we find problems: these problems are basically due to the fact that for the modal component the identity of two individuals living on different worlds is an important information which is completely out of the control of the extensional component. The idea is to include 'trans-world' identification into the semantics as an explicit data, following the classical suggestion of counterpart theory [42]. Since we want our alternative models to provide a semantics which is equivalent to the constant domains semantics, the most elegant solution seems to be that of representing individual domains as descent data.

An Alternative Translation.

Fix a set W; we call *descent data* for W[75] a triple (E, p, θ) where $p : E \longrightarrow W$ and $\theta : E \times W \longrightarrow E$ are functions satisfying the following three requirements for all $e \in E, w, w_1, w_2 \in W$:[76]

$$\begin{aligned} p(\theta(e, w)) &= w & (6) \\ \theta(e, p(e)) &= e & (7) \\ \theta(\theta(e, w_1), w_2) &= \theta(e, w_2). & (8) \end{aligned}$$

To understand this definition from a modal point of view, we may think of E as the domain of all possible individuals and of p as the function that associates with an individual e the

[74]The class of Kripke frames on which Φ_{1M} is based is taken to be closed under disjoint unions, but this assumption is not really relevant for relativized satisfiability (i.e. for conditional word problems): notice that a constraint containing at most one negative literal is satisfied in a disjoint union iff it is satisfied in a component, hence one can always close under disjoint unions the class of Kripke frames under consideration, without loss of generality, as far as relativized satisfiability is concerned.

[75]One should better say 'descent data for the unique map $W \longrightarrow \{*\}$'. Descent theory is a powerful and deep theory in pure mathematics (see e.g [35], [34]).

[76]From now on, we shall use the letters $e, w, \ldots$ both for (sorted) variables in a logical language and for concrete elements of given structures (and for the names of these concrete elements in expanded languages, like in Subsection 2.4): the context will carefully clarify, case by case, the intended use.

world $p(e)$ where e lives: in this sense, $\theta(e, w)$ has to be thought as the counterpart of e in the world w. Since we want to mimic a constant domain semantics, we would like from conditions (6)-(8) to follow that counterparts behave in such a way that fibers over W are in fact 'constant'. This is true (provided W is not empty, which means that the map $W \to \{*\}$ is onto) by a 'descent theorem' that holds for instance in exact categories: we shall explain (and check) what we need just for the very easy special case we are interested in.

We call *canonical descent data* the descent data of the kind $(D \times W, p_W, \theta_W)$ where D is a set, p_W is the projection on the second component and $\theta_W(\langle d, w' \rangle, w) = \langle d, w \rangle$. Now the descent theorem says the following: *every triple* (E, p, θ) *forming descent data for a non empty set* W *is isomorphic to a canonical descent data.* This is proved as follows: take the equivalence relation over E given by $e_1 \simeq e_2$ iff (there is $w \in W$ such that $\theta(e_1, w) = e_2$). Let now D be the quotient set of E under $\simeq$. To check that the canonical descent data $(D \times W, p_W, \theta_W)$ are isomorphic to (E, p, θ), consider the bijective function $h : E \to D \times W$ associating with $e \in E$ the pair $h(e) := \langle [e], p(e) \rangle$ and observe that this bijection commutes with p and θ (in the sense that we have $p_W(h(e)) = p(e)$ and $\theta_W(h(e), w) = h(\theta(e, w))$).

The idea is now that of using descent data to define a new translation for modal first order formulae: this translation (corresponding to an alternative equivalent 'descent' semantics) has the first advantage that we do not need to modify the original first order signature $\mathcal{L}$ by altering the types of the relational symbols in it (as it happened for the definition of $\mathcal{L}^W$ in the case of the standard translation).

Let $\mathcal{L}$ be a first-order relational one-sorted signature; the signature $\mathcal{L}^d$ is obtained by changing the name of the unique sort of $\mathcal{L}$ from D to E and then by adding to $\mathcal{L}$ a new sort W, a binary relation $R : WW \to \Omega$ and function symbols $p : E \to W$, $\theta : EW \to E$.

For a modal $\mathcal{L}$-formula $\varphi[e_1, \ldots, e_n]$ (here $e_1, \ldots, e_n$ are just individual variables of the unique sort D of $\mathcal{L}$) and for a variable $w : W$, we define the (non modal) $\mathcal{L}^d$-formula $DT(\varphi, w)$

as follows:

$$
\begin{aligned}
DT(\top, w) &= \top \\
DT(\bot, w) &= \bot \\
DT(X(e_i), w) &= X(e_i) \\
DT(P(e_{i_1}, \dots, e_{i_m}), w) &= P(e_{i_1}, \dots, e_{i_m}) \\
DT(\neg\psi, w) &= \neg DT(\psi, w) \\
DT(\psi_1 \vee \psi_2, w) &= DT(\psi_1, w) \vee DT(\psi_2, w) \\
DT(\psi_1 \wedge \psi_2, w) &= DT(\psi_1, w) \wedge DT(\psi_2, w) \\
DT(\Diamond\psi, w) &= \exists v^W (R(w, v) \wedge DT(\psi, w)[\theta(e_1, v), \dots, \theta(e_n, v), v]) \\
DT(\exists e^D \psi, w) &= \exists e^E (p(e) = w \wedge DT(\psi, w)),
\end{aligned}
$$

where, according to our usual conventions, the notation $DT(\psi, w)[\theta(e_1, v), \dots, \theta(e_n, v), v]$ means the formula obtained from $DT(\psi, w)$ by replacing w by v and the free variables e_i by the $\mathcal{L}^d$-terms $\theta(e_i, v)$. In the following, we shall use notations like $DT(\psi[u_1, \dots, u_n], t)$, where $u_1, \dots, u_n$ are $\mathcal{L}^d$-term of type E and t is an $\mathcal{L}^d$-term of type W, to mean the formula obtained from $DT(\psi[x_1, \dots, x_n], w)$ by appling it the substitution $x_1 \mapsto u_1, \dots, x_n \mapsto u_n, w \mapsto t$.

We now *reformulate* the notions of a one variable modal i.a.f. and of a monadically suitable i.a.f. into this equivalent alternative descent semantics. We shall be interested in the terms of the kind $\{e \mid DT(\varphi[e], p(e))\}$, where φ is a first order modal formula. However, these terms are not precisely closed under substitutions for second order variables: to get closure under substitution, some equational reasoning (partially based on the descent equations) would be needed, in addition to $\beta\eta$-equivalence. Since we are not precisely interested here in investigating the related technical details, we prefer to close under substitution the set of terms we need to build our i.a.f.'s. Such an operation of taking the smallest substitution closed set $cl(T)$ of $\mathcal{L}^d$-terms extending a given set T of $\mathcal{L}^d$-terms is rather harmless for our purposes: for instance, in case T contains the relevant variables (as it will always be the case in this Section), we can use a mechanism similar to the purification procedure explained in Subsection 4.1 and assume without loss of generality that the $cl(T)$-terms appearing in constraints are in fact terms from the original set T.

We first translate one-variable modal fragments into descent semantics:

Definition 5.26. *Let $\Phi_{1M} = \langle \mathcal{L}_{1M}, T_{1M}, \mathcal{S}_{1M} \rangle$ be the one-variable modal fragment induced by the frame class $\mathcal{F}_{1M}$. The i.a.f. $\Phi^d_{1M} = \langle \mathcal{L}^d_{1M}, T^d_{1M}, \mathcal{S}^d_{1M} \rangle$ is so defined:*

(i) *T^d_{1M} is the substitution closure of the set of the terms of the kind $\{e^E \mid DT(\varphi[e], p(e))\}$, where $\varphi[x]$ is a one-variable modal predicate formula;*

(ii) $\mathcal{S}^d_{1M}$ *contains the* $\mathcal{L}^d_{1M}$*-structures* $\mathcal{A}$ *such that* $[\![E]\!]_\mathcal{A}$, $\mathcal{I}_\mathcal{A}(p)$, $\mathcal{I}_\mathcal{A}(\theta)$ *satisfy the descent data equations* (6)-(8), $\mathcal{I}_\mathcal{A}(p)$ *is surjective*[77] *and the Kripke frame* $([\![W]\!]_\mathcal{A}, \mathcal{I}_\mathcal{A}(R))$ *belongs to* $\mathcal{F}_{1M}$.

Suppose now we are given a monadically suitable i.a.f. $\Phi_e = \langle \mathcal{L}_e, T_e, \mathcal{S}_e \rangle$; this is a purely extensional fragment, hence the DT-translations of the terms in it do not contain the descent multiplication θ: this is why we can translate Φ_e into the signature $\mathcal{L}^{-d}_e$ (the latter is like $\mathcal{L}^d_e$ except that θ and R are omitted). The related definition is the following one:

Definition 5.27. *Let* $\Phi_e = \langle \mathcal{L}_e, T_e, \mathcal{S}_e \rangle$ *be a monadically suitable i.a.f.. Let* Φ^{-d}_e *be the i.a.f.* $\langle \mathcal{L}^{-d}_e, T^{-d}_e, \mathcal{S}^{-d}_e \rangle$, *where:*

(i) T^{-d}_e *is the substitution closure of the set of the terms* $\{e^E \mid DT(\varphi[e], p(e))\}$ *such that* $\{x \mid \varphi[x]\} \in T_e$;

(ii) $\mathcal{S}^{-d}_e$ *contains the* $\mathcal{L}^{-d}_e$*-structures* $\mathcal{A}$ *which are isomorphic to* $[\![W]\!]_\mathcal{A}$*-indexed disjoint unions of structures* $\mathcal{A}_w$ *from* $\mathcal{S}_e$ *(these* $\mathcal{A}_w$, *varying* $w \in [\![W]\!]_\mathcal{A}$, *are called the* fiber components *of* $\mathcal{A}$*).*

In full details, condition (ii) from Definition 5.27 means the following. Notice first that it makes sense to consider $\mathcal{L}_e$ as a subsignature of $\mathcal{L}^{-d}_e$ (modulo the change of name of the unique sort of $\mathcal{L}_e$ from D to E). Thus in (ii) we are asking that $\mathcal{A} \in \mathcal{S}^{-d}_e$ if and only if (up to isomorphism): (a) the $\mathcal{L}_e$-reduct of $\mathcal{A}$ is the disjoint union $\sum_{w \in [\![W]\!]_\mathcal{A}} \mathcal{A}_w$ of some $\mathcal{L}_e$-structures $\mathcal{A}_w$, varying $w \in [\![W]\!]_\mathcal{A}$; (b) the $\mathcal{L}_e$-structures $\mathcal{A}_w$ all belong to $\mathcal{S}_e$; (c) $\mathcal{I}_\mathcal{A}(p)$ is given by $\mathcal{I}_\mathcal{A}(p)(d, w) = w$, for every $(d, w) \in [\![E]\!]_\mathcal{A}$[78] (notice that since the supports of the $\mathcal{A}_w$'s are not empty, $\mathcal{I}_\mathcal{A}(p)$ is surjective).

In other words, this means that *the following quite simple schema builds (up to isomorphism) precisely the structures in* $\mathcal{S}^{-d}_e$: take a family $\{\mathcal{A}_w \mid w \in I\}$ of structures from $\mathcal{S}_e$, interpret W as the index set I, E as the disjoint union of the supports of the $\mathcal{A}_w$, all predicates as in standard disjoint union of relational structures and p as the functions associating with an element the index of its support.

We now make an important observation, to be fixed in Lemma 5.28 below. If the $\mathcal{L}^{-d}_e$-structure $\mathcal{A}$ is isomorphic to the $[\![W]\!]_\mathcal{A}$-indexed disjoint union of the $\mathcal{L}_e$-structures $\mathcal{A}_w$, *satisfiability in* $\mathcal{A}$ *of* Φ^{-d}_e*-constraints is fiberwise*, in the following sense. For every $w \in [\![W]\!]_\mathcal{A}$, for every $d_1, \dots, d_n$ in the support of $\mathcal{A}_w$ and for every (non modal) $\mathcal{L}_e$-formula $\varphi[x_1, \dots, x_n]$,

[77] This requirement corresponds to the fact that the domain of the constant domain semantics is not empty and, on the basis of the descent equations, it is equivalent to the fact that $[\![E]\!]_\mathcal{A}$ is not empty.

[78] By (a), the elements of $[\![E]\!]_\mathcal{A}$ can be represented as pairs (d, w), where $w \in [\![W]\!]_\mathcal{A}$ and d is from the support of $\mathcal{A}_w$.

we have

$$\mathcal{A} \models DT(\varphi[(d_1,w),\ldots,(d_n,w)],w) \quad \text{iff} \quad \mathcal{A}_w \models \varphi[d_1,\ldots,d_n] \tag{9}$$

(a trivial induction is sufficient to establish this fact). Notice that (9) holds also in case $\mathcal{L}_e$ is expanded by free constants $\underline{c}$ whose type is a Φ_e-type (there is only one Φ_e-type, namely the type of the subsets of the domain). Since $[\![E]\!]_{\mathcal{A}}$ is the disjoint union of the supports of the $\mathcal{A}_w$, the interpretation of the $\underline{c}$'s in $\mathcal{A}$ is obtained by gluing their restrictions to the supports of the $\mathcal{A}_w$. Thus, if we consider $\mathcal{A}$ as a $\mathcal{L}_e(\underline{c})$-structure, it is still isomorphic to the $[\![W]\!]_{\mathcal{A}}$-indexed disjoint unions of $\mathcal{L}_e(\underline{c})$-structures (that we still call $\mathcal{A}_w$). Now a $\Phi_e^{-d}(\underline{c})$-closed constraint is formed by positive literals

$$\{e^E \mid DT(\varphi_i[e],p(e))\} = \{e^E \mid DT(\varphi_i'[e],p(e))\} \quad (i=1,\ldots,n)$$

and by negative literals

$$\{e^E \mid DT(\psi_j[e],p(e))\} \neq \{e^E \mid DT(\psi_j'[e],p(e))\} \quad (j=1,\ldots,m);$$

according to (9), such a constraint is satisfied in $\mathcal{A}$ iff (i) for every $w \in [\![W]\!]_{\mathcal{A}}$, we have $\mathcal{A}_w \models \bigwedge_i(\{x \mid \varphi_i[x]\} = \{x \mid \varphi_i'[x]\})$[79] and (ii) for every j there is $w_j \in [\![W]\!]_{\mathcal{A}}$ such that $\mathcal{A}_{w_j} \models \{x \mid \psi_j[x]\} \neq \{x \mid \psi_j'[x]\}$. Of course, the same observation applies to generalized constraints too.

To sum up, we introduce the following notion: if a generalized $\Phi_e^{-d}(\underline{c})$-closed constraint Γ is given, a $\Phi_e(\underline{c})$-positive literal of Γ is a positive literal of the form $\{x \mid \varphi[x]\} = \{x \mid \varphi'[x]\}$ such that $\{e^E \mid DT(\varphi[e],p(e))\} = \{e^E \mid DT(\varphi'[e],p(e))\} \in \Gamma$ (the definition of a $\Phi_e(\underline{c})$-negative literal of Γ is analogous). The above observation now reads as:

Lemma 5.28. *Let $\Phi_e = \langle \mathcal{L}_e, T_e, \mathcal{S}_e \rangle$ be a monadically suitable i.a.f. and let Φ_e^{-d} be the i.a.f. of Definition 5.27. Given free constants $\underline{c}$, suppose that the $\mathcal{L}_e^{-d}(\underline{c})$-structure $\mathcal{A}$ is isomorphic to the $[\![W]\!]_{\mathcal{A}}$-indexed disjoint union of the $\mathcal{L}_e(\underline{c})$-structures $\mathcal{A}_w \in \mathcal{S}_e(\underline{c})$. Now $\mathcal{A}$ satisfies a generalized $\Phi_e^{-d}(\underline{c})$-closed constraint Γ iff the $\Phi_e(\underline{c})$-positive literals of Γ hold in all fiber components $\mathcal{A}_w$ and every $\Phi_e(\underline{c})$-negative literal of Γ holds in at least one fiber component $\mathcal{A}_w$.*

Remark 5.29. The following strong consequence of the above Lemma will be repeatedly used in the following: we know that a structure $\mathcal{A} \in \mathcal{S}_e^{-d}(\underline{c})$ is isomorphic to a $[\![W]\!]_{\mathcal{A}}$-indexed disjoint union of structures $\mathcal{A}_w$ from $\mathcal{S}_e(\underline{c})$. Suppose now that we replace, in such a disjoint

[79] In more detail: we have that $\mathcal{A} \models \{e^E \mid DT(\varphi_i[e],p(e))\} = \{e^E \mid DT(\varphi_i'[e],p(e))\}$ iff for every $w \in [\![W]\!]_{\mathcal{A}}$ and for every d in the support of $\mathcal{A}_w$, we have $\mathcal{A} \models DT(\varphi[(d,w)] \leftrightarrow \varphi'[(d,w)],w)$. By (9), this is the same as $\mathcal{A}_w \models \varphi[d] \leftrightarrow \varphi'[d]$ for all w and d in $\mathcal{A}_w$, which means that the $\Phi_e(\underline{c})$-atom $\{x \mid \varphi_i[x]\} = \{x \mid \varphi_i'[x]\}$ is true in all fiber structures $\mathcal{A}_w$.

union, the structures $\mathcal{A}_w$ by some structures $\mathcal{A}'_w \in \mathcal{S}_e(\underline{c})$ such that $\mathcal{A}_w \equiv_{\Phi_e(\underline{c})} \mathcal{A}'_w$: call $\mathcal{A}'$ the $[\![W]\!]_{\mathcal{A}'}$-indexed structure obtained in this way (here the interpretation of W has not changed, namely we have $[\![W]\!]_{\mathcal{A}'} := [\![W]\!]_{\mathcal{A}}$). Clearly $\mathcal{A}' \in \mathcal{S}_e^{-d}(\underline{c})$ and Lemma 5.28 implies that we have $\mathcal{A} \equiv_{\Phi_e^{-d}(\underline{c})} \mathcal{A}'$: thus $\Phi_e^{-d}(\underline{c})$*-equivalence is preserved, whenever we apply fiberwise a* $\Phi_e(\underline{c})$*-equivalence preserving construction.*

Proof of the Monodic Decidability Transfer Result.

Let now Φ_{1M} be a one-variable modal i.a.f. and Φ_e be a monodically suitable i.a.f.. Our plan is the following: *we first check that* $\Phi_e^W \oplus \Phi_{1M}$ *can be equivalently replaced by* $\Phi_e^{-d} \oplus \Phi_{1M}^d$ *and then we apply Theorem 5.6 to the latter.*

The first part of the plan just consists of unwinding the definitions we gave. In fact Φ_e^W and Φ_e^{-d} (as well as Φ_{1M} versus Φ_{1M}^d, and $\Phi_e^W \oplus \Phi_{1M}$ versus $\Phi_e^{-d} \oplus \Phi_{1M}^d$), are basically the same i.a.f.; however for our purposes the statement of the following lemma is sufficient:

Lemma 5.30. *Satisfiability of pure constraints in* $\Phi_e^W \oplus \Phi_{1M}$ *can be reduced to satisfiability of pure constraints in* $\Phi_e^{-d} \oplus \Phi_{1M}^d$*, and vice versa. Constraint satisfiability for* Φ_e^W *(resp. for* Φ_{1M}*) can also be reduced to constraint satisfiability for* Φ_e^{-d} *(resp. for* Φ_{1M}^d*), and vice versa.*

Proof. A pure $\Phi_e^W \oplus \Phi_{1M}$-constraint contains equations and inequations among terms of the kind $\{w^W, x^D \mid ST(\varphi, w)\}$, where $\varphi[x]$ is either a one-variable modal predicate formula or it is such that $\{x \mid \varphi\}$ is a Φ_e-term. On the other hand, a pure $\Phi_e^{-d} \oplus \Phi_{1M}^d$-constraint contains equations and inequations among terms of the kind $\{e^E \mid DT(\varphi[e], p(e))\}$, where $\varphi[x]$ is either a one-variable modal predicate formula or it is such that $\{x \mid \varphi\}$ is a Φ_e-term. Hence it is clear how to convert a pure $\Phi_e^W \oplus \Phi_{1M}$-constraint Γ into a pure $\Phi_e^{-d} \oplus \Phi_{1M}^d$-constraint Γ^d, and vice versa: it remains to show the equi-satisfiability.

To any $\Phi_e^W \oplus \Phi_{1M}$-structure $\mathcal{A}$ we can associate a $\Phi_e^{-d} \oplus \Phi_{1M}^d$-structure $\mathcal{A}^d$ as follows: let $[\![W]\!]_{\mathcal{A}^d} := [\![W]\!]_{\mathcal{A}}$ and $\mathcal{I}_{\mathcal{A}^d}(R) := \mathcal{I}_{\mathcal{A}}(R)$; the symbols E, p, θ are interpreted as canonical descent data for $[\![W]\!]_{\mathcal{A}}$ relatively to $[\![D]\!]_{\mathcal{A}}$. Thus by definition $[\![E]\!]_{\mathcal{A}^d} = [\![D]\!]_{\mathcal{A}} \times [\![W]\!]_{\mathcal{A}}$, so it makes sense to put $\mathcal{I}_{\mathcal{A}^d}(P) := \{\langle (d_1, w), \dots, (d_n, w)\rangle \mid \langle d_1, \dots, d_n, w\rangle \in \mathcal{I}_{\mathcal{A}}(P)\}$ for every predicate symbol P having type $D^n \to \Omega$ in $\mathcal{L}$.[80] Actually every $\Phi_e^{-d} \oplus \Phi_{1M}^d$-structure is isomorphic to one of the kind $\mathcal{A}^d$, for some $\Phi_e^W \oplus \Phi_{1M}$-structure $\mathcal{A}$: in fact, the $\mathcal{L}_e^{-d}$-reduct of a $\Phi_e^{-d} \oplus \Phi_{1M}^d$-structure $\mathcal{B}$ is a $[\![W]\!]_{\mathcal{B}}$-disjoint union of $\mathcal{L}_e$-structures $\mathcal{B}_w$ (see Definition 5.27(ii)), whereas *by the descent theorem* we can assume that, up to isomorphism, the descent data in $\mathcal{B}$ are canonical (the combination of these two facts means that $\mathcal{B} \simeq \mathcal{A}^d$ for some $\mathcal{A}$).[81]

[80] Recall that, in correspondence to a $P : D^n \to \Omega$, the signature $\mathcal{L}^{WR}$ contains $P : D^n W \to \Omega$, whereas $\mathcal{L}^d$ contains $P : E^n \to \Omega$.

[81] In full detail: since descent data in $\mathcal{B}$ are canonical, we can assume that $[\![E]\!]_{\mathcal{B}} = \tilde{D} \times [\![W]\!]_{\mathcal{B}}$, for some $\tilde{D}$

That Γ is satisfied in $\mathcal{A}$ iff Γ^d is satisfied in $\mathcal{A}^d$ is straightforward: to see it, just check by induction that, for every modal formula $\varphi[x_1, \dots, x_n]$,

$$\mathcal{A} \models ST(\varphi[d_1, \dots, d_n], w) \quad \text{iff} \quad \mathcal{A}^d \models DT(\varphi[(d_1, w), \dots, (d_n, w)], w) \tag{10}$$

holds for all $d_1 \dots, d_n \in [\![D]\!]_{\mathcal{A}}$ and all $w \in [\![W]\!]_{\mathcal{A}}$ (under the corresponding assignments to the second order variables).[82]

The statement for the fragments Φ_{1M} and Φ^d_{1M} is shown in the same way. For fragments Φ^W_e and Φ^{-d}_e, we need a preliminary observation, because the supports of the fiber components in a Φ^{-d}_e-structure may not be isomorphic (there are no full descent data here). The observation is the following: Φ_e is monadically suitable, so by taking suitably large disjoint I-copies we can expand the cardinalities of the non empty supports in the fiber components of a Φ^{-d}_e-structure and make such supports to coincide, up to renaming of their elements (see Remark 5.29). At this point, canonical descent data exists and the above argument based on (10) applies. $\square$

In view of Lemma 5.30, to complete the proof of Theorem 5.25 it is sufficient now to show the following:

Proposition 5.31. *If the one variable modal i.a.f. Φ_{1M} and the monadically suitable i.a.f. Φ_e have decidable constraint satisfiability problems, then the combined fragment $\Phi^{-d}_e \oplus \Phi^d_{1M}$ also has decidable constraint satisfiability problems.*

Proof. Since by Lemmas 5.24, 5.30 the component fragments Φ^{-d}_e and Φ^d_{1M} have decidable constraint satisfiability problems, we can try to apply Theorem 5.6, by checking the remaining conditions.

Notice that both i.a.f.'s have $E \to \Omega$ as the only type for their terms and that the shared signature $\mathcal{L}_0$ contains the two sorts E, W and the function symbol p of type $E \to W$. According to the definitions of a one variable modal and of a monadically suitable fragment, the set of terms T_0 in the shared fragment $\Phi_0 = \langle \mathcal{L}_0, T_0, \mathcal{S}_0 \rangle$ contains the terms of the kind $\{e^E \mid DT(\varphi[e], p(e))\}$, where $\varphi[x]$ is a one-variable (non modal) first-order formula in the empty one-sorted signature $\mathcal{L}_\emptyset$: this implies, in particular, that Φ_0 *is effectively locally finite*

and that the descent symbols p, θ are interpreted in the canonical way. We define $\mathcal{A}$ by taking $[\![W]\!]_{\mathcal{A}} := [\![W]\!]_{\mathcal{B}}$, $\mathcal{I}_{\mathcal{A}}(R) := \mathcal{I}_{\mathcal{B}}(R)$, $[\![D]\!]_{\mathcal{A}} := \tilde{D}$ and $\mathcal{I}_{\mathcal{A}}(P) := \{\langle d_1, \dots, d_n, w\rangle \mid \langle (d_1, w), \dots, (d_n, w)\rangle \in \mathcal{I}_{\mathcal{B}}(P)\}$, for every n-ary predicate symbol P (the index structures of $\mathcal{A}$ are now isomorphic to the corresponding fiber components of $\mathcal{B}$, hence $\mathcal{A} \in \mathcal{S}^W_e$). Since the $\mathcal{L}_e$-reduct of $\mathcal{B}$ is the disjoint union of the $\mathcal{B}_w$, it turns out that $\mathcal{B} \simeq \mathcal{A}^d$.

[82] By 'corresponding' we obviously mean here that (d, w) belongs to the subset assigned to X in $\mathcal{A}^d$ iff (d, w) belongs to the subset assigned to X in $\mathcal{A}$. To understand this and to check inductively the above condition, recall that descent data in $\mathcal{A}^d$ are canonical (thus, e.g. elements of $[\![E]\!]_{\mathcal{A}^d}$ are pairs (d, w), the symbol p is interpreted as the second projection, etc.).

(only second order variables for subsets, Boolean connectives and the quantifier $\forall x, \exists x$ can be used to build φ). We do not have complete information about the $\mathcal{L}_0$-structures $\mathcal{A} \in \mathcal{S}_0$, but we know that $\mathcal{I}_{\mathcal{A}}(p)$ is always surjective in them (because the interpretation of p is a surjective function in structures coming from both $\mathcal{S}^d_{1M}$ and $\mathcal{S}^{-d}_e$).

We need to identify suitable structural operations on Φ_0 to apply Theorem 5.6,[83] but first we study invariants for $\mathcal{L}_0(\underline{c}_0)$-structures. To this aim, suppose we are given free constants $\underline{c}_0 := \{P_1, \dots, P_n\}$ (all having type $E \to \Omega$, which is the only Φ_0-type) and a $\mathcal{L}_0(\underline{c}_0)$-structure $\mathcal{A}$. For $w \in [\![W]\!]_{\mathcal{A}}$, we denote by $\mathcal{A}_w$ the fiber component over w: this is the $\mathcal{L}_\emptyset(\underline{c}_0)$-structure whose support is given by the $e \in [\![E]\!]_{\mathcal{A}}$ such that $\mathcal{I}_{\mathcal{A}}(p)(e) = w$ (in $\mathcal{A}_w$ the predicates P_j are interpreted by taking the restriction of $\mathcal{I}_{\mathcal{A}}(P_j)$ to the support of $\mathcal{A}_w$).

Consider, as in the proof of Proposition 5.12, the boolean combinations of the form $\varepsilon(e) = Q_1(e) \wedge \cdots \wedge Q_n(e)$ where $Q_j \equiv P_j$ or $Q_j \equiv \neg P_j$. With each $w \in [\![W]\!]_{\mathcal{A}}$ we associate the 2^n cardinal invariants of Proposition 5.12 for the fiber component $\mathcal{A}_w$. That is: for $\varepsilon = 1, \dots, 2^n$, let $\kappa_\varepsilon(w)$ be the cardinality of the set of the $e \in [\![E]\!]_{\mathcal{A}}$ living in the support of the w-fiber component such that $\mathcal{A}_w \models \varepsilon(e)$; also, let $\mu(w)$ be equal to the set of the ε such that $\kappa_\varepsilon(w) > 1$. Finally, let $\mu(\mathcal{A})$ be the set of sets formed by the $\mu(w)$, varying $w \in [\![W]\!]_{\mathcal{A}}$. Notice that the fact that $\mu(\mathcal{A}) = J$ is equivalent to

$$\begin{aligned} \mathcal{A} \models \quad & \textstyle\bigwedge_{S \in J} \exists w \, [\bigwedge_{\varepsilon \in S} \exists e \, (p(e) = w \wedge \varepsilon(e)) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists e \, (p(e) = w \wedge \varepsilon(e))] \wedge \\ & \wedge \textstyle\bigwedge_{S \notin J} \neg \exists w \, [\bigwedge_{\varepsilon \in S} \exists e \, (p(e) = w \wedge \varepsilon(e)) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists e \, (p(e) = w \wedge \varepsilon(e))] \end{aligned}$$

i.e. to

$$\begin{aligned} \mathcal{A} \models \quad & \textstyle\bigwedge_{S \in J} \exists w \, DT[\bigwedge_{\varepsilon \in S} \exists x \, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x \, \varepsilon(x), w] \wedge \\ & \wedge \textstyle\bigwedge_{S \notin J} \neg \exists w \, DT[\bigwedge_{\varepsilon \in S} \exists x \, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x \, \varepsilon(x), w] \end{aligned}$$

Since $\mathcal{I}_{\mathcal{A}}(p)$ is surjective, this is the same as

$$\begin{aligned} \mathcal{A} \models \quad & \textstyle\bigwedge_{S \in J} \exists e \, DT[\bigwedge_{\varepsilon \in S} \exists x \, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x \, \varepsilon(x), p(e)] \wedge \\ & \wedge \textstyle\bigwedge_{S \notin J} \neg \exists e \, DT[\bigwedge_{\varepsilon \in S} \exists x \, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x \, \varepsilon(x), p(e)]; \end{aligned}$$

the latter simply says that the $\Phi_0(\underline{c}_0)$-closed constraint

$$\begin{aligned} & \textstyle\bigwedge_{S \in J} (\{e \mid DT[\bigwedge_{\varepsilon \in S} \exists x \, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x \, \varepsilon(x), p(e)]\} \neq \{e \mid DT(\bot, p(e))\}) \wedge \\ & \wedge \textstyle\bigwedge_{S \notin J} (\{e \mid DT[\bigwedge_{\varepsilon \in S} \exists x \, \varepsilon(x) \wedge \bigwedge_{\varepsilon \notin S} \neg \exists x \, \varepsilon(x), p(e)]\} = \{e \mid DT(\bot, p(e))\}) \end{aligned}$$

is satisfied in $\mathcal{A}(\underline{c}_0)$.

[83]In the statement of Theorem 5.6, we take $\Phi_0^\star$ equal to Φ_0, so only condition 5.6 (4) has not yet been checked.

Suppose now that the $\mathcal{L}_0(\underline{c}_0)$-structures $\mathcal{A}_1$ and $\mathcal{A}_2$ are $\Phi_0(\underline{c}_0)$-equivalent: then we have $\mu(\mathcal{A}_1) = \mu(\mathcal{A}_2)$, as explained above. *We will make $\mathcal{A}_1$ and $\mathcal{A}_2$ $\Phi_0(\underline{c}_0)$-isomorphic in two steps: each step needs a Φ_0-structural operation that will be proved to be extensible both to Φ_e^{-d} and to Φ_{1M}^d* (notice that the composition of extensible structural operations is extensible).

The *first structural operation* is taking disjoint I-copies $\sum_I$: notice that taking disjoint I-copies operation applies to structures over a multi-sorted first-order relational language having also *unary* function symbols (like our p).[84] Consequently, the operation applies to the language $\mathcal{L}_e^{-d}$ (and, in particular, to $\mathcal{L}_0$); it is a Φ_e^{-d}-structural operation (hence its restriction to $\mathcal{L}_0$ is in particular Φ_e^{-d}-extensible), because the fiber components in $\sum_I \mathcal{A}$ are the same as the fiber components in $\mathcal{A}$ (we just made more copies of each fiber component) and Lemma 5.28 applies, guaranteeing that $\mathcal{A} \equiv_{\Phi_e^{-d}(\underline{c}_0)} \sum_I \mathcal{A}$.

The operation $\sum_I$ is Φ_{1M}^d-extensible: to see this, first observe that the frame class $\mathcal{F}_{1M}$ defining Φ_{1M} is closed under disjoint unions. If $\mathcal{A}$ is now a structure from $\mathcal{S}_{1M}(\underline{c}_0)$, in order to make $\sum_I \mathcal{A}$ a $\mathcal{S}_{1M}(\underline{c}_0)$-structure as well, we only need to introduce a descent multiplication on $\sum_I \mathcal{A}$. This is done as follows: put $\mathcal{I}_{\sum_I \mathcal{A}}(\theta)((e,i),(w,j)) := (\mathcal{I}_{\mathcal{A}}(\theta)(e,w), j)$. The descent equations (6)-(8) can be checked in a straightforward way. In addition, to show that truth of closed $\Phi_{1M}^d(\underline{c}_0)$-atoms is preserved, one simply check by induction that, for every modal one-variable $\mathcal{L}_\emptyset$-formula $\varphi[x]$ (in which second order variables have been replaced by the free constants $\underline{c}_0$), for every index $i \in I$, for every $e \in [\![E]\!]_{\mathcal{A}}$, for every $w \in [\![W]\!]_{\mathcal{A}}$, if $\mathcal{I}_{\mathcal{A}}(p)(e) = w$, then we have

$$\sum_I \mathcal{A} \models DT(\varphi[(e,i)], (w,i)) \quad \text{iff} \quad \mathcal{A} \models DT(\varphi[e], w).$$

If we apply $\sum_I$ for sufficient large I to our $\Phi_0(\underline{c}_0)$-equivalent structures $\mathcal{A}_1$ and $\mathcal{A}_2$, then we get $\mathcal{L}_0(\underline{c}_0)$-structures $\mathcal{A}_1'$ and $\mathcal{A}_2'$[85] such that for every S the cardinality of the $w_1 \in [\![W]\!]_{\mathcal{A}_1'}$ with $\mu(w_1) = S$ is the same as the cardinality of the $w_2 \in [\![W]\!]_{\mathcal{A}_2'}$ with $\mu(w_2) = S$. Thus we have a bijection $\iota_W : [\![W]\!]_{\mathcal{A}_1'} \to [\![W]\!]_{\mathcal{A}_2'}$, preserving the invariant μ. To make $\mathcal{A}_1'$ and $\mathcal{A}_2'$ $\Phi_0(\underline{c}_0)$-isomorphic, we need the fiber components over w and $\iota(w)$ to be isomorphic (for all w): to that aim, since $\mu(w)$ is equal to $\mu(\iota(w))$, it is sufficient to apply 'fiberwise' the argument of Proposition 5.12, provided we are allowed to take suitably large disjoint I-copies *of the sets $[\![E]\!]_{\mathcal{A}_1'}$ and $[\![E]\!]_{\mathcal{A}_2'}$ only.* This will be achieved through the *second structural operation* we are going to introduce.

Let $\mathcal{A}$ be a $\mathcal{L}_0(\underline{c}_0)$-structure and let I be a non empty set of indices; we call $\sum_I^E(\mathcal{A})$ the $\mathcal{L}_0(\underline{c})$-structure so defined: (i) we interpret the sort W as in $\mathcal{A}$ and the sort E as the disjoint union $\sum_I [\![E]\!]_{\mathcal{A}}$; (ii) we interpret the symbol p as the function mapping (e,i) to $\mathcal{I}_{\mathcal{A}}(p)(e)$; (iii)

[84] This is because one can put, for unary p and $i \in I$, $\mathcal{I}_{\sum_I \mathcal{A}}(p)(e,i) := (\mathcal{I}_{\mathcal{A}}(p)(e), i)$.

[85] $\mathcal{A}_1'$ and $\mathcal{A}_2'$ are still $\Phi_0(\underline{c}_0)$-equivalent (in fact, we observed that, more generally, truth of closed $\Phi_{1M}^d(\underline{c}_0)$- and $\Phi_e^{-d}(\underline{c}_0)$-literals is preserved).

we interpret the unary predicate $P \in \underline{c}_0$ as the set of all (e, i) such that $e \in \mathcal{I}_{\mathcal{A}}(P)$. That $\sum_I^E(\mathcal{A})$ is $\Phi_0(\underline{c}_0)$-equivalent to $\mathcal{A}$ will be checked below (directly for the stronger cases of $\Phi_{1M}^d(\underline{c}_0)$- and of $\Phi_e^{-d}(\underline{c}_0)$-equivalence).

The operation $\sum_I^E$ is Φ_e^{-d}-extensible: this comes from Remark 5.29 and from the fact that Φ_e is a monadically suitable fragment (so that taking disjoint copies is Φ_e-extensible at each fiber component separately).

Finally, the operation $\sum_I^E$ is Φ_{1M}^d-extensible: we can interpret the descent multiplication symbol θ in $\sum_I^E(\mathcal{A})$ into the function associating $(\mathcal{I}_{\mathcal{A}}(\theta)(e, w), i)$ with the pair $((e, i), w)$ (equations (6)-(8) are easily checked).[86] Of course, the accessibility relation R is interpreted in $\sum_I^E(\mathcal{A})$ as it was interpreted in $\mathcal{A}$. Thus it remains to prove that $\mathcal{A} \equiv_{\Phi_{1M}^d(\underline{c}_0)} \sum_I^E(\mathcal{A})$: to this aim, it is sufficient to check inductively that for every one-variable modal formula $\varphi[x]$, for $w \in [\![W]\!]_{\mathcal{A}}, e \in [\![E]\!]_{\mathcal{A}}$ (such that $\mathcal{I}_{\mathcal{A}}(p)(e) = w$) and for $i \in I$, we have that

$$\sum_I^E(\mathcal{A}) \models DT(\varphi[(e, i)], w) \quad \text{iff} \quad \mathcal{A} \models DT(\varphi[e], w).$$

This completes the proof because, as already pointed out, for sufficiently large I the structures $\sum_I^E(\mathcal{A}'_1)$ and $\sum_I^E(\mathcal{A}'_2)$ are now $\Phi_0(\underline{c}_0)$-isomorphic. $\square$

6 Conclusions

In this paper we introduced a type-theoretic machinery in order to deal with the combination of decision problems of various nature. Higher order type theory has been essentially used as a *unifying specification language*. We have also seen how the *types interplay* can be used in a rather subtle way to design combined fragments and consequently appropriate decision problems.

Decision problems are at the heart of logic and of its applications, that's why they are so complex and irregularly behaved. Given that it is very difficult (and presumably impossible) to get satisfying general results in this area, the emphasis should concentrate on *methodologies* which are capable of solving entiere classes of concrete problems. Among methodologies, we can certainly include *methodologies for combination*: these may be very helpful when the solution of a problem can be modularly decomposed or when the problem itself appears to be heterogeneous in its nature.

In this paper, we took into consideration *Nelson-Oppen methodology* (which is probably the simplest combination methodology) and tried to push it as far as possible. Surprisingly, it

[86] One may also use the complete formulation of the descent theorem here, saying that the category of sets is equivalent to the category of descent data for the non empty set W, and realize that the above definition is just the definition of an I-indexed coproduct.

turned out that it might be quite powerful, when *joined to strong model theoretic results* (the isomorphism theorems). Thus, we tried to give the reader a gallery of different applications that can be solved in a *uniform way* by this methodology. Some of these applications are new, some other summarize recent work by various people. Each of the concrete applications we found can probably be strengthened: in particular, we only partially cared, case by case, to perfectly match all the existing results and we certainly did not attain the strongest possible statements which might have been at hand in the new applications. This was not the point, as our main emphasis was on the development of a general methodology.

New problems certainly arise now: they concern both further applications of Nelson-Oppen schema and the individuation or more sophisticated schemata, for the problems that cannot be covered by the Nelson-Oppen approach. We hope that the higher order framework and the model theoretic techniques we introduced in this paper may give further contributions within this research perspective.

Aknowledgments

We thank Manfred Schmidt-Schauss for comments on the preliminary version of this report.

References

[1] H. Andréka, I. Nemeti, and J. VanBenthem. Modal languages and bounded fragments of predicate logics. *Journal of Philosophical Logic*, 27:217–274, 1998.

[2] Peter B. Andrews. *An introduction to mathematical logic and type theory: to truth through proof*, volume 27 of *Applied Logic Series*. Kluwer Academic Publishers, Dordrecht, second edition, 2002. ISBN 1-4020-0763-9.

[3] Peter B. Andrews. Classical type theory. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume II, pages 966–1007. Elsevier/MIT, 2001.

[4] Franz Baader and Silvio Ghilardi. Connecting many-sorted theories. In *Proceedings of the 20th International Conference on Automated Deduction (CADE-05)*, Lecture Notes in Artificial Intelligence, Tallinn (Estonia), 2005. Springer-Verlag.

[5] Franz Baader and Silvio Ghilardi. Connecting many-sorted structures and theories through adjoint functions. In *Proceedings of the 5th International Workshop on Frontiers of Combining Systems (FROCOS-05)*, Lecture Notes in Artificial Intelligence, Wien (Austria), 2005. Springer-Verlag.

[6] Franz Baader and Tobias Nipkow. *Term rewriting and all that*. Cambridge University Press, Cambridge, 1998. ISBN 0-521-45520-0; 0-521-77920-0.

[7] Franz Baader and Cesare Tinelli. A new approach for combining decision procedures for the word problem, and its connection to the Nelson-Oppen combination method. In W. McCune, editor, *Proceedings of the 14th International Conference on Automated Deduction (Townsville, Australia)*, volume 1249 of *Lecture Notes in Artificial Intelligence*, pages 19–33. Springer-Verlag, 1997.

[8] Franz Baader and Cesare Tinelli. Deciding the word problem in the union of equational theories. *Information and Computation*, 178(2):346–390, December 2002.

[9] Franz Baader, Silvio Ghilardi, and Cesare Tinelli. A new combination procedure for the word problem that generalizes fusion decidability results in modal logics. *Information and Computation*. (to appear).

[10] Franz Baader, Carsten Lutz, Holger Sturm, and Frank Wolter. Fusions of description logics and abstract description systems. *Journal of Artificial Intelligence Research*, 16: 1–58, 2002.

[11] Franz Baader, Silvio Ghilardi, and Cesare Tinelli. A new combination procedure for the word problem that generalizes fusion decidability results in modal logics. In *Proceedings of the Second International Joint Conference on Automated Reasoning (IJCAR'04)*, volume 3097 of *Lecture Notes in Artificial Intelligence*, pages 183–197, Cork (Ireland), 2004. Springer-Verlag.

[12] Leo Bachmair, Harald Ganzinger, and Uwe Waldmann. Theorem proving for hierarchic first-order theories. In *Algebraic and logic programming (Volterra, 1992)*, volume 632 of *Lecture Notes in Comput. Sci.*, pages 420–434. Springer, Berlin, 1992.

[13] H. P. Barendregt. *The lambda calculus*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, revised edition, 1984. ISBN 0-444-86748-1; 0-444-87508-5. Its syntax and semantics.

[14] Peter Baumgartner, Ulrich Furbach, and Uwe Petermann. A unified approach to theory reasoning. Research Report 15–92, Universität Koblenz-Landau, Koblenz, Germany, 1992. URL `http://www.uni-koblenz.de/fb4/publikationen/gelbereihe/RR-15-92.ps.gz`. Fachberichte Informatik.

[15] Egon Börger, Erich Grädel, and Yuri Gurevich. *The classical decision problem*. Universitext. Springer-Verlag, Berlin, 2001. ISBN 3-540-42324-9. Reprint of the 1997 original.

[16] Torben Bräuner and Silvio Ghilardi. First-order modal logic. In Johan van Benthem, Patrick Blackburn, and Frank Wolter, editors, *Handbook of Modal Logic*. 2005. (to appear).

[17] Chen-Chung Chang and H. Jerome Keisler. *Model Theory*. North-Holland, Amsterdam-London, IIIrd edition, 1990.

[18] Hans de Nivelle and Ian Pratt-Hartmann. A resolution-based decision procedure for the two-variable fragment with equality. In *Automated reasoning*, volume 2083 of *Lecture Notes in Comput. Sci.*, pages 211–225. Springer, Berlin, 2001.

[19] G. Dowek. Higher order unification and matching. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume II, pages 1009–1062. Elsevier/MIT, 2001.

[20] C. Fermüller, A. Leitsch, T. Tammet, and N. Zamov. *Resolution methods for the decision problem*, volume 679 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1993. ISBN 3-540-56732-1. Lecture Notes in Artificial Intelligence.

[21] Camillo Fiorentini and Silvio Ghilardi. Combining word problems through rewriting in categories with products. *Theoretical Computer Science*, 294:103–149, 2003.

[22] Harvey Friedman. Equality between functionals. In *Logic Colloquium (Boston, Mass., 1972-1973)*, pages 22–37. Lecture Notes in Math., Vol. 453. Springer, Berlin, 1975.

[23] D. M. Gabbay and V.B. Shehtman. Undecidability of modal and intermediate first-order logics with two individual variables. *Journal of Symbolic Logic*, 58:800–823, 1993.

[24] D.M. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyaschev. *Many-Dimensional Modal Logics: Theory and Applications*, volume 148 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 2003.

[25] Harald Ganzinger and Hans de Nivelle. A superposition decision procedure for the guarded fragment with equality. In *14th Symposium on Logic in Computer Science (Trento, 1999)*, pages 295–303. IEEE Computer Soc., Los Alamitos, CA, 1999.

[26] Harald Ganzinger, Harald Rueß, and Natarajan Shankar. Modularity and refinement in inference systems. In *Notes of the Third Workshop on Pragmatics of Decision Procedures in Automated Reasoning (PDPAR'05)*, Edinburgh (Scotland), July 2005.

[27] Silvio Ghilardi. Quantifier elimination and provers integration. In Ingo Dahn and Laurent Vigneron, editors, *Proceedings of the 4th International Workshop on First-Order Theorem*

Proving, FTP 2003, Valencia, Spain, 2003. Published as Volume 86, Issue 1, of the Electronic Notes in Theoretical Computer Science.

[28] Silvio Ghilardi. Model theoretic methods in combined constraint satisfiability. *Journal of Automated Reasoning*, 33(3-3):221–249, 2005.

[29] Silvio Ghilardi and Luigi Santocanale. Algebraic and model theoretic techniques for fusion decidability in modal logics. In Moshe Vardi and Andrei Voronkov, editors, *Proceedings of the 10th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2003)*, volume 2850 of *Lecture Notes in Computer Science*, pages 152–166. Springer-Verlag, 2003.

[30] Silvio Ghilardi, Enrica Nicolini, and Daniele Zucchelli. A comprehensive framework for combined decision procedures. In *Proceedings of the 5th International Workshop on Frontiers of Combining Systems (FROCOS-05)*, Lecture Notes in Artificial Intelligence, Wien (Austria), 2005. Springer-Verlag.

[31] Jean-Yves Girard, Paul Taylor, and Yves Lafont. *Proofs and types*, volume 7 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, 1989. ISBN 0-521-37181-3.

[32] Valentin Goranko and Martin Otto. Model theory of modal logic. In Johan van Benthem, Patrick Blackburn, and Frank Wolter, editors, *Handbook of Modal Logic*. 2005. (to appear).

[33] E. Grädel. Decision procedures for guarded logics. In *Proceedings of CADE-16*, volume 1632 of *Lecture Notes in Computer Science*, pages 31–51. Springer, 1999.

[34] George Janelidze and Walter Tholen. Facets of descent. I. *Appl. Categ. Structures*, 2(3): 245–281, 1994. ISSN 0927-2852.

[35] André Joyal and Myles Tierney. An extension of the Galois theory of Grothendieck. *Mem. Amer. Math. Soc.*, 51(309):vii+71, 1984. ISSN 0065-9266.

[36] Emanuel Kieronski and Martin Otto. Small substructures and decidability issues for firstorder logic with two variables. In *Twentieth Annual IEEE Symposium on Logic in Computer Science (LICS 2005)*. IEEE Computer Soc., Los Alamitos, CA, 2005. (in print).

[37] R. Kontchakov, A. Kurucz, and M. Zakharyaschev. Undecidability of first-order intuitionistic and modal logics with two variables. 2004. manuscript.

[38] S. Kripke. The undecidability of monadic modal quantificational theory. *Z. Math. Logik Grundlag. Math.*, 8:113–116, 1962.

[39] J. Lambek and P. J. Scott. *Introduction to higher order categorical logic*, volume 7 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1988. ISBN 0-521-35653-9. Reprint of the 1986 original.

[40] F. William Lawvere. Functorial semantics of algebraic theories. *Proc. Nat. Acad. Sci. U.S.A.*, 50:869–872, 1963.

[41] Philippe Le Chenadec. *Canonical forms in finitely presented algebras*. Research Notes in Theoretical Computer Science. Pitman Publishing Ltd., London, 1986. ISBN 0-273-08721-5.

[42] D. Lewis. Counterpart theory and quantified modal logic. *Journal of Philosophy*, (65): 113–126, 1968.

[43] L. Löwhenheim. Über Möglichkeiten im Relativkalkül. *Math. Annalen*, 76:228–251, 1915.

[44] Saunders MacLane and Garrett Birckhoff. *Algebra*. MacMillan, New York, 1965.

[45] Michael Makkai and Gonzalo E. Reyes. *First-Order Categorical Logic*, volume 611 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1977.

[46] M. Marx. Tolerance logic. *Journal of Logic, Language and Information*, 10:353–374, 2001.

[47] M. Mortimer. On languages with two variables. *Z. Math. Logik Grundlag. Math.*, 21: 135–140, 1975.

[48] Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. on Programming Languages and Systems*, 1(2):245–257, October 1979.

[49] Piergiorgio Odifreddi. *Classical recursion theory*, volume 125 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, 1989. ISBN 0-444-87295-7. The theory of functions and sets of natural numbers, With a foreword by G. E. Sacks.

[50] Derek C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12:291–302, 1980.

[51] Michael O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Math. Soc.*, 141:1–35, 1969. ISSN 0002-9947.

[52] D. Scott. A decision method for for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.

[53] K. Segerberg. Two-dimensional modal logic. *Journal of Philosophical Logic*, 2:77–96, 1973.

[54] V.B. Shehtman. On some two-dimensional modal logics. In *8th Congress on Logic Methodology and Philosophy of Science, vol. 1*, pages 326–330. Nauka, Moskow, 1987.

[55] Cesare Tinelli. Cooperation of background reasoners in theory reasoning by residue sharing. *Journal of Automated Reasoning*, 30(1):1–31, January 2003.

[56] Cesare Tinelli and Mehdi T. Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In F. Baader and K.U. Schulz, editors, *Frontiers of Combining Systems: Proceedings of the 1st International Workshop (Munich, Germany)*, Applied Logic, pages 103–120. Kluwer Academic Publishers, March 1996.

[57] Johan van Benthem. *Modal logic and classical logic.* Indices: Monographs in Philosophical Logic and Formal Linguistics, III. Bibliopolis, Naples, 1985. ISBN 88-7088-113-X.

[58] F. Wolter and M. Zakharyaschev. Decidable fragments of first-order modal logics. *Journal of Symbolic Logic*, 66:1415–1438, 2001.

[59] Frank Wolter. Fusions of modal logics revisited. In M. Kracht, M. de Rijke, H. Wansing, and M. Zakharyaschev, editors, *Advances in Modal Logic*. CSLI, Stanford, CA, 1998.

Index