



HAL
open science

Namely, Painless: A unifying approach to safe programming with first-order syntax with binders

Nicolas Pouillard

► **To cite this version:**

Nicolas Pouillard. Namely, Painless: A unifying approach to safe programming with first-order syntax with binders. Programming Languages [cs.PL]. Université Paris-Diderot - Paris VII, 2012. English. NNT: . tel-00759059

HAL Id: tel-00759059

<https://theses.hal.science/tel-00759059>

Submitted on 1 Dec 2012

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

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

UNIVERSITÉ PARIS DIDEROT (PARIS 7)
ÉCOLE DOCTORALE : SCIENCES MATHÉMATIQUES DE
PARIS CENTRE

DOCTORAT
INFORMATIQUE

**Une approche unifiante pour programmer
sûrement avec de la syntaxe du premier
ordre contenant des lieurs**

english title:
Namely, Painless

**A unifying approach to safe programming with
first-order syntax with binders**

Nicolas Pouillard

Thèse dirigée par François Pottier

et soutenue le 13 Janvier 2012 devant le Jury composé de :

Président	M. Roberto Di Cosmo
Rapporteurs	M. Andrew Pitts M. Dale Miller
Examineurs	M. Daniel Hirschhoff M. Conor McBride
Directeur	M. François Pottier

À mon papa...

Abstract

This dissertation describes a novel approach to safe meta-programming. A meta-program is a program which processes programs or similar data. Compilers and theorem provers are prime examples of meta-programs which could benefit from this approach. To this end, this work focuses on the representation of names and binders in data structures.

Programming errors are really easy to make with usual techniques. We propose an abstract interface to names and binders that rules out these errors. This interface is implemented as a library in Agda. It allows defining and manipulating term representations in nominal style. Thanks to abstraction, other styles are supported as well: the de Bruijn style, the combinations of these styles, and more.

Whereas indexing the types of names and terms with a natural number is a well-known technique to better control de Bruijn indices, we index them with worlds. Worlds are at the same time more precise and more abstract than natural numbers. Via logical relations and parametricity, we are able to demonstrate in what sense our library is safe, and to obtain theorems for free about world-polymorphic functions. For instance, we prove that a world-polymorphic term transformation function must commute with any renaming of the free variables. The proof is entirely carried out in Agda.

The usability of our technique is shown on several examples including normalization by evaluation which is known to be challenging. We show that our world-indexed approach can express a wide range of data types by embedding several definition languages from the literature.

Résumé

Cette thèse décrit une nouvelle approche pour la méta-programmation sûre. Un méta-programme est un programme qui manipule des programmes ou assimilés. Les compilateurs et systèmes de preuves sont de bons exemples de méta-programmes qui bénéficieraient de cette approche. Dans ce but, ce travail se concentre sur la représentation des noms et des lieux dans les structures de données.

Les erreurs de programmation étant courantes avec les techniques usuelles, nous proposons une interface abstraite pour les noms et les lieux qui élimine ces erreurs. Cette interface est implémentée sous forme d’une bibliothèque en Agda. Elle permet de définir et manipuler des représentations de termes dans le style nominal. Grâce à l’abstraction, d’autres styles sont aussi disponibles : le style de De Bruijn, les combinaisons de ces styles, et d’autres encore.

Nous indiquons les noms et les termes par des mondes. Les mondes sont en même temps précis et abstraits. Via les relations logiques et la paramétricité, nous pouvons démontrer dans quel sens notre bibliothèque est sûre, et obtenir des “théorèmes gratuits” à propos des fonctions monde-polymorphiques. Ainsi une fonction monde-polymorphique de transformation de termes doit commuter avec n’importe quel renommage des variables libres. La preuve est entièrement conduite en Agda.

Notre technique se montre utile sur plusieurs exemples, dont la normalisation par évaluation qui est connue pour être un défi. Nous montrons que notre approche indiquée par des mondes permet d’exprimer un large panel de type de données grâce à des langages de définition embarqués.

Remerciements

Trois années de noms et de termes. Des liens puis des noms, des noms et des indices, puis des indices seulement, puis finalement encore des noms. Aujourd'hui j'écris ces remerciements et ce sont vos noms qui me viennent à l'esprit.

Je porte mes remerciements tout d'abord à ma famille, mes parents, pour m'avoir soutenu dans mes longues études. Merci aussi à mon épouse, Gaëlle, qui m'a non seulement soutenu mais épaulé, écouté, poussé, aidé, ... Bref, elle est aussi responsable de cet aboutissement.

Je tiens particulièrement à remercier mes relecteurs pour leur travail remarquable. Leurs relectures de ce document m'ont permis de nombreuses améliorations rendant le présent document plus accessible et plus complet.

Je remercie le LRDE qui m'a aiguillé vers la recherche, je remercie en particulier Akim Demaille pour son cours de compilation qui m'a initié à la conception de langages.

Je remercie d'avance tous ceux que je ne mentionne pas individuellement ici : membres de ma famille, camarades, collègues, ami(e)s, simples connaissances. Vous qui justement lisez ces lignes, je vous en remercie.

Un grand merci à toute l'équipe Gallium de l'INRIA Rocquencourt et tous les membres que j'y ai rencontrés. Cette équipe m'a accueilli jeune ingénieur et m'a rendu jeune chercheur. Plus particulièrement je remercie : Michel Mauny pour m'avoir accepté en stage mais aussi sur le contrat qui a suivi ; j'y ai appris une foule d'anecdotes sur l'évolution des dialectes successifs de Caml ; grâce à lui et malgré lui j'ai appris à apprécier la programmation paresseuse (sans jeux de mots) ce qui m'a ouvert de nouveaux horizons. Puis c'est au tour de François Pottier, qui m'a accepté en thèse et m'a habilement guidé et soutenu pendant ces années. Il a su être disponible et à l'écoute. Outre une montagne de connaissances scientifiques, j'ai appris de lui l'art d'écrire des articles bien que j'ai encore beaucoup à apprendre à ce sujet. Xavier Leroy pour son accueil et sa bienveillance quant à la bonne marche de l'équipe ; Damien Doligez pour son intarissable culture et ses théories sur tout ; Didier Rémy pour ses discussions sur les détails du fonctionnement de $\text{T}_{\text{E}}\text{X}$:) ; Alain Frisch pour m'avoir enseigné tellement de choses sur OCaml et la programmation fonctionnelle en général pendant les mois durant lesquels j'ai partagé son bureau ; Yann Régis-Gianas pour avoir tracé

une voie qu'il m'a suffit d'emprunter ; Berke Durak pour m'avoir accompagné dans le développement d'ocamlbuild ; Benoit Razet pour toutes nos discussions sur les automates et les structures de flots de données ; Gérard Huet pour ses anecdotes variées qui m'ont beaucoup appris ; Jean-Baptiste Tristan pour avoir participé et accueilli chez lui l'équipe des "jeunes" de Gallium au concours ICFP ; Tous ceux qui ont participé aux concours ICFP ; Arthur Charguéraud pour son secret de l'efficacité : les axiomes COQ et [notepad++](#) :) ; Zaynah Dargaye, Boris Yakobowski et Paolo Herms pour les discussions et les bons moments passés avec eux ; Benoît Montagu pour m'avoir ouvert son bureau et avoir partagé des discussions fructueuses sur nos thèses respectives ; Julien Cretin pour sa singularité mais aussi pour tous les sujets sur lesquels nous pouvons discuter agréablement ; Alexandre Pilkiewicz de m'avoir entraîné jusqu'au parcours sportif un certain nombre de fois ; Jonathan Protzenko car grâce à lui le web évolue de jour en jour :) ; Tahina Ramananandro pour nos discussions politiques et nos concours de jeux de mots fumeux ; Dana Xu pour ses encouragements ; Gabriel Scherer et Valentin Robert pour leurs nombreux commentaires sur des parties de ce document.

À Luc Maranget pour les discussions nautiques et ses slogans devenus cultes : "Les programmeurs veulent des types plus riches!", "Les types existentiels existent-ils?", "Ils nous parlent seulement des types, où sont les termes?" ; Jean-Jacques Levy pour son talent à poser tout haut les questions auxquelles tout le monde pense.

J'ai une pensée aussi pour tous les "camarades de navette" avec qui j'ai pu discuter pendant un des nombreux trajets entre Paris et Rocquencourt.

À Jean-Philippe Bernardy, pour avoir rendu la paramétricité accessible et claire dans un cadre simple et élégant. Ses travaux ont amplement aidé les miens en apportant l'outil de preuve adapté à ce problème.

Contents

1	Introduction	13
1.1	Programs and Programming Languages	14
1.2	Program syntax	15
1.3	Typing	17
1.4	Program representation, meta-programming	19
1.5	Names and local bindings	23
1.6	Scoping	25
1.7	Empowering our language	28
2	The nominal approach	41
2.1	Introduction to the nominal approach	41
2.1.1	Warm-up: the bare nominal approach	41
2.1.2	Using well-formedness judgements	44
2.1.3	Well-scoped terms	47
2.2	NOMPA: A programming interface in nominal style	48
2.2.1	All we need to define nominal syntax	50
2.2.2	Building binders and names	51
2.2.3	World widening: Name weakening	53
2.2.4	Comparing and refining names	56
3	Programming on top of NomPa	59
3.1	Various examples	59
3.1.1	Example: computing free variables	59
3.1.2	Example: working with environments	60
3.1.3	Example: Term comparison	62
3.2	Kits and Traversals	63
3.2.1	Traversal Kits	63
3.2.2	Coercing kit	64
3.2.3	Renaming kits	65
3.2.4	Substitution kit	67
3.2.5	Other kits and combinators	68
3.2.6	Reusable traversal	68
3.2.7	Reusing the traversal	69

3.2.8	Building any λ -term	71
3.3	Towards elaborate uses of worlds	72
3.3.1	Data type of “one hole contexts”	72
3.3.2	Patterns a la ML	73
3.3.3	Term and types: System F	73
3.4	Advanced example: normalization by evaluation	73
4	Behind the scene of the NomPa library	77
4.1	Implementation of NOMPA	77
4.2	Soundness: logical relations and parametricity	80
4.2.1	Recap of the framework	80
4.2.2	An example: Boolean values represented by numbers	85
4.3	Relations for NOMPA	87
4.3.1	Relations for NOMPA types	88
4.3.2	NOMPA values fit the relation	91
4.3.3	An example not fitting the relation	92
5	The de Bruijn approach	93
5.1	Introduction to de Bruijn indices	93
5.1.1	bare : The original approach	94
5.1.2	Maybe : The nested data type approach	94
5.1.3	The Fin approach	95
5.2	An interface for de Bruijn indices	96
5.2.1	Singleton worlds!	101
5.3	Examples and advanced operations	101
5.3.1	Some convenience functions	101
5.3.2	Building terms	104
5.3.3	Computing free variables	104
5.3.4	Generic traversal	105
5.3.5	Nameless term comparison	108
5.4	Logical Relations for de Bruijn indices	109
5.4.1	Relations for the de Bruijn world operations	109
5.4.2	De Bruijn functions fit the relation	110
5.4.3	On the strength of free theorems	111
5.4.4	Using logical relations and parametricity	114
6	Variations and Related Work	119
6.1	Dynamically stratified representations	120
6.1.1	Locally closed terms	121
6.1.2	Building terms	122
6.1.3	Performances	123
6.1.4	Non-structural recursion	124
6.1.5	Free atoms using <code>openTm</code>	125
6.1.6	Common parts of the internal representation	126

6.1.7	Locally Nameless	126
6.1.8	Locally Named	128
6.1.9	When does this stratified technique pay off?	130
6.2	de Bruijn levels	133
6.2.1	Term comparison	133
6.2.2	Kits and traversals	134
6.2.3	Derived functions	136
6.2.4	New inclusion rules	137
6.2.5	Addition and subtraction kits	137
6.2.6	Conversion with nominal terms	138
6.3	Links: Binders as World Relations	139
6.3.1	Terms and examples	140
6.3.2	Implementations	141
6.3.3	Building links and terms	143
6.3.4	Kits and traversals	144
6.4	Combining nominal and nameless styles	148
6.5	Nominal types	151
6.5.1	FRESHML & FRESH OCAML types	151
6.5.2	Caml types	159
6.5.3	Connection with BINDERS UNBOUND	163
6.6	Related work	167
7	Conclusion	175
7.1	Contributions	175
7.2	Future work	176
	Bibliography	177
	List of Figures	183
	Index	185

Chapter 1

Introduction

Foreword

This document describes my work during these past three years which has been done under the direction of François Pottier. This work extends the research topic of safe meta-programming by providing a novel approach to the representation of programs with names and binders. – Nicolas Pouillard

Safe meta-programming is a research topic which is part of the broader topic of programming language design. One of the goals of the research in programming language design is to provide programmers with better tools allowing to produce correct software applications in a reasonable amount of time. Since correct software applications (programs for short) are not supposed to crash, language designs have to incorporate safety features to detect programming errors as soon as possible. Modern language designs often come with some static disciplines such as lexical scoping, strong typing and the like. These static disciplines are used to reject the programs we do not want to run.

A program may process various forms of data such as numbers, text, spreadsheets, images, sound, and many more. Some programs process other programs, they are called meta-programs. Meta-programs cover a wide class of programs ranging from compilers, to static analyzers, code generators, etc. Proof systems, proof assistants and theorem provers are also meta-programs since propositions and proofs share the same kind of structure as programs. Moreover since language designs come with static disciplines, we claim that meta-programs should follow these disciplines when processing programs. For instance if a program **A** written in a programming language **P** generates a program **B** written in a programming language **Q**, we want **A** to follow the static discipline of **P**, moreover we want **B** which does not exist yet to follow the static discipline of **Q**.

This first chapter aims at introducing all the notions to understand

the research problem that we are focusing on. This chapter requires little knowledge in the design of programming languages, but does require some knowledge about computers, mathematics and programming. Hence this introduction chapter can be safely skipped by readers accustomed with the research topic.

1.1 Programs and Programming Languages

To introduce the notions of programs and programming languages, let us start with an example of a simple program:

```
print "Hello! 2 times 21 is equal to " >>
print (show (2 * 21))
```

Most generally, programs are written to be executed on a computer. Running a program often results in the execution of different tasks and yields a final result. In our previous example the program is meant to print on the screen: “Hello! 2 times 21 is equal to 42”. Printing on the screen is done with the construct called `print`. This printing construct is used twice. The two uses are *sequenced* using the construct `>>`, which will perform the action on the left and then the action on the right. The number `42` is computed from the arithmetic expression `2 * 21` which is then turned into a printable form using the construct called `show`.

This program, such as any program, is written in a programming language. Programming languages are described by a finite set of constructs, rules explaining how the constructs can be used and rules giving the behavior of programs. In our further examples the programming language contains constructs such as `print`, `>>`, `show`, `*`, character strings such as `"Hello! 2 times 21 is equal to "`, numbers such as `2` and `21`, and parentheses. With these constructs our programming language is not much better than a simple calculator.

This tiny introductory programming language is built to be a subset of the AGDA language. AGDA [Norell, 2007] is a total functional programming language, based on dependent type theory. AGDA is used thoroughly in this document. First it is used to introduce the research topic of meta-programming. Second it is used to implement our solution to the problem. Third it is used to formalize our system and show its soundness. Alas, advanced languages such as AGDA are far from simple to completely describe. The first set of programming constructs we introduce is roughly common to most languages and is described simply in great detail. The second set of constructs follows the development of functional programming languages (such as HASKELL, ML, or LISP). They are described more quickly and may require further reading about these topics. An excellent introduction

to functional programming in HASKELL is Graham Hutton’s “Programming in Haskell” [Hutton, 2007]. Moreover the third set of constructs is specific to dependently typed and total functional languages or even restricted to AGDA only. To fully grasp these constructs we recommend further reading on these topics. In particular Ulf Norell’s [2007] PhD thesis on the theoretical and practical development of AGDA is recommended.

1.2 Program syntax

Examples The following program makes use of the four basic arithmetic operators to compute the value 42. Operators are attributed with their usual semantics, addition is `+` and multiplication is `*`. Subtraction and division (`-` and `÷`) are restricted to natural numbers, meaning that they return 0 when going out of range.

```
(1 + 1) * (28 ÷ 14 ÷ 2)
```

The following program illustrates character string literals and their concatenation with `++`. This program prints “Hello world!” on the screen:

```
print ("Hello" ++ " " ++ "world!")
```

Our next program illustrates the use of `show` to turn a number into a character string containing its decimal representation. This program prints “show 1 renders as: 1” on the screen:

```
print ("show 1 renders as: " ++ show 1)
```

The following program illustrates the use of `>>` to sequence two actions from left to right. This program prints “Print me first.” and then prints “Then print me.” on the screen.

```
print "Print me first." >> print "Then print me."
```


We have lists in our programming language. The simplest list value is the empty list, written `[]`. The following program is a list with three elements, namely `"a"`, `"bc"` and `"5"`:

```
[ "a" , "b" ++ "c" , show (2 + 3) ]
```

We go from lists to trees by using the construct called `node`. The first argument of `node` is the label of the node and the second is the list of children of the node. The following program represents a six-node tree

depicted hereafter:

```
node "A"
  [ node "B" []
  , node "C" [ node "D" [] , node "E" [] ]
  , node "F" [] ]
```



Syntax of our arithmetic calculator The syntax of a programming language describes when a program has a valid “shape”. In particular it does not check if the program has a valid “meaning”. For instance, the program `1 + "Hello!"` has a valid syntax but no valid meaning in our tiny language.

Any non-empty sequence of digits (0, 1, 2, 3, 4, 5, 6, 7, 8, 9) is a valid literal number and a valid program as well. A valid literal character string is a sequence of characters different from the double quote character (‘”’) and surrounded by two double quote characters (we avoid talking about escape sequences because we do not need them here). If two programs `p1` and `p2` are syntactically valid, then for any operator `•` among `+`, `-`, `*`, `/`, `++`, `>>` the program `p1 • p2` is syntactically valid. These operators are said to be *infix*, meaning that they appear between the two operands. If `p` is a syntactically valid program, then `print p`, `show p` and `(p)` are syntactically valid. If `p1`, `p2`, ..., `pn` are syntactically valid programs, then `[p1 , p2 , ... , pn]` is syntactically valid. Finally if `p1` and `p2` are syntactically valid programs then `node p1 p2` is syntactically valid.

Of course not every textual program is syntactically valid. Consider this example: `3 + * 4`. This program does not follow the rules we just described, hence it is rejected as syntactically invalid.

Syntax also comprises some details about layout and comments. Spaces and newlines are important to some extent. Spaces or parentheses have to be used to separate the words in a program. However, as long as they are separated, how many spaces or newlines are used to separate two words does not matter. There is special support in the syntax to embed comments in our programs. There are two syntaxes used in our language. A comment can start with `{-` and stop with `-}` or can start with `--` and stop at the

end of line. Here is a syntactically valid program to illustrate layout and comments:

```
print "Hello.." >> -- this prints Hello...
print  {- extra spaces and comments are ok -} "... world!"
```

1.3 Typing

Typing captures, ahead of time, programs that may go wrong. Typing might be seen as a vigilant companion giving advice on our programs. When typing says that our program is ill-typed, then this is a program we may not want to even try running. If the typing says our program is fine, then we know for sure that the program will not crash.

Such a technology is to be applied systematically before running a program. It is an affordable safety measure which massively helps program development, maintenance, testing, and verification.

However, verifying that a program cannot crash without running it is a difficult problem to say the least. In particular we want the typing part of our system to always terminate, even if the given program does not terminate. We not only want the typing to terminate but to answer if the program is well-typed or not. If so we say that such a type system is decidable.

A type system is said to be *sound* if and only if every well-typed program is a non-crashing program. We expect every well-designed type system to be sound.

A type system is said to be *complete* if and only if every non-crashing program is well-typed. The decidability constraint entails that a type system will never be complete for non-trivial languages. For further reading on the topic of type systems we recommend reading “Types and Programming Languages” [Pierce, 2002] and “Advanced Topics in Types and Programming Languages” [Pierce, 2005].

Typing our arithmetic language The first goal of type systems is to prevent us from running programs that may crash. To do so, they define *types*, an abstraction to classify pieces of data. By extension, types also classify programs computing pieces of data. Types help to define a common “contract” for producers and consumers of a piece of data. This kind of “contracts” in on the shape of data. The more shapes are allowed by a type, the weaker the type is. The strength of a type affects producers and consumers in different ways. If a type is weak, it is easy for producers of this type to be well-typed. If a type is strong, the producers have to meet all the constraints imposed by the type. This is the opposite for consumers. If a type is strong, the data has a known precise shape, and the consumer

has an easy job. If a type is weak, then the consumer has to consider all the shapes that the data can take.

To illustrate the notion of typing, we equip our little programming language with such a type discipline. Step by step we introduce new types and how the constructs of our language are typed.

We have only a few types of values. We introduce the type \mathbb{N} as the type of numbers (0, 1, 2, ...) and of programs computing numbers (such as $2 * 21$).

Given a type τ (such as \mathbb{N}) and a program p (such as $6 * 7$), we write p has type τ to formally assert that the program p is of type τ . We call these type assertions *typing judgements*. The following paragraphs define the type system for our tiny language by giving a typing judgement for each construct.

For each number n , n has type \mathbb{N} . Given two programs p_1 and p_2 , if both have type \mathbb{N} , then $p_1 + p_2$, $p_1 * p_2$, $p_1 \div p_2$, and $p_1 \div p_2$ all have type \mathbb{N} as well. For instance with these rules one can formally assert: $6 * 7$ has type \mathbb{N} .

Our second type of basic values is text values. To this end we introduce the type `String` as the type of text values such as "Hello!" and of programs computing text values such as "Hello " ++ show 42. More precisely for each syntactically valid text value s , s has type `String`. If two programs p_1 and p_2 have type `String` then the program $p_1 ++ p_2$ has type `String` as well. If a program p has type \mathbb{N} then `show p` has type `String`.

It is now time to show an ill-typed program: "Hello!" ++ 42. While this program is syntactically valid, it does not follow our typing rules. There is indeed only one rule about the construct ++. This rule imposes programs on each side of ++ to be of type `String`. While "Hello!" is of type `String`, 42 is not. The only rule that applies to 42 says that it is of type \mathbb{N} . Since `String` and \mathbb{N} are different, this program is rejected as being ill-typed.

Our little programming language has list values. However there is no single type for all lists. The type of a list depends on the type of its elements. We require all the elements of a list to have the same type. Given any type A , if p_1, p_2, \dots, p_n are programs of type A then $[p_1 , p_2 , \dots , p_n]$ is of type `List A`. Thus there is no single type for lists, but for each type A there is a type for lists where elements are of type A , namely `List A`. For instance $[1 , 2]$ has type `List \mathbb{N}` , and $[[1 , 2] , [3]]$ has type `List (List \mathbb{N})`.

Our next type is the type of trees. Like for lists, the type of trees is parameterized by the type of its elements, namely the labels of nodes. For each type A , `Tree A` is the type of trees with labels of type A . To build a tree we use the construct `node`: given a type A , if p_1 has type A and p_2 has type `List (Tree A)` then `node p1 p2` has type `Tree A`.

The following program has type `Tree String`:

```
node "A" [ node "B" [] , node "C" [] ]
```

This one has type `Tree N`:

```
node 1 [ node 2 [] , node 3 [] ]
```

This one is ill-typed, though:

```
node 1 [ node "B" [] , node "C" [] ]
```

Finally the type `Interactive` is used for programs interacting with the environment. We currently introduced only two forms of interaction, namely `print` and `>>`. Given a program `p` of type `String`, `print p` has type `Interactive`. This interactive program computes an interactive value which is meant to be triggered. Once the interactive value is triggered it prints the value of the text `p` on the screen.

The second form of interaction enables to sequence two interactive values. Given two programs `p1` and `p2` of type `Interactive` the program `p1 >> p2` has type `Interactive` as well.

1.4 Program representation, meta-programming

When reading, writing, and editing programs we often work with the textual form of programs. We write programs in text files and so their representation is a sequence of characters in a file.

Machines do not directly accept this kind of program for execution. Even languages very close to the hardware such as assembly languages are not directly understood by the machine. Hence programs are processed by other programs. These programs processing other programs are called meta-programs and are described in greater detail in the following section.

Up to here our simple programming language can process numbers, text, lists, trees and interactive values. What about programs? What is required to process programs themselves? If a program is simply a text in a file then our text value can represent programs. Here are two programs, the second is the textual representation of the first:

```
2 * (10 + 11) {- this should compute 42 -}
```


```
"2 * (10 + 11) {- this should compute 42 -}"
```

This is common knowledge nowadays that the textual representation is not adapted for any non-trivial processing. A first step called *lexing* gets rid of the lexical issues of layout and comments. The lexing step turns

the program text to a list of words, called *tokens*. These tokens are often annotated by a *token class* to distinguish numbers, operators, constructs. . . Here is the same program as a simple list of tokens:

```
[ "2" , "*" , "(" , "10" , "+" , "11" , ")" ]
```

This representation is still unworkable for non-trivial processing. In particular there is too little structure in a list to reflect the program structure. To solve this issue, syntax trees have been introduced. Here is our same program, graphically depicted as a tree:



We picked names for some constructs such as “num” for numbers, “text” for text values, and “list” for lists. An exception is made for parentheses which no longer add any useful information and hence are not represented in the tree.

The process of translating text or tokens into a tree is called *parsing*. We will not discuss it more.

These syntax trees are generally called *Abstract Syntax Trees* or *AST* for short. Syntax trees are often said “abstract” because they no longer depend on some details of the “concrete” syntax.

Happily we have enough constructs in our language to build trees. We can hence show another representation of the same program as a tree built in our language:


```
node "*"
  [ node "num" [ node "2" [] ]
  , node "+"
    [ node "num" [ node "10" [] ]
    , node "num" [ node "11" [] ] ] ] ]
```

Representing programs of our language as values of our language itself is good to have but not necessary. However we can expect a programming language good at processing programs to be good at representing as many programs as possible.

An example to illustrate lists:

```
[ "a" , "b" ++ "c" , show (2 + 3) ]
```

The syntax tree, graphically:




The syntax tree as a program:

```
node "list"
  [ node "text" [ node "a" [] ]
  , node "++"
    [ node "text" [ node "b" [] ]
    , node "text" [ node "c" [] ] ]
  , node "show"
    [ node "+"
      [ node "num" [ node "2" [] ]
      , node "num" [ node "3" [] ] ] ] ] ] ]
```

An example to illustrate trees:

```
node "A" [ node "B" [] , node "C" [] ]
```

The syntax tree, graphically:



The syntax tree as a program:

```
node "node"
  [ node "text" [ node "A" [] ]
  , node "list"
    [ node "node"
      [ node "text" [ node "B" [] ] , node "list" [] ]
    , node "node"
      [ node "text" [ node "C" [] ]
      , node "list" [] ] ] ]
```

Meta-programming The broad topic of “programs processing other programs” is called meta-programming. Sometimes meta-programming is given narrower definitions, but we find this one to better account for the different parts of the research field. Meta-programming comprises the generation, analysis, transformation of programs or similar objects such as formulae and proofs. Some programming languages are completely designed around meta-programming to support run-time code generation (METAML [Taha, 1999], ‘C [Engler et al., 1996]). Program translations and optimizations, as done by compilers, are prime examples of meta-programming

1.5 Names and local bindings

Our programming language, while capable of arithmetic computations, interactions, and program representation, lacks the concept of names. We

hence introduce a construct to give a name to part of a program and refer to this part using the name. This enables an important goal of re-usability. In a program we should avoid to repeat ourselves. Hence if two parts are the same they may benefit from being written only once. Maybe more importantly, changes to this part of the software logic, during later evolution of the code, only have to be done in one place instead of being duplicated. The syntax of this new construct relies on three keywords `let`, `=` and `in`. If two programs `p1` and `p2` are syntactically valid and `x` is a name then `let x = p1 in p2` is syntactically valid. Moreover, names (such as `x`) are valid program themselves. If we want to print the same text twice we can write:

```
let x = print "Hello!" in
x >> x
```

If we replace the occurrences of `x` by its definition (here `print "Hello!"`) we obtain a program with the same behavior:

```
print "Hello!" >> print "Hello!"
```

The gain brought by sharing is significant. To illustrate this fact, we build an artificial example where each added `let` would double the size of the program if we could not use `let`. Albeit artificial this example conveys the fact that sharing makes a big difference. The following example has 54 syntax nodes while the expanded version has 220 syntax nodes (only 124 if we reduce `x0` to `0` first). Note also that the small computation done in `x0` can be performed only once instead of 16 times.

```
let x0 = 42 ÷ 2 * 21 in
let x1 = node x0 [] in
let x2 = node 1 [ x1 , x1 ] in
let x3 = node 2 [ x2 , x2 ] in
let x4 = node 3 [ x3 , x3 ] in
node 4 [ x4 , x4 ]
```



```

let x = "1" in
print x >>                                -- this prints 1
let x = "2" in
print x >>                                -- this prints 2
print (let x = "3" in x) >>               -- this prints 3
(let x = "4" in print x) >>              -- this prints 4
print x >>                                -- this prints 2 again
let x = x ++ x in
print x >>                                -- this prints 22
let x = let x = "5" in x ++ x in
print x                                    -- this prints 55

```

We now introduce some vocabulary about names. In the program `let x = p1 in p2`, the name `x` is said to be *bound*. In particular it is bound in `p2`. We call the name `x` in this position a *binder*. In the program `print x`, the name `x` is said to be *free*. We also call it an *occurrence* of `x`.

Sometimes names are also called *variables*. We try to prefer the term *name* over *variable* in this document. More precisely we use the term *variable* to represent the construct which holds just a name. For instance the program `x + x` contains two variables but only one name.

A program without any free names/variables is said to be closed. The set of free names can be defined inductively as follows. The set of free names of a variable `x` is the singleton set with the name `x`. The free names of an operation such as `print p` or `show p` are the free names of `p`. The free names of an operation such as `p1 + p2` are the union of free names of `p1` and of `p2`. Finally the free names of `let x = p1 in p2` are the free names of `p2` minus `x`, union the free names of `p1`.


A name `x` is said to be *fresh* for a program `p` if the name `x` is not a member of the free variables of `p`. For example the name `y` is fresh for the program `let y = 42 in x + y`, but the name `x` is not.

Two programs `p1` and `p2` are said to be *α-equivalent* if they differ only by a consistent renaming of the bound names. Note that this is not a formal description of α-equivalence, but an informal definition appealing to our intuition of “name irrelevance”. Defining this relation precisely is one of the important part of this thesis. For instance the programs `let x = 21 in x + x` and `let y = 21 in y + y` are α-equivalent. The name of a variable should not have any importance except being a tool to reference a position in the program without ambiguities. α-equivalence is an equivalence relation. This means that the relation is reflexive (every program is α-equivalent to itself), symmetric (if `p1` is α-equivalent to `p2`, then `p2` is α-equivalent to `p1`), and transitive (if `p1` is α-equivalent to `p2` and `p2` is α-equivalent to `p3` then `p1` is α-equivalent to `p3`). The α-equivalence is also a congruence, meaning that if `p1` is α-equivalent to `p2` then if we put the programs in the *context* `C[]` to produce bigger programs then `C[p1]` is α-equivalent to `C[p2]`.

Typing our `let` construct We have introduced local definitions and variables and we now extend our typing discipline to these constructs. Like we have done for the previous constructs, we give the typing informally using a textual description. The program `let x = p1 in p2` has type σ if and only if `p1` has some type τ and that `p2` has type σ *assuming* that `x` has type τ . To type a variable we use the assumptions we gathered so far. There are several ways to precisely describe how to manage these assumptions, but we do not detail them in this introduction.

Let us try to check the typing with an example. We check that the program `let x = 42 in let y = show x in print y` has type `Interactive`. To do so we first check that `42` has type `N` which is true, then we must check that `let y = show x in print y` has type `Interactive` assuming that `x` has type `N`. To do so we first check that `show x` has type `String`, which given the rule for `show` amounts to check that `x` has type `N`, which is true by assumption. We then check that `print y` has type `Interactive` assuming that `y` has type `String`. Given the rule for `print` this amounts to check that `y` has type `String` which is true by assumption.


Meta-programming and representation of variables How do we represent variables and the `let` construct? It is reasonable to start with names being values of type `String`, each variable being a tree node that we call `var`, and the `let` construct being a node with three subtrees. The program `let x = 6 in let y = x + 1 in x * y` can thus be depicted as:



This first approach is the root of the *nominal* approach that we describe in greater detail in chapter 2.

A different approach is to use de Bruijn indices [de Bruijn, 1972]. This representation is said to be nameless because variables are no longer identified by a name but a notion of “distance” to the binding point. This nameless approach solves part of the problem by providing a canonical representation. However a major issue with this nameless representation is its

arithmetic flavor. Indeed properties about names and binders are turned into arithmetic formulae. This second approach is covered in chapter 5. Here is the same example using the de Bruijn style:



1.7 Empowering our language

Declarations and definitions We now fast-forward from our tiny subset of AGDA to AGDA itself as we need it in the remainder of this document. We start with declarations, which allow us to declare a symbol that can be used globally. This generally differs from the `let` construct whose scope is local. Moreover the declaration just specifies the type of the symbol. Subsequent phrases have to define this declared symbol. The declarations make use of the character `:` to separate the declared symbol from its type. The following line is read “Dear AGDA, let us declare `hello-world` of type `Interactive`”:

```
hello-world : Interactive
```

After the declaration must come the definition. The name of the defined symbol is recalled and a program is given for its definition:

```
hello-world = print "Hello World!"
```

Functions We introduce a new sort of types, namely function types. If σ and τ are types then $\sigma \rightarrow \tau$ is a type as well. This type represents functions whose domain is σ and co-domain is τ . Happily AGDA functions are like mathematical functions and no surprise whatsoever will trouble this. The syntax for definitions enables the definitions of functions in a very simple way: we just give a name to the argument before the equal sign. To apply a

user defined function `f` to an argument `x`, the syntax is the lightest syntax possible: `f x`.

```
double : ℕ → ℕ
double n = n + n

-- This program prints: 42
prog1 : Interactive
prog1 = print (show (double 21))
```

The name `n` before the equal sign is a binder and scopes over the sub-program found after the equal sign. The scoping works pretty much like with the `let` construct. Here is another example:

```
hello : String → Interactive
hello s = print ("Hello " ++ s ++ "!")

-- This program prints: Hello Functional World!
prog2 : Interactive
prog2 = hello "Functional World"
```

In order to receive multiple arguments one simply has to make functions return functions. Indeed the type $\mathbb{N} \rightarrow (\mathbb{N} \rightarrow \mathbb{N})$ is the type of a function taking a first number argument and returning a function taking a second number argument to finally deliver its result number. This is a pattern so common that we make the arrow type associate on right such that we can write the type this way: $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}$. The same shortcut goes for function applications. Instead of writing `(f x) y` or worse `(f(x))(y)` we simply make the application associate on left, thus we can write `f x y`. Here is an example:

```
hello2 : String → ℕ → Interactive
hello2 s n = print ("Hello " ++ s ++ " " ++ show n ++ "!")

-- This program prints: Hello World 42!
prog3 : Interactive
prog3 = hello2 "World" (2 * 21)
```

Data-types Among the types presented so far, `ℕ`, `List`, `Tree` are user-defined in AGDA. The general mechanism used to define these types is called *inductive families*. Inductive families generalize various forms of definition mechanisms such as *sum and product types*, *regular tree types*, *algebraic*

data types, and *GADTs* (generalized algebraic data-types). To define an inductive family, we need to declare its name and type, and declare as many *data constructors* as we want. Each data constructor has a name and a type as well. Let us study the definition for `Bool`, one of the simplest data type possible:

```
data Bool : Set where
  false : Bool
  true  : Bool
```

The first line declares `Bool` to have type `Set`. Indeed `Set` is a special type which is the type of basic types. AGDA treats types like other values of the language. Then we define two data constructors called `false` and `true`. These two constructors are the two only values of type `Bool`.

We now focus on the type of natural numbers, namely `N`:

```
data N : Set where
  zero : N
  suc  : N → N
```

The first line declares `N` to have type `Set`. Then we define two data constructors called `zero` and `suc`. The constructor `zero` is introduced as a value of type `N` to represent the number 0. The constructor `suc` (for successor) is introduced as a function from `N` to `N` which might seem like a surprising beast. Two points are striking. First we are using the type `N` in its own definition: this type is said to be *recursive*. The second point is to declare a function without giving it a definition. Indeed `suc` has no other definition, and thus giving it an argument does not trigger a computation. For instance `suc zero` stays stuck like that and does not reduce further. The language is built to take advantage of data constructors to define functions by *pattern matching*.

Note that the same name can be used for different data constructors of distinct data types. AGDA makes use of type annotations to resolve ambiguities.

To illustrate definitions made with pattern matching the following program defines a function called `not` which negates a Boolean value:

```
not : Bool → Bool
not true  = false
not false = true
```

This definition is made of two equations: one for each constructor of the inspected argument. In the case for `true` we return `false` and in the case for `false` we return `true`.

We now go on a more interesting definition made with pattern matching. The following program defines a function called `triple` which multiplies by three its argument:

```
triple : ℕ → ℕ

-- 3 * 0 = 0
triple zero    = zero

-- 3 * (1 + n) = 3 + 3 * n
triple (suc n) = suc (suc (suc (triple n)))
```

While this definition is longer than we could expect (`triple n = 3 * n`), it has a pedagogical interest. This definition is made of two equations: one for each constructor of the inspected argument. In the successor case the function `triple` is used in the definition of `triple` itself! This function is said to be recursive.

We can now reveal that all the operations on natural numbers we have seen so far are actually defined within the language using recursive definitions. Here are for instance the definitions for `+` and `*`. To express that we want these operators to be infix, we declare `+` as `_+_`. Each `_` indicates special places where operands should go. Both functions `_+_` and `_*_` pattern-match on their first argument:

```
_+_ : ℕ → ℕ → ℕ
zero + n = n          -- 0 + n = n
suc m + n = suc (m + n) -- (1 + m) + n = 1 + (m + n)

*_ _ : ℕ → ℕ → ℕ
zero * n = zero      -- 0 * n = 0
suc m * n = n + m * n -- (1 + m) * n = n + m * n
```

The notation we used for literal numbers such as `2` or `4` is actually strictly equivalent to `suc (suc zero)` and `suc (suc (suc (suc zero)))` respectively.

We define a last operation on natural numbers, namely division by two. We do not show the definition for `÷`. Indeed, division by a known constant is indeed much simpler to define recursively and illustrates subtler pattern matching. The function is declared as `_/2` and hence has to be used either prefix as `_/2 4` or postfix as `4 /2`. You can notice how spaces are of prime importance in AGDA since they allow to freely choose meaningful names for functions. The function `_/2` proceeds as follows. The two base cases for zero and one are handled in the first two equations. The last equation deals with all numbers strictly greater than one which makes a simple recursive call to

obtain our result:

```

_/2 : ℕ → ℕ

-- 0 / 2 = 0
zero /2 = zero

-- 1 / 2 = 0
suc zero /2 = zero

-- (2 + n) / 2 = 1 + n / 2
suc (suc n) /2 = suc (n /2)

```

Polymorphism Sometimes functions do not need to know the nature of parts of their arguments. This means that many types could be acceptable for such functions. An extreme case is the identity function which simply returns its argument. What type should we give it? $\mathbb{N} \rightarrow \mathbb{N}$, $\text{String} \rightarrow \text{String}$, or $\text{List } \mathbb{N} \rightarrow \text{List } \mathbb{N}$ and the list goes on indefinitely. The solution to this issue is called *polymorphism*. We generalize the type of the function type to $\forall\{A\} \rightarrow A \rightarrow A$, which reads “for all type A a function from A to A ”. The polymorphic identity function (`id`) can thus be written:

```

id : ∀{A} → A → A
id x = x

-- The function id can be used at different types
prog4 : Interactive
prog4 = print (id "Hello " ++ show (id 42))

```

Lists We define two new data types that we assumed to be special so far, namely `List` and `Tree`. The type for lists –such as the type \mathbb{N} – is made of two data constructors. A base constructor named `[]` and pronounced “nil” represents the empty list. Another constructor named `_::_` and pronounced “cons” appends one element to the front of a list. The notation `[n1 , n2 , ...]` is simply a shorthand for the less familiar `n1 :: n2 ... :: []`.

```

data List (A : Set) : Set where
  [] : List A
  _:_ : A → List A → List A

```

Plenty of interesting functions can be defined recursively on lists. We

show one of them which while being simple can be insightful. The function `length` takes a list and returns its length as a natural number. The function is defined with two equations: one for the empty list whose length is zero, and one for any constructed list whose length is the successor of the length of the *tail* of the list. In short this function replaces the lists constructors by the constructors of natural numbers. This highlights the fact that both types share the same structures. Natural numbers are lists of meaningless elements and lists are natural numbers whose constructors are annotated by elements. The function `length` is polymorphic. It takes a list, forgets the elements to reveal the bare structure behind any list: a natural number.

```
length : ∀{A} → List A → ℕ
length []      = zero
length (x :: xs) = suc (length xs)
```

Let us remark that since we make no use of the name `x` in the second equation for `length` we could have use the special *wildcard* pattern. The wildcard pattern is noted `_` and is used to replace a name to state that we do not use this part of the value.

Another common function on lists is the function `map`. The function `map` takes a function `f` from a type `A` to a type `B`, and also takes a list of type `List A` on which it applies `f` on every element to build the resulting list of type `List B`. Since `map` takes a function as argument it is what we call an *higher-order* function. Moreover, the function `map` is also recursive and polymorphic. Here is the definition of `map`:

```
map : {A B : Set} → (A → B) → List A → List B
map f []          = []
map f (x :: xs) = f x :: map f xs
```

We now introduce an extension to pattern matching equations, namely the `with` construct. This construct extends a pattern-matching-based definition with new columns. This construct is of great effect when combined with dependent pattern-matching. However, we present it here on a simpler example, the function to filter a lists. The function `filter`, takes a predicate `p` and a list `xs` and keeps only the elements of `xs` which satisfy the predicate `p`. The `with` construct is used here to select a branch according to the result of the predicate:

```

filter : {A : Set} → (A → Bool) → List A → List A
filter p [] = []
filter p (x :: xs) with p x
filter p (x :: xs) | true  = x :: filter p xs
filter p (x :: xs) | false = filter p xs

```

Moreover, an ellipsis ... can be used to elide a redundant equation prefix. Hence we can write `filter`, this way:

```

filter : {A : Set} → (A → Bool) → List A → List A
filter p [] = []
filter p (x :: xs) with p x
...          | true  = x :: filter p xs
...          | false = filter p xs

```

Trees and forests We continue with trees the exploration of our previously assumed basic types. Trees can be defined with a new data type. This data type `Tree` has a single data constructor named `node` with two arguments: the node label of type `A` and the children as a list of subtrees (`List (Tree A)`).

```

data Tree (A : Set) : Set where
  node : A → List (Tree A) → Tree A

```

Commonly we call forests the lists of trees. The name is suggestive and the type is shorter to write. Its definition shows how types are treated such as other values:

```

Forest : Set → Set
Forest A = List (Tree A)

```

Trees and forests are interdependent. They finally build a pair of *mutually recursive* types. This is the occasion to define a pair of mutually recursive functions. We define `sumTree` and `sumForest` which respectively compute the sum of the labels in a tree of natural numbers and in a forest of natural numbers. For AGDA to accept these definitions, the two declarations must come before the two definitions.

```

sumTree : Tree N → N

```

```

sumForest : Forest N → N

sumTree (node n forest) = n + sumForest forest
sumForest [] = 0
sumForest (tree :: forest) = sumTree tree + sumForest forest

```

Robust program representations So far we have seen three distinct types to represent programs: text values (`String`), lists of tokens (`List String`), and trees (`Tree String`). These three types are based on the type `String` and thus they are loose representations. For instance the tokens can be malformed tokens and the tree labels could be malformed as well. We could introduce data types for tokens and tree labels which rule out the malformed text values. However we can do even better. Here is a data type for our tiny programming language of the beginning:

```

data Program : Set where
  'num   : N → Program
  'text  : String → Program
  'list  : List Program → Program
  'node  : Program → Program → Program
  '+'    : Program → Program → Program
  '-     : Program → Program → Program
  '*'    : Program → Program → Program
  '÷     : Program → Program → Program
  '++    : Program → Program → Program
  'print : Program → Program
  'show  : Program → Program
  '>>    : Program → Program → Program
  'var   : String → Program
  'let   : String → Program → Program → Program

```

The type `Program` is a recursive type inducing a specialized tree structure. Each construct is modeled by a single data constructor which is named with the corresponding label we have for trees plus an extra `'` character. The data constructor also specifies how many subtrees are expected and what kind of subtree is expected. Here are two of our previous examples represented with our type `Program`:

```

prog4 : Program
prog4 = '* ('num 2) ('+ ('num 10) ('num 11))

```

```

prog5 : Program
prog5 = '>>
        ('print ('text "Hello! 2 times 21 is equal to "))
        ('print ('show ('* ('num 2) ('num 21))))

```

Here are now a few programs that cannot be represented with the data type `Program`. Indeed while all correct programs can be represented, this simple data type already rules out plenty of wrong programs. Here is first a list of syntactically wrong programs represented as texts:

```

[ "2 *"      -- missing operand
, "(2"      -- missing parenthesis
, "2 {- "   -- non-closed comment
, "bla"     -- unknown construct or variable
]

```

Here is now a list of still syntactically wrong programs represented as trees:

```

[ node "+" [ p ]           -- missing operand
, node "+" [ p1 , p2 , p3 ] -- extra operand
, node "bla" []           -- unknown construct
, node "var" [ node "bla" [] ] -- unknown variable
]

```

Finally, there are still wrong programs accepted by all of our representations. Those are the ill-scoped programs and the ill-typed programs. Here is a list of ill-scoped or ill-typed programs using the type `Program`:

```

[ 'var "bla"                -- unknown variable
, '+ ('num 1) ('text "Hello!") -- ill-typed
, 'let "x" ('var "x") ('num 1) -- unknown variable
]

```

The rest of this document focuses on how to improve program representation to better account for the handling of variables in programs.

Declaring the last operations We have shown how various constructs of our tiny programming language can be defined in AGDA. We now quickly declare the remaining constructs that can be AGDA functions. Thus we declare `_÷_`, `_÷÷_`, `show`, `_++_`, `print`, and `_>>_`. While we give no definitions for these functions the declarations describe both their syntax and their


typing very concisely:

```

_÷_      : ℕ → ℕ → ℕ
_÷_      : ℕ → ℕ → ℕ
show     : ℕ → String
_+_      : String → String → String
print    : String → Interactive
->>-    : Interactive → Interactive → Interactive

```

The core constructs from our tiny language that we do not define are literals (numbers, texts, lists), variables and the `let` construct. However, in order to use all these functions and constructors we introduced a fairly-discrete construct, namely application. Since application is written as a simple juxtaposition (using spaces or parentheses) we might forget it. In our tree representation we use the label "app" for these application nodes. We can now view the example program `2 * (10 + 11)` as a tree where the application nodes replace the special nodes of operations:



Agda types In AGDA, the usual function space is written $A \rightarrow B$, while the dependent function space is written $(x : A) \rightarrow B$ or $\forall (x : A) \rightarrow B$. An implicit parameter, introduced via $\forall\{x : A\} \rightarrow B$, can be omitted at a call site if its value can be inferred from the context. There are shortcuts for introducing multiple arguments at once or for omitting a type annotation, as in $\forall\{A\} \{i\ j : A\} x \rightarrow e$.

In AGDA, `Set` (or `Set0`) is the type of small types such as `ℕ`, `List String`, and `Maybe (Bool × ℕ)`. `Set1` is the type of `Set`, `Set → Bool`, `ℕ → Set`, and `Set → Set`.

There is no specific sort for propositions in AGDA: everything is in `Set ℓ` for some ℓ . The unit type is a record type with no fields named `⊤`. It also represents the `True` proposition. The empty type is named `⊥` and is an

(inductive) data type with no constructors. It also represents the `False` proposition. The negation $\neg A$ is defined as $A \rightarrow \perp$.

Lexical conventions We recall that AGDA is strict about whitespace: $x \leq y$ is an identifier, whereas $x \leq y$ is an application. This allows naming a variable after its type (deprived of any whitespace). For example: $x \leq y$ might be a variable of type $x \leq y$, that is, a proof of $x \leq y$.

Various notions of equality Advanced logics and proof systems exhibit various forms of equality. The *definitional equality* is a first form of equality which is deeply rooted in the computation rules of the programming language itself, here, AGDA. Two terms are *definitionally equal* if and only if they can both reduce to a common term. In AGDA, the reduction is a combination of β -reductions, η -conversions, and application of the definition equations.

Here are a few examples of definitionally equal terms. We informally use `=` for this equality since this symbol is the one used for definitions.

```
(λ x → x) suc zero = suc zero
suc = λ x → suc x
zero + n = n
```

Here is an example of two terms which are not definitionally equal even if this equality seems natural:

```
n + zero ≠ n
```

Definitional equality is completely automatic and thus requires no help from the user. However if we want to include some reasoning steps we can use the *propositional equality*. The symbol used in AGDA and further in this document is \equiv . To simplify matters, here is the definition of the propositional equality specialized to values of small types (`Set0`), namely \equiv_0 :

```
data _≡₀_ {A : Set₀} (x : A) : A → Set₀ where
  refl : x ≡₀ x

_≠₀_ : {A : Set₀} → A → A → Set₀
x ≠₀ y = ¬(x ≡₀ y)
```

We do not intend to explain the subtleties of such a definition. We only mention that there is a single constructor for this type requiring both sides to be same, hence definitionally equal.

```
-- This is definitionally true, hence immediately proven by refl
zero + n ≡ n
```

```
-- This is provable by a simple induction on n
∀ n → n + zero ≡ n
```

Sometimes, two functions `f` and `g` cannot be proved propositionally equal but still produce the same output on every possible input. In this case, the functions `f` and `g` are said to be *equal pointwise*. The symbol for pointwise equality is $\overset{\circ}{=}$. Here is the definition of pointwise equality in AGDA, specialized to small types (`Set0`), non-dependent functions:

```
 $\overset{\circ}{=}_0$  : {A B : Set0} (f g : A → B) → Set0
f  $\overset{\circ}{=}_0$  g = ∀ x → f x ≡0 g x
```

To illustrate pointwise equality we define the function `not2` which iterates the function `not` two times.

```
not2 : Bool → Bool
not2 x = not (not x)
```

We then show that the function `not2` is equal pointwise to the identity function. Given the definitions for $\overset{\circ}{=}_0$ and `not2` this exactly amounts to show: $\forall x \rightarrow \text{not} (\text{not } x) \equiv x$.

```
not2 $\overset{\circ}{=}$ id : not2  $\overset{\circ}{=}_0$  id
not2 $\overset{\circ}{=}$ id true  = refl -- true = not (not true) definitionally
not2 $\overset{\circ}{=}$ id false = refl -- false = not (not false) definitionally
```

Equalities are said to be either intensional or extensional. Let us use sorting functions as an example to illustrate the difference. If we have two different sorting functions each implementing one sorting algorithm such as bubble sort and quick sort. The two functions are said to be extensionally equal since they both return the same sorted list for any given list. However bubble sort and quick sort are not intentionally equal. The *intention* is actually opposite. When we write a quick sort function, we intentionally want to make it different than bubble sort!

In AGDA both the definitional and the propositional equalities are currently intensional. There is an ongoing will to make the propositional equality more extensional. While this sounds like a radical change, the propositional equality is actually already compatible with extensionality. Indeed, there is no way to contradict extensionality with the current propositional equality. In short there is no proof than quick sort is different from bubble sort and thus making them provably equal introduce no contradiction.

The *pointwise equality* is by construction extensional and in section 4.2 we explain a generalized definition to relate functions which sends related inputs to related outputs.

Full development online We use some definitions from AGDA’s standard library: natural numbers, booleans, lists, and applicative functors (`pure` and `_*`).

For the sake of conciseness, the code fragments presented in this document are sometimes not perfectly self-contained. However, a complete AGDA development is available online [Pouillard, 2011a].

Outline This document is organized as follows.

The next three chapters delve into the nominal approach. We introduce the nominal approach (chapter 2) and present our system: a safe programming interface for the nominal style. We then present how to use it (chapter 3), and what happens behind the scene to implement it and show its safety (chapter 4).

The next chapter (5) covers a nameless approach known as “de Bruijn indices”. We present how to seamlessly extend our system to this new approach. This chapter covers the same aspects: the interface and its implementation, its usage, its soundness.

The remaining chapter (6) discusses different variations and combinations of these variations. At the same time this chapter covers the related work and concludes.

Chapter 2

The nominal approach

This chapter is organized as follows. The first section informally introduces and presents several techniques to represent data structures with bindings in a nominal style. Section 2.2 describes our solution, an interface to program with names and binders.

2.1 Introduction to the nominal approach

2.1.1 Warm-up: the bare nominal approach

The bare approach to abstract syntax with names and binders in nominal style requires very little infrastructure to start with. Names and binders are represented by so-called atoms. The set of atoms is countably infinite and the only required operation is an equality test. Using natural numbers as a concrete representation for atoms is a common and sensible choice.

```
-- A set of atoms (could be  $\mathbb{N}$ )
Atom : Set

-- Atom is countably infinite; here are some atoms:
-- x,y,z... could be represented by 0,1,2...
x y z f g {...} : Atom

-- The equality test on atoms
_==^A_ : (x y : Atom) → Bool
```

Given the type `Atom` we can readily define algebraic data types for abstract syntax with names and binders. Our running example is the untyped λ -calculus defined below. The type `TmA` (`Tm` for “term” and ^A for “atom”) is made of three data constructors. The constructor `V` is for variables and simply holds an atom. The constructor `._.` is the function application construct made of two subterms. The constructor `λ` takes an atom and a subterm

in which this atom is considered bound. This means that the construct λ is introducing a variable. On contrary the atom in the construct V is said to be free. Indeed this atom is supposed to refer to a bound atom. The constructor `Let` also takes an atom but two subterms. The atom is bound in the second subterm only.

```
data TmA : Set where
  V      : (x : Atom) → TmA
  _·_    : (t u : TmA) → TmA
  λ     : (b : Atom) (t : TmA) → TmA
  Let   : (b : Atom) (t u : TmA) → TmA
```

It is striking that there is no formal distinction between the atoms that represent binders and those that represent occurrences. Neither is there any indication of the scope of the binders. This calls for improvement.

We consider it very important to highlight that atoms are used for two distinct purposes whether they are in a binding position or a free position. Starting from section 2.2, we embrace that distinction and provide distinct types for these usages.

Here are two term examples, the identity function and the application function:

```
-- λx. x
idTmA : TmA
idTmA = λ x (V x)

-- λf. λx. f x
apTmA : TmA
apTmA = λ f (λ x (V f · V x))
```

The strength of this approach resides in its simplicity. The representation of terms closely follows the concrete syntax of the language. The main issue is adequacy: there are multiple equivalent representations of the “same” term.

Indeed the choice of atoms is fairly arbitrary. Two terms can represent the same piece of syntax when they differ only by a consistent renaming of bound names. In this situation two such terms are said to be α -equivalent. Here are for example two α -equivalent terms:

```
tx : TmA
tx = λ x (V f · V x)
ty : TmA
ty = λ y (V f · V y)
```

The term t^x is α -equivalent to t^y since consistently renaming the bound name x by y in the term t^x yields the term t^y . Respectively consistently renaming y by x in t^y yields t^x . However α -equivalence can be subtle, let us observe a third term t^f which is *not* α -equivalent to neither t^x nor t^y :

```
tf : TmA
tf = λ f (V f · V f)
```

An *inconsistent* renaming of the bound name x by f in t^x yields t^f , while a consistent renaming of the bound name f by x in t^f does *not* yield t^x .

We do not delve more into the details of α -equivalence yet. We focus on the fact that the concrete naming of a term is a representation issue and should not be relevant for the computation. In particular, good functions should be independent of this representation issue. This property can be stated as follows:

A function is well-behaved if, when applied to α -equivalent arguments, it produces α -equivalent results.

To illustrate well-behaved functions we give a few examples. The following functions `rmA` (removes an atom from a list) and `fv` (lists the free variables/atoms of a term) are well-behaved:

```
rmA : Atom → List Atom → List Atom
rmA [] = []
rmA x (y :: ys) =
  if x ==A y then rmA x ys
  else y :: rmA x ys
```

```
-- Since rmA behaves well, this holds:
test-rmA : rmA x [ x ] ≡ rmA y [ y ]
test-rmA = refl -- both reduces to []
```

```
fv : TmA → List Atom
fv (V x) = [ x ]
fv (t · u) = fv t ++ fv u
fv (λ x t) = rmA x (fv t)
fv (Let x t u) = fv t ++ rmA x (fv u)
```

```
-- Since fv behaves well, this holds:
test-fv : fv tx ≡ fv ty
test-fv = refl -- both reduces to [ f ]
```

We now illustrate misbehaving functions with two examples. The function `ba` computes the list of bound atoms. The function `cmp-ba` takes two terms and compare their bound variables if both are λ constructs.

```

ba : TmA → List Atom
ba (λ x t)      = x :: ba t
ba (Let x t u) = x :: ba t :: ba u
ba (t · u)      = ba t ++ ba u
ba (V -)       = []

[x]≠[y] : [ x ] ≠ [ y ]
[x]≠[y] = {! omitted !}

-- Since ba does not behave well:
test-ba : ba tx ≠ ba ty
test-ba = [x]≠[y]

cmp-ba : TmA → TmA → Bool
cmp-ba (λ x -) (λ y -) = x ==A y
cmp-ba - - - - - = false

-- Since cmp-ba does not behave well:
test-cmp-ba : cmp-ba tx tx ≠ cmp-ba tx ty
test-cmp-ba () -- true ≠ false

```

In traditional informal developments, α -equivalence is identified with equality. The upside is that every definition respects α -equivalence. Hence functions automatically map α -equivalent inputs to α -equivalent outputs. The downside is that building functions is more difficult. Indeed what were misbehaving functions before, are not functions anymore when we identify α -equivalence and equality. They are not functions simply because they map equal inputs to different outputs.

One of the goals of our system is to guarantee that well-typed functions are well-behaved. A function like `cmp-ba` is ill-typed in our system. A function like `ba` is typed differently than the function `fv`, effectively disarming the function `ba`.

2.1.2 Using well-formedness judgements

One way to define the scoping rules is to define a well-formedness judgement. It can be done using a recursive predicate over the structure of terms. It is a well-known technique used in the definition of type systems. The standard presentation makes use of a set of inference rules where judgements are

made of an environment, a term and a type. The environment tracks the type of each introduced variable. To specify scoping rules one follows the same presentation but without the types. We directly focus on a formal presentation in AGDA since inference rules (and grammars) have a direct translation into inductive types. First one defines the environments which are nothing more than lists of atoms:

```
data Env : Set where
  ε      : Env
  _,_   : (Γ : Env) (x : Atom) → Env
```

Then comes the definition of the membership predicate. It states when an atom is a member of some environment. This is done through two rules: one applies when the atom is present at the first position of the environment, and a second rule applies when is present in the tail of the environment. Here is the definition in AGDA where lines drawn with comments give constructors a flavor of inference rules (a comment starts with `--` and stops at the end of line):

```
data _∈_ x : (Γ : Env) → Set where

  here  : ∀ {Γ}      → -----
                    x ∈ (Γ , x)

  there : ∀ {Γ y}    →  x ∈ Γ
                    → -----
                    x ∈ (Γ , y)
```

Finally one can state the scoping rules for the λ -calculus. The definition boils down to using an environment to track bound atoms, using the membership predicate at variable occurrences and pushing the binders onto the environment:

```

data ⊢_ Γ : TmA → Set where
  V : ∀ {x}           → x ∈ Γ
                        → -----
                        Γ ⊢ V x

  ·_ : ∀ {t u}        → Γ ⊢ t
                        → Γ ⊢ u
                        → -----
                        Γ ⊢ t · u

  λ : ∀ {t b}         → Γ , b ⊢ t
                        → -----
                        Γ ⊢ λ b t

  Let : ∀ {t u b}     → Γ ⊢ t
                        → Γ , b ⊢ u
                        → -----
                        Γ ⊢ Let b t u

```

We can now state that our example terms are well-scoped in the empty environment. Thanks to implicit arguments and to the definition of the membership predicate, the proof that a term is well-scoped is the term itself in de Bruijn notation (`zero` is `here` and `suc` is `there`).

```

⊢id : ε ⊢ idTmA
⊢id = λ (V here)

⊢ap : ε ⊢ apTmA
⊢ap = λ (λ (V (there here) · V here))

```

The use of AGDA implicit arguments keep these examples concise. The curious reader who is unfamiliar with AGDA might find interesting to look at the same (second) example with all arguments made explicit:

```

⊢ap' : ε ⊢ apTmA
⊢ap' = λ {ε} {λ x (V f · V x)} {f}
      (λ {ε , f} {V f · V x} {x}
        (V {ε , f , x} {f}
          (there {f} {ε , f} {x}
            (here {f} {ε}))) ·
        V {ε , f , x} {x}
        (here {x} {ε , f})))

```

However this technique is not directly what we are looking for. Indeed

these scoping properties are explicitly stated on the side of each definition. We are looking for something more integrated into the types. This would enable well-formedness to be enforced in a more pervasive and automatic manner. Happily this technique can be adapted to merge the scoping information and the term.

2.1.3 Well-scoped terms

Merging terms and scoping rules is not a new technique and amounts to removing the term index of the scoping predicate. The resulting term type is indexed by an environment. This environment is extended at binders and queried at variable occurrences:

```
data TmJ Γ : Set where
  V    : ∀ {x} → x ∈ Γ → TmJ Γ
  _·_  : TmJ Γ → TmJ Γ → TmJ Γ
  λ    : ∀ b → TmJ (Γ , b) → TmJ Γ
  Let  : ∀ b → TmJ Γ → TmJ (Γ , b) → TmJ Γ
```

One can define our two term examples in this style as well:

```
idTmJ : TmJ ε
idTmJ = λ x (V here)

apTmJ : TmJ ε
apTmJ = λ f (λ x (V (there here) · V here))
```

The function `fv` can easily be adapted to this style, while the function `rmA` can be reused:

```
fv : ∀ {Γ} → TmJ Γ → List Atom
fv (V {x} _) = [ x ]
fv (t · u)   = fv t ++ fv u
fv (λ x t)   = rmA x (fv t)
fv (Let x t u) = fv t ++ rmA x (fv u)
```

Alas the functions `ba` and `cmp-ba` are not rejected by this style. Their types becomes:

```
ba : ∀ {Γ} → TmJ Γ → List Atom
cmp-ba : ∀ {Γ1 Γ2} → TmJ Γ1 → TmJ Γ2 → Bool
```

We have described the scoping rules but have not made much progress. Indeed two different but α -equivalent terms can still distinguished. Worse,

by introducing environments and membership proofs as explicit objects, we have introduced new issues. Functions receiving terms now also receive or know the environment in which these terms are scoped. This extra information can be used by functions making them less agnostic to choices of free variables. Here are two functions using the environment by pattern-matching:

```

fast-fv :  $\forall \{\Gamma\} \rightarrow \text{Tm}^J \Gamma \rightarrow \text{List Atom}$ 
fast-fv  $\{\epsilon\}$  _ = [] -- for sure the term is closed
fast-fv      t = fv t

env-length : Env  $\rightarrow \mathbb{N}$ 
env-length  $\epsilon$  = 0
env-length ( $\Gamma$  , _) = 1 + env-length  $\Gamma$ 

env-length-TmJ :  $\forall \{\Gamma\} \rightarrow \text{Tm}^J \Gamma \rightarrow \mathbb{N}$ 
env-length-TmJ  $\{\Gamma\}$  _ = env-length  $\Gamma$ 

```

If one consider the term as the only input of `env-length-TmJ`, then this function does not behave well since the same term can be scoped in two environments of different length. The issue is that Γ is a concrete input to the functions like `fv`, `ba`, `cmp-ba`, `fast-fv` and `env-length-TmJ`.

In order to fix these issues we make use of abstract types and yet provide an interface with precise types for a safe usage. Environments (the type `Env`) are abstracted as so-called *worlds*. Atoms used in binding positions are abstracted as so-called *binders* and cannot even be compared. Atoms in free positions are glued together with the proof of their membership in the environment and are called *names*. Finally a few extra operations allow to deal with these abstract types in a safe way.

2.2 NomPa: A programming interface in nominal style

Our approach is called NOMP and proposes a programming interface whose types are abstract. This interface (Figure 2.1) admits a concrete implementation (section 4.1). Hiding this implementation keeps the programmer away from the internal details and gives rise to interesting properties via parametricity.

While the complete interface is given in Figure 2.1, each ingredient of the interface is presented in turn. First here are the building blocks needed to define data types with names and binders in a nominal style.

2.2. NOMPA: A PROGRAMMING INTERFACE IN NOMINAL STYLE 49

```

record NomPa : Set1 where
  constructor mk

infixr 5 _<_
infix 3 _⊆_
infix 2 _#_

field
  -- minimal kit to define types
  World   : Set
  Name    : World → Set
  Binder  : Set
  _<_     : Binder → World → World

  -- An infinite set of binders
  zeroB : Binder
  sucB  : Binder → Binder

  -- Converting names and binders back and forth
  nameB  : ∀ {α} b → Name (b < α)
  binderN : ∀ {α} → Name α → Binder

  -- There is no name in the empty world
  ∅       : World
  ¬Name∅  : ¬ (Name ∅)

  -- Names are comparable and exportable
  _==N_  : ∀ {α} (x y : Name α) → Bool
  exportN : ∀ {α b} → Name (b < α) → Name (b < ∅) ⊔ Name α

  -- The fresh-for relation
  _#_     : Binder → World → Set
  _#∅_    : ∀ b → b # ∅
  suc#    : ∀ {α b} → b # α → (sucB b) # (b < α)

  -- inclusion between worlds
  _⊆_     : World → World → Set
  coerceN : ∀ {α β} → (α ⊆ β) → (Name α → Name β)
  ⊆-refl  : Reflexive _⊆_
  ⊆-trans : Transitive _⊆_
  ⊆-∅     : ∀ {α} → ∅ ⊆ α
  ⊆-<     : ∀ {α β} b → α ⊆ β → (b < α) ⊆ (b < β)
  ⊆-#    : ∀ {α b} → b # α → α ⊆ (b < α)

```

Figure 2.1: The NOMPA interface

2.2.1 All we need to define nominal syntax

To replace environments in the definition of well-scoped terms, we introduce an abstract notion of worlds. Hence the interface starts by an abstract type for worlds:

```
World : Set
```

We consider it very important to highlight that atoms are used for distinct purposes whether they are in a binding position or a free position. We embrace that distinction and provide distinct types for these usages.

We first introduce a type for binders. Binders are atoms to be used only in binding positions. While these binders are explicitly accessible in terms it does not mean one can distinguish two α -equivalent terms by simply looking at the binders. Indeed the interface does not provide any means to distinguish two binders. Internally binders are really atoms but the interface keeps them abstract:

```
Binder : Set
```

For atoms to be put in free position, we introduce a type called `Name`. This type is indexed by a world in which the name is scoped. A name glues together an atom and a proof of its membership in the world.

```
Name : World → Set
```

Worlds can be thought of as a set of atoms (names or binders) but will be thought of as relations between names in section 4.3.

Now that we have binders and worlds, we can introduce a way to extend a world with a binder. This is analogous to the constructor `_,_` of section 2.1.2 for extending an environment with an atom.

```
_<_ : Binder → World → World
```

At this point, we have all we need to build data types with names and binders. The new definition of `Tm` is very close to the previous one. All we have to do is to use `_<_` instead of `_,_`, rename Γ into α , and use `Name` instead of an atom plus a membership proof:

```
data Tm  $\alpha$  : Set where
  V      : Name  $\alpha$  → Tm  $\alpha$ 
  _·_    : Tm  $\alpha$  → Tm  $\alpha$  → Tm  $\alpha$ 
   $\lambda$   :  $\forall b \rightarrow Tm (b < \alpha) \rightarrow Tm \alpha$ 
  Let    :  $\forall b \rightarrow Tm \alpha \rightarrow Tm (b < \alpha) \rightarrow Tm \alpha$ 
```

Here binding constructs rely on dependent types, since the binder \mathbf{b} is used in the type of the subterm. In our previous work [Pouillard and Pottier, 2010], we avoided dependent types on purpose to make the results more widely usable. The non-dependent approach based on links can be built on top of the one we are presenting here. A weak link [Pouillard and Pottier, 2010] from α to β is represented by a binder \mathbf{b} such that β is equal to $\mathbf{b} \triangleleft \alpha$. Hence we want to first study in detail the dependent approach with its logical relation.

Here is a trivial example of a function that traverses a term and measures its size. It is remarkable for its simplicity: name abstractions are traversed without fuss. The code would be exactly the same in the bare nominal approach of section 2.1.1. In particular this unaltered induction suggests that the expressiveness with respect to the bare nominal approach is kept. It is efficient: no renaming, substitution, or shifting is involved. Polymorphic recursion is exploited: the call to `size t` in the λ (and `Let`) case is at some inner world.

```
size :  $\forall \{\alpha\} \rightarrow \text{Tm } \alpha \rightarrow \mathbb{N}$ 
size (V _)           = 1
size (t · u)         = 1 + size t + size u
size ( $\lambda$  _ t)      = 1 + size t
size (Let _ t u)     = 1 + size t + size u
```

2.2.2 Building binders and names

In order to build terms, we need binders and names. Binders are introduced in our system using two simple primitive operations called `zeroB` and `sucB`:

```
zeroB : Binder
sucB  : Binder  $\rightarrow$  Binder
```

We introduce a cheap convenience function to turn any natural number into a binder:

```
_B :  $\mathbb{N} \rightarrow$  Binder
zeroB = zeroB
(suc n)B = sucB (nB)
```

In effect, binders are natural numbers with a limited interface. Here we give only zero and successor but all of the operations on natural numbers could be exposed as well. One key limitation, however, is that no information must leak out of a binder. Hence no equality test over binders is provided.

This is clarified when building the logical relation in section 4.3.

To build new names we provide a function, called `nameB`, to turn binders into names. For completeness we introduce the converse function `binderN`. However, this function, without any properties on it, is of very little use.

```
nameB  : ∀ {α} b → Name (b < α)
binderN : ∀ {α} → Name α → Binder
```

While `nameB` can turn any binder into a name, it imposes the shape of the world that the produced name inhabits. Indeed `nameB b` returns a name in any world whose last introduced binder is `b`. We will soon see that this is an important limitation and how to overcome it. For the moment, we have enough blocks to define a representation of the identity function!

```
idTm : ∀ {α} → Tm α
idTm = λ x (V (nameB x))
      where x = 0B
```

Our second example is a representation of the λ -term for false ($\lambda x. \lambda x. x$). We use the same binder twice on purpose to show a use of name shadowing. To type check the variable occurrence of `x`, the world `x < α` is implicitly given to `nameB` which returns a name in the world `x < x < α`. Then each of the two “ λx ” takes away one “`x <`” to leave the final term in `α`.

```
falseTm : ∀ {α} → Tm α
falseTm = λ x (λ x (V (nameB x)))
        where x = 0B
```

One however faces an issue when building a representation of the λ -term for true ($\lambda x. \lambda y. x$). The naïve approach does not type-check:

```
-- this does not type-check
trueTm : ∀ {α} → Tm α
trueTm = λ x (λ y (V (nameB x)))
        where x = 0B
              y = 1B
```

While `nameB x` inhabits any world of the form `x < β`, the context expects a name in the world `y < x < α`. We introduce means to move a name from one world to another in section 2.2.3.

Closed terms The above well-typed terms (`idTm` and `falseTm`) have no free variables. They are said to be closed. They both have the type:

$\forall \{\alpha\} \rightarrow \mathbf{Tm} \alpha$. This world-polymorphic type reflects this closedness property. If any world fits, the empty world (noted \emptyset) fits as well. So \mathbf{idTm} also has the type $\mathbf{Tm} \emptyset$. Conversely, because the empty world is included in all worlds, we will be able to move any term from the empty world to an arbitrary world. In short $\mathbf{Tm} \emptyset$ and $\forall \{\alpha\} \rightarrow \mathbf{Tm} \alpha$ are *isomorphic* types.

To allow exploiting the fact that any term in the empty world is closed, we introduce $\neg\mathbf{Name}\emptyset$, which witnesses that there is no name in the empty world:

```
 $\emptyset$            : World
 $\neg\mathbf{Name}\emptyset$  :  $\neg(\mathbf{Name} \emptyset)$ 
```

This last primitive operation is provided to secure the system against a buggy notion of closed terms but not only! In various situations it enables arguing that certain cases are impossible, like looking up a name in an empty environment.

2.2.3 World widening: Name weakening

It is often necessary to widen the world which a name inhabits. One example was the \mathbf{trueTm} definition we tried earlier: to type-check it, we need to widen the world of the name x . Widening a world causes a loss of static information. For instance we have seen that a term in the empty world is closed: if one widens its world, it is no longer statically known to be closed. We call this operation weakening because it loses some static information. The name “weakening” is also vastly used for the same purpose on typing environments.

To account for the multiple ways in which we could widen a world we introduce a type to represent the inclusion between two worlds. The type $_ \subseteq _$ is for witnesses of world inclusion. If a world α is included in a world β then it is permitted to transport a name (and a term as well, as we shall see) from α to β . The primitive operation \mathbf{coerce}^N takes an inclusion witness, a name and it returns the same name at a wider world. We also introduce an alias to \mathbf{coerce}^N called $_ \langle \mathbf{-because-} _ \rangle$ which is useful to keep the code separated from the typing/proof: the proof, which appears between the angle brackets, can be safely skipped by the reader.

```
 $\_ \subseteq \_$       : World  $\rightarrow$  World  $\rightarrow$  Set
 $\mathbf{coerce}^N$     :  $\forall \{\alpha \beta\} \rightarrow \alpha \subseteq \beta \rightarrow (\mathbf{Name} \alpha \rightarrow \mathbf{Name} \beta)$ 

infix 0  $\_ \langle \mathbf{-because-} \_ \rangle$ 
```

```

_⟨-because_-⟩ : ∀ {α β} → Name α → α ⊆ β → Name β
_⟨-because_-⟩ n pf = coerceN pf n
-- We can now write:
--   x ⟨-because some proof -⟩
-- Instead of the more noisy:
--   coerceN (some proof) x

```

World inclusion rules Rules for world inclusion are given in figure 2.1. World inclusion is reflexive and transitive (\subseteq -*refl* and \subseteq -*trans*). The empty world is a least element for world inclusion (\subseteq - \emptyset). For any binder \mathbf{b} , $\mathbf{b} \triangleleft _$ is covariant, meaning it preserves world inclusion (\subseteq - \triangleleft).

Finally the rule \subseteq - $\#$ states that a world α is included into $\mathbf{b} \triangleleft \alpha$ under the condition that \mathbf{b} is not a member of α . This condition is necessary to ensure the soundness of our library. An interpretation of worlds as sets suggests that this condition might be unnecessary. This shows the limitations of this way of thinking. When we introduce logical relations in section 4.3, we explain that interpreting worlds as *relations* provides a richer viewpoint and explains why this condition is necessary. The following code snippet demonstrates that without this condition one can write wrong programs:

```

wrong : Binder → Binder → Bool
wrong x y = nameB x ==N nameB y ⟨-because wrongProof -⟩
  where postulate
        wrongProof : y < ∅ ⊆ x < y < ∅

```

As this code snippet shows, in the absence of this condition it would be permitted to compare names at different worlds and so to compare binders and finally to make a distinction between two α -equivalent terms.

The fresh-for relation The last inclusion rule (\subseteq - $\#$), uses a fresh-for relation called $_ \# _$. We introduce two rules to produce witnesses of this relation. The first rule tells that any binder is fresh for the empty world ($_ \# \emptyset$). The second rule takes the “successor” of a fresh binder (*suc#*). This successor-based way of building new fresh-for witnesses is both simple and efficient. Indeed with these two rules, a binder is not only fresh-for a world but strictly greater than any element of it, allowing for a constant time *suc#* operation. These rules with their types can be found in figure 2.1 as well.

Emptiness of worlds The inclusion relation can express emptiness as well. A world α is empty if it is included in the empty world. We use this definition of emptiness as opposed to an intensional equality with the empty world for two reasons. First $\emptyset \subseteq \alpha$ is always true, hence the converse

is enough. Second we introduce in section 5.2 another operator on worlds called `_+1` for which $\emptyset +1$ is empty but does not reduce to \emptyset .

With a backward reasoning, combining `coerceN` and `¬Name \emptyset` turns any contradiction on names into an inclusion problem, reusing any automation done on inclusion proofs.

```
¬Name : ∀ {α} → α ⊆ ∅ → ¬(Name α)
¬Name α ⊆ ∅ = ¬Name $\emptyset$  ∘ coerceN α ⊆ ∅
```

Eliminating absurd names We introduce convenient eliminators for absurd names which instead of going to the empty type, goes to any type of our choice:

```
Name-elim : ∀ {A : Set} {α} → α ⊆ ∅ → Name α → A
Name-elim pf x with ¬Name pf x
...          | ()
```

```
Name $\emptyset$ -elim : {A : Set} → Name  $\emptyset$  → A
Name $\emptyset$ -elim = Name-elim ⊆-refl
```

Relational reasoning Sometimes one has to build complex inclusion witnesses. While inference would be of great effect here and is currently kept as future work, we propose a modest syntactic tool to build them, namely the `⊆-Reasoning` module. It gives access to the transitivity rule `⊆-trans` in a style which focuses more on the intermediate states of the reasoning rather than on the steps. The syntax is a list of worlds interspersed with inclusion witnesses, with a `■` after the last world. We soon present an example of its use with the function `trueTm`. The code for `⊆-Reasoning` is given below for reference and can be safely skipped.

```
module ⊆-Reasoning where
  infix 2 _■
  infixr 2 _⊆(-)-

  _⊆(-)- : ∀ α {β γ} → α ⊆ β → β ⊆ γ → α ⊆ γ
  _⊆(-)- _ = ⊆-trans

  _■ : ∀ α → α ⊆ α
  _■ _ = ⊆-refl
```


Building any term We can now finally build all nominal terms (up to α -equivalence). In section 3.2.8 we precisely show how to build all λ -terms by converting nominal λ -terms from section 2.1.1 to the type `Tm`. Moreover we show an encoding for nominal types in section 6.5. In the meantime we focus on the term `trueTm`. The two fresh-for rules above are the strict minimum for usability. More rules could be safely added but we focus on the minimal set of required rules here. The logical relation (section 4.2) serves as a criterion to judge what can or cannot be added. Having only these two rules imposes a particular style. A variable occurrence crosses only binders with a lower value to reach its own binder. This does not hinder much the expressiveness but requires to freshen some binders if one cannot prove they are fresh-for a given world. As stated previously removing this minor restriction would require adding more rules to the fresh-for and inclusion relations. One can finally build the term `trueTm`:

```

trueTm :  $\forall \{ \alpha \} \rightarrow \text{Tm } \alpha$ 
trueTm {  $\alpha$  } =  $\lambda x (\lambda y (\forall x^N))$ 
  where x      = 0B
        y      = 1B
        xN    = nameB x  $\langle$ -because  $x \triangleleft \emptyset \quad \subseteq \langle \subseteq \# (\text{suc} \# (x \# \emptyset)) \rangle$ 
                                      $y \triangleleft x \triangleleft \emptyset \quad \subseteq \langle \subseteq \triangleleft y (\subseteq \triangleleft x \subseteq \emptyset) \rangle$ 
                                      $y \triangleleft x \triangleleft \alpha \quad \blacksquare \ - \rangle$ 

```

The proof required to “move x in the correct world” is involved in comparison with the simplicity of the example. By choosing the empty world instead of world polymorphism we could cut in half the proof. To build larger examples, we build specialized building functions that requires only to specify the distance to the binding site [Pouillard, 2011a].

2.2.4 Comparing and refining names

While two binders cannot be compared, our interface allows comparing two names that inhabit a common world. This may seem contradictory, since one can turn binders into names. In fact, binder comparison cannot be implemented in terms of name comparison because two arbitrary binders can be turned only into names in distinct worlds.

```

_==N_ :  $\forall \{ \alpha \} (x y : \text{Name } \alpha) \rightarrow \text{Bool}$ 

```

While world inclusion gives a means to weaken the type of a name, we also need a means to strengthen the type of a name by refining its world. Assume a name x is known to be in the scope of a binder b : x is in the world $b \triangleleft \alpha$. The function `exportN` tests whether x is equal to b ; if so, x is refined to the world $b \triangleleft \emptyset$ (which only b inhabits), otherwise it is refined

2.2. NOMPA: A PROGRAMMING INTERFACE IN NOMINAL STYLE 57

to α . In short given a name of type $\text{Name } (\mathbf{b} \triangleleft \alpha)$, the function $\text{export}^{\mathbf{N}}$ returns the same name with a refined type telling whether the name stands on the \mathbf{b} side or on the α side.

```
 $\text{export}^{\mathbf{N}} : \forall \{ \mathbf{b} \ \alpha \} \rightarrow \text{Name } (\mathbf{b} \triangleleft \alpha) \rightarrow \text{Name } (\mathbf{b} \triangleleft \emptyset) \uplus \text{Name } \alpha$   
 $\text{export}^{\mathbf{N}} = \text{maybe } \text{inj}_2 (\text{inj}_1 (\text{name}^{\mathbf{B}} \_)) \circ' \text{export}^{\mathbf{N}^?}$ 
```

Actually the primitive operation in the interface is the function $\text{export}^{\mathbf{N}^?}$ which returns a result of type $\text{Maybe } (\text{Name } \alpha)$. The function $\text{export}^{\mathbf{N}}$ is then built on top of $\text{export}^{\mathbf{N}^?}$.

```
-- A  $\rightarrow?$  B is the type of partial functions from A to B  
A  $\rightarrow?$  B = A  $\rightarrow$  Maybe B
```

```
 $\text{export}^{\mathbf{N}^?} : \forall \{ \mathbf{b} \ \alpha \} \rightarrow \text{Name } (\mathbf{b} \triangleleft \alpha) \rightarrow? \text{Name } \alpha$ 
```

On top of $\text{export}^{\mathbf{N}^?}$ we also build a convenient eliminator for names which is simply the elimination of the result of $\text{export}^{\mathbf{N}^?}$.

```
 $\text{exportWith} : \forall \{ \mathbf{b} \ \alpha \ \mathbf{A} \} \rightarrow \mathbf{A} \rightarrow (\text{Name } \alpha \rightarrow \mathbf{A}) \rightarrow \text{Name } (\mathbf{b} \triangleleft \alpha) \rightarrow \mathbf{A}$   
 $\text{exportWith } v \ f = \text{maybe } f \ v \circ' \text{export}^{\mathbf{N}^?}$ 
```


Chapter 3

Programming on top of NomPa

3.1 Various examples

3.1.1 Example: computing free variables

We now have enough tools to present a more interesting example, namely a function that constructs a list of the free variables of a term. At variables and applications, the code is straightforward. At a name abstraction, one easily collects the free variables of the body via a recursive call. However, this yields a list of names that inhabit the inner world of the abstraction—a value of type `List (Name (b α))`. This list cannot be directly returned, and this is fortunate, since doing so would let the bound name leak out of its scope! We define an auxiliary function, `rm`, which removes all occurrences of a binder `b` in a list of names and at the same time performs type refinement in the style of `exportN/exportN?`.

```
rm : ∀ {α} b → List (Name (b <math>\alpha</math>)) → List (Name α)
rm b [] = []
rm b (x :: xs) with exportN x -- b is implicit
... {- bound: x≡b -} | inj1 - = rm b xs
... {- free: x≠b -} | inj2 x' = x' :: rm b xs

fv : ∀ {α} → Tm α → List (Name α)
fv (V x) = [ x ]
fv (fct . arg) = fv fct ++ fv arg
fv (λ b t) = rm b (fv t)
fv (Let b t u) = fv t ++ rm b (fv u)
```

The function `rm` applies `exportN` to every name `x` in the list and builds

a list of only those x 's that successfully export to the world α . It exhibits a typical way of using `exportN` to perform a name comparison together with a type refinement. This idiom is recurrent in the programs that we have written.

3.1.2 Example: working with environments

Here is another example, where we introduce the use of an environment.

```

occurs : ∀ {α} → Name α → Tm α → Bool
occurs x0 = occ (λ y → x0 ==N y)
where
  OccEnv : World → Set
  OccEnv α = Name α → Bool
  extend  : ∀ {α b} → OccEnv α → OccEnv (b < α)
  extend  = exportWith false

  occ : ∀ {α} → OccEnv α → Tm α → Bool
  occ Γ (V x)      = Γ x
  occ Γ (t · u)    = occ Γ t ∨ occ Γ u
  occ Γ (λ _ t)    = occ (extend Γ) t
  occ Γ (Let _ t u) = occ Γ t ∨ occ (extend Γ) u

```

The function `occurs` tests whether the name x_0 occurs free in a term. An environment Γ is carried down, augmented when a binder is crossed, and looked up at variable occurrences. Here, this environment is represented as a function of type `Name α → Bool`. Although this is a simple and elegant representation, others exist. For instance, we could represent the environment as a list of binders: the code for this variant is online [Pouillard, 2011a]; see also below.

The definition of `extend` states that, in order to look up x in the environment `extend Γ`, one must first compare x and b and, only if they differ, one must look up x in Γ . If x and b are equal, then this occurrence of x is not free, so `occ Γ (V x)` must return `false`. This is concisely implemented by using the function `exportWith` built on top of `exportN` (Again thanks to the precise typing using worlds, b is passed as an implicit argument to the function `extend` and `exportWith`).

We claim that this code is standard and uncluttered. There is no hidden cost: no renaming is involved. Admittedly, neither functions nor lists are the most efficient representation of environments. It would be nice to be able to implement environments using, say, balanced binary search trees, while preserving well-typedness. We leave this issue to future study.

The type system forces us to use names in a sound way. For instance, in the definition of `occ`, forgetting to extend the environment when crossing a

binder (that is, writing Γ instead of `extend Γ`) would cause a type error. In the definition of `extend`, attempting to check whether x occurs in Γ without first comparing x and b would cause a type error. Remember that in this nominal implementation scheme it is permitted for newer bindings to shadow earlier ones; our type discipline guarantees that the code works also in that case.

As suggested previously, one may wish to represent environments as an explicit data structure (an association list keyed by binders) rather than as an opaque object (a lookup function).

```
data DataEnv (A : Set) : ( $\alpha$   $\beta$  : World)  $\rightarrow$  Set where
   $\epsilon$  :  $\forall$  { $\beta$ }  $\rightarrow$  DataEnv A  $\beta$   $\beta$ 
   $-, \mapsto -$  :  $\forall$  { $\alpha$   $\beta$ } ( $\Gamma$  : DataEnv A  $\alpha$   $\beta$ )  $b$  ( $x$  : A)
     $\rightarrow$  DataEnv A ( $b < \alpha$ )  $\beta$ 

lookup :  $\forall$  {A  $\alpha$   $\beta$ }  $\rightarrow$  DataEnv A  $\alpha$   $\beta$   $\rightarrow$  Name  $\alpha$   $\rightarrow$  Name  $\beta$   $\uplus$  A
lookup  $\epsilon$  = inj1
lookup ( $\Gamma$  ,  $- \mapsto v$ ) = exportWith (inj2 v) (lookup  $\Gamma$ )
```

The type `DataEnv A α β` is the type of an environment, or environment fragment, where α is called the inner world, where β is called the outer world, and every name in the environment is associated with a datum of type A . The inner world is the world where all the binders of the environment are in scope. The outer world is the world where none of the binders are in scope. The expression `lookup Γ x` looks up the name x in the environment Γ . The name x must make sense in the scope of Γ , that is, x must inhabit the world α . If x is found among the bindings, then the information associated with x can be returned. If x is not found among the bindings, then x is returned, with a more precise type: indeed, since x is not among the names introduced by Γ , it must make sense outside Γ , that is, in the world β .

We illustrate the use of `DataEnv` with an alternative definition of the function `fv` (here the payload type parameter A is instantiated with the unit type \top). This variant avoids the need to take the bound atoms off the list by not inserting them in the first place. At variables, we use `lookup` to check whether the name is free or bound. If the name is free, we return it as a singleton list (using `[_]`). If it is bound, we ignore it and return an empty list. At every other node, we simply carry out a recursive traversal. Whenever a name abstraction is entered, the current environment Γ is extended with the bound name b .

```

fv' : ∀ {α β} → DataEnv ⊤ α β → Tm α → List (Name β)
fv' Γ (V x)      = [ [-] , const [] ]' (lookup Γ x)
fv' Γ (t · u)    = fv' Γ t ++ fv' Γ u
fv' Γ (λ b t)    = fv' (Γ , b ↦ -) t
fv' Γ (Let b t u) = fv' Γ t ++ fv' (Γ , b ↦ -) u

```

```

fv : ∀ {α} → Tm α → List (Name α)
fv = fv' ε

```

3.1.3 Example: Term comparison

Our next example focuses on comparison of terms. We first define `|Cmp|`, where `|Cmp| F` is the type of functions comparing `F`-structures. In our case the index type `Ix` is the type of worlds: `World`.

```

|Cmp| : ∀ {Ix} (F : Ix → Set) (i j : Ix) → Set
|Cmp| F i j = F i → F j → Bool

```

To compare two terms we carry down an environment, telling how to compare two names. At name occurrences we ask our environment, at application nodes we carry down our environment, and we need to extend the environment at name-abstractions. We extend the environment with the function `extendNameCmp` which receives a name comparator `f` for worlds α_1 and α_2 , a name x_1 in $b_1 \triangleleft \alpha_1$, and a name x_2 in $b_2 \triangleleft \alpha_2$. Then, the function `extendNameCmp` exports both names x_1 and x_2 simultaneously: when the two exports succeed we can use the comparator `f`, when both fail we learn that x_1 is b_1 and that x_2 is b_2 we hence return `true`, when one succeeds and the other fails we return `false`:

```

extendNameCmp : ∀ {α1 α2 b1 b2} → |Cmp| Name α1 α2
                                     → |Cmp| Name (b1 < α1) (b2 < α2)
extendNameCmp f x1 x2
  with exportM? x1 | exportM? x2
... | just x'1      | just x'2   = f x'1 x'2
... | nothing       | nothing     = true
... | -             | -           = false

```

```

cmpTm : ∀ {α1 α2} (Γ : |Cmp| Name α1 α2) → Tm α1 → Tm α2 → Bool

```

```

cmpTm Γ (V x1)      (V x2)      = Γ x1 x2
cmpTm Γ (t1 · u1)  (t2 · u2)  = cmpTm Γ t1 t2 ∧ cmpTm Γ u1 u2
cmpTm Γ (λ _ t1)   (λ _ t2)   = cmpTm (extendNameCmp Γ) t1 t2
cmpTm Γ (Let b1 t1 u1) (Let b2 t2 u2)
  = cmpTm Γ t1 t2 ∧ cmpTm (extendNameCmp Γ) u1 u2
cmpTm - -           -           = false

```

Notice that the code above does not need to compare neither names nor binders from the first term with those from the second term. Two bound names are said to be equal when they are bound at the same time. Bound names are said to be compared positionally. This also explains why we do not need a function to compare binders. Notice also that the function `cmpTm` has to accept two terms in different worlds to carry the recursion successfully. In the end we can give the function `_==N_` as an initial environment yielding an homogeneous comparison on terms, which coincides with α -equivalence:

```

_==Tm_ : ∀ {α} → Tm α → Tm α → Bool
_==Tm_ = cmpTm _==N_

```

In this final setting bound names are compared positionally and free names are compared for equality.

3.2 Kits and Traversals

We have seen that working with worlds requires explicitly moving names from world to world using operations like `coerceN` and `exportN?`. These operations quickly become necessary at the level of user-defined types like our type `Tm`. More generally any operation on names would benefit from being lifted to user-defined types.

From our experimentations with this library we outline two points. First, in most of the operations on terms we present, only the binding structure is relevant. The code for the traversal can hence be written only once. Second, the parts specific to each operation can be made reusable to work not only with the generic traversal but with custom traversals as well.

To share code and separate concerns we introduce some infrastructure. While abstract this infrastructure is then concretely applied to our example type `Tm`.

3.2.1 Traversal Kits

We first introduce the notion of *traversal kits*, which encompasses: a type of environment `Env α β` to be carried down in the traversal going from the world α to β ; a function `trName` to apply the environment at names; a function `trBinder` for binders; and finally a function `extEnv` to push

an environment under a binder. For more flexibility the result type `Res` of `trName` is a parameter as well, we shall see its use later on in this section. We build traversal kits for: coercing (applying a world inclusion witness), exporting, applying general effectful functions over names, freshening, and combinations of these.

```
record TrKit (Env : (α β : World) → Set)
  (Res : World → Set) : Set where
  constructor mk
  field
    trName    : ∀ {α β} → Env α β → Name α → Res β
    trBinder  : ∀ {α β} → Env α β → Binder → Binder
    extEnv    : ∀ {α β} b (Δ : Env α β)
               → Env (b < α) (trBinder Δ b < β)
```

3.2.2 Coercing kit

Our first kit, called `coerceKit`, is simple and to the point. The “environment” type is `_⊆_`. The result type is simply `Name`. The operation to call at names is `coerceN`. The action at binders is the identity, meaning that this kit does not perform any renaming or freshening. Finally the inclusion rule `⊆-<` fits the job of `extEnv`.

```
coerceKit : TrKit _⊆_ Name
coerceKit = mk coerceN (const id) ⊆-<
```

To illustrate the usage of `coerceKit`, we use it to lift `coerceN` from names to terms. In this situation the reuse is negligible, but this code serves a pedagogical purpose.

```
module CoerceTmWithCoerceKit where
  open TrKit coerceKit

  coeTm : ∀ {α β} → α ⊆ β → Tm α → Tm β
  coeTm Δ (V x)      = V (trName Δ x)
  coeTm Δ (t · u)    = coeTm Δ t · coeTm Δ u
  coeTm Δ (λ b t)    = λ (trBinder Δ b) (coeTm (extEnv b Δ) t)
  coeTm Δ (Let b t u) = Let (trBinder Δ b) (coeTm Δ t)
                          (coeTm (extEnv b Δ) u)
```

The function `coeTm` takes an inclusion witness and an input term. The inclusion witness is carried down the traversal, used at names with the function `trName` (equal to `coerceN` in the `coerceKit`) and extended at abstrac-

tions with the function `extEnv` (equal to \subseteq - in the `coerceKit`). The function `trBinder` is used to choose the binders of the output term which are the ones from the input in the `coerceKit`. This traversal is actually independent of the kit we use.

Our next kit requires the concept of name supplies. A name supply for the world α is just a pair of a binder called `seed` and a proof that `seed` is fresh for α :

```
record Supply  $\alpha$  : Set where
  constructor _,_
  field
    seedB : Binder
    seed# : seedB #  $\alpha$ 
```

We also introduce two helpers called `zeros` and `sucs`. `zeros` is the initial name supply and the function `sucs` increments both the seed and the fresh-for proof:

```
zeros : Supply  $\emptyset$ 
zeros = 0B , 0B # $\emptyset$ 

sucs :  $\forall \{\alpha\} \rightarrow (s : Supply \alpha) \rightarrow Supply (Supply.seedb s  $\triangleleft$   $\alpha$ )
sucs (seedB , seed#) = sucB seedB , suc# seed#$ 
```


3.2.3 Renaming kits

Our second kit, the “renaming kit”, is more involved and shows how to apply any function over names on the binding structure. To avoid captures we have to freshen the binders in the result and the system indeed imposes that we do so. We first define the record type `SubstEnv α β` . This record type holds: a function `trName` to apply at names and a name supply for β :

```
record SubstEnv (Res : World  $\rightarrow$  Set)  $\alpha$   $\beta$  : Set where
  constructor _,_
  field
    trName : Name  $\alpha$   $\rightarrow$  Res  $\beta$ 
    supply : Supply  $\beta$ 

open Supply supply public

-- each binder is translated to a fresh binder
trBinder : Binder  $\rightarrow$  Binder
trBinder _ = seedB
```

Figure 3.1: Lifting `trName`

The so-called `renameKit` is then defined by providing `trName`, `trBinder` and `extEnv`. This last one is the most involved part of the kit and depicted in figure 3.1. The function `trName` is lifted using `exportWith`. The lifted function `trName` takes a name, if this name is bound then the `seedB` is returned, otherwise this name (exported) is given to `trName` to produce a name which is then imported back using `coerceN`. This call to `coerceN` is valid only because `seedB` is known to be fresh for the output world. Here again `SubstEnv` is parameterized by the result type of `trName` for greater flexibility:

```

RenameEnv : (α β : World) → Set
RenameEnv = SubstEnv Name

renameKit : TrKit RenameEnv Name
renameKit = mk trName trBinder extEnv
  where
    open SubstEnv
    extEnv : ∀ {α β} b (Δ : RenameEnv α β)
              → RenameEnv (b < α) (- < β)
    extEnv - (trName      , (seedB      , seed#β))
            = (trName'   , (sucS (seedB , seed#β)))
    where
      trName' = exportWith
                (nameB seedB) -- bound
                (coerceN (⊆-# seed#β) ∘ trName) -- free

```

Our renaming kit works for any *total* function over names. However, to lift the function `exportN?` from names to terms, we need to deal with partial functions as well. We can actually build another renaming kit parameterized over a notion of effectful computation, namely an applicative functor. An applicative functor [McBride and Paterson, 2008] is halfway

between a functor and a monad. Like a monad, an applicative functor has a unit called `pure`. The function `pure` allows to embed any pure value as a potentially effectful one. The second operation, called `_*`, is an effectful application, taking an effectful function and argument and resulting in an effectful result. Here is how we can use applicative functors to apply an effectful function to a list:

```
module MapA {E} (E-app : Applicative E) where
  open Applicative E-app

  mapA : {A B : Set} → (A → E B) → List A → E (List B)
  mapA _ []          = pure []
  mapA f (x :: xs) = pure :-: * f x * mapA f xs
```

To build a more general renaming kit, we reuse the type `SubstEnv` but make use of the flexibility on the result type to plug in the effect type `E` (the applicative functor). The code for `renameAKit` is similar to `renameKit`, so we omit it here and provide only its type:

```
RenameAEnv : (E : Set → Set) (α β : World) → Set
RenameAEnv E = SubstEnv (E ∘ Name)

renameAKit : ∀ {E} → Applicative E →
              TrKit (RenameAEnv E) (E ∘ Name)
renameAKit = {! code similar to renameKit omitted !}
```

3.2.4 Substitution kit

We reuse our `SubstEnv` once more to generalize in another direction the renaming kit to the so-called `substKit` which helps build substitution functions. Substitution functions are built for a structure `F` (like our type `Tm`), provided there is a way of injecting names to `F`-structures and of coercing an `F`-structure from one world to another. Here again we choose another result type for `SubstEnv`, namely `F`.

```
-- Index-respecting functions
F |→| G = ∀ {i} → F i → G i

-- The type for 'coerce' on an F-structure
Coerce F = ∀ {α β} → α ⊆ β → F α → F β

substKit : ∀ {F}
```

```

      (V      : Name |→| F)
      (coerceF : Coerce F)
      → TrKit (SubstEnv F) F
substKit = {! code similar to renameKit omitted !}

```

3.2.5 Other kits and combinators

We can also build others kits and combinators [Pouillard, 2011a]. For instance `o-Kit` can compose two kits by pairing the two environments and composing their operations. Another combinator, `starKit`, takes a kit `k` on environments of type `Env` and builds a kit working on `Star Env` (the reflexive and transitive closure operator in AGDA). Finally with `mapKit` one can pre-compose and post-compose a function with a kit to get a new kit. This last one is given below:

```

mapKit : ∀ {Env F G} (f : Name |→| Name) (g : F |→| G)
        → TrKit Env F → TrKit Env G
mapKit f g kit = mk (λ Δ → g ∘ trName Δ ∘ f) trBinder extEnv
  where open TrKit kit

```

3.2.6 Reusable traversal

Putting all these kits to work is a matter of writing a traversal function working for any kit and any effect. It is essentially a map function where the free names are transformed by a user-supplied function. More precisely the function `trTm` traverses the term, carrying an environment. It uses `trName` at variable occurrences without putting back the constructor `V`. In every other case, it just rebuilds the same structure using the corresponding constructor and the operations of our applicative functor. To carry the environment under bindings `extEnv` is used:

```

module TraverseTm {E} (E-app : Applicative E)
                    {Env} (trKit : TrKit Env (E o Tm)) where

open Applicative E-app
open TrKit trKit

trTm :  $\forall \{\alpha \beta\} \rightarrow \text{Env } \alpha \beta \rightarrow \text{Tm } \alpha \rightarrow \text{E } (\text{Tm } \beta)$ 
trTm  $\Delta$  (V x)      = trName  $\Delta$  x
trTm  $\Delta$  (t · u)    = pure _·_  $\otimes$  trTm  $\Delta$  t  $\otimes$  trTm  $\Delta$  u
trTm  $\Delta$  ( $\lambda$  b t)  = pure ( $\lambda$  _)  $\otimes$  trTm (extEnv b  $\Delta$ ) t
trTm  $\Delta$  (Let b t u) = pure (Let _)  $\otimes$  trTm  $\Delta$  t
                                      $\otimes$  trTm (extEnv b  $\Delta$ ) u

```

3.2.7 Reusing the traversal

We can now collect the fruit of our work, by combining the traversal and a few kits. For the sake of simplicity we directly do so on the type `Tm`, but in our development [Pouillard, 2011a] we further abstract over `Tm` (as `F`) and `trTm` by defining a sequence of parameterized modules.

```

open TraverseTm

-- Like 'trTm' but restores the constructor 'V'.
-- The result type of 'trKit' is refined.
trTm' :  $\forall \{E\}$  (E-app : Applicative E)
        {Env} (trKit : TrKit Env (E o Name))
         $\{\alpha \beta\} \rightarrow \text{Env } \alpha \beta \rightarrow \text{Tm } \alpha \rightarrow \text{E } (\text{Tm } \beta)$ 
trTm' E-app trKit
  = trTm E-app (mapKit id (Applicative.<$> E-app V) trKit)

```

While `coerceN` can be lifted to terms using a renaming function, this would impose freshening. Using `trTm'` with the `coerceKit` yields a more efficient implementation.

```

-- The identity on types froms an applicative functor
id-app : Applicative id
id-app = {! definition omitted !}

coerceTm :  $\forall \{\alpha \beta\} \rightarrow \alpha \subseteq \beta \rightarrow \text{Tm } \alpha \rightarrow \text{Tm } \beta$ 
coerceTm = trTm' id-app coerceKit

```

Building the renaming function amounts to use one of our renaming kits

and to pick one applicative functor:

```

renameTmA : ∀ {E} → Applicative E →
            ∀ {α β} → Supply β → (Name α → E (Name β))
                               → (Tm α → E (Tm β))
renameTmA E-app s f = trTm' E-app (renameAKit E-app) (f , s)

renameTm : ∀ {α β} → Supply β → (Name α → Name β)
                               → (Tm α → Tm β)
renameTm s f = trTm' id-app renameKit (f , s)

renameTm? : ∀ {α β} → Supply β → (Name α →? Name β)
                               → (Tm α →? Tm β)
renameTm? = renameTmA Maybe.applicative

```

Exporting is built by lifting the function `exportN?` using the function `renameTm?`:

```

exportTm? : ∀ {b α} → Supply α → Tm (b < α) →? Tm α
exportTm? s = renameTm? s exportN?

```

Another special case of `renameTm?` is the so-called `closeTm?`. This function takes a term in any world and checks if the term is closed. If so, the same term is returned in the empty world. Otherwise the function fails by returning `nothing`:

```

closeTm? : ∀ {α} → Tm α →? Tm ∅
closeTm? = renameTm? (0B , (0B)#∅) (const nothing)

```

Building capture avoiding substitution amounts to use the substitution kit with the constructor `V` and the function `coerceTm` as arguments:

```

substTm : ∀ {α β} → Supply β → (Name α → Tm β) → Tm α → Tm β
substTm (s , s#) f = trTm id-app (substKit V coerceTm) (f , s , s#)

```

To illustrate the use of `substTm`, here is a simple function `β-red` which performs a β -reduction when a β -redex appears at the root of the term. The function `exportWith a V` associate `a` to `b` and `x` to `V x` when `x ≠ b`:

```

β-red : ∀ {α} → Supply α → Tm α → Tm α
β-red s (λ b f · a) = substTm s (exportWith a V) f
β-red _ t           = t

```

3.2.8 Building any λ -term

A way to illustrate that every λ -term can be encoded using our type `Tm` is to define a conversion function from another type for λ -terms to the type `Tm`. We do so by choosing the type `TmA` from section 2.1.1 as the source language.

The process is very close to the combination of a specific renaming kit and traversal function. The kit is specific since the source names are of type `Atom` and not `Name`. The traversal function is specific since the source and target types are not the same and we picked the identity functor for simplicity.

First, we introduce the environment type which holds a mapping from free atoms to free names and a name supply:

```
module Conv-TmA→Tm where
  record Env  $\alpha$  : Set where
    constructor --,--
    field
      trAtom : Atom → Name  $\alpha$ 
      supply : Supply  $\alpha$ 
    open Supply supply public
  open Env
```

We then define how an environment is extended. It works similarly to the one for the renaming kit:

```
extEnv :  $\forall \{\alpha\} \rightarrow$  Atom  $\rightarrow$  ( $\Delta$  : Env  $\alpha$ )  $\rightarrow$  Env (seedB  $\Delta$   $\triangleleft$   $\alpha$ )
extEnv bA  $\Delta$  = trN, sucS (supply  $\Delta$ )
  where trN =  $\lambda$  xA  $\rightarrow$  if bA ==A xA
    then nameB (seedB  $\Delta$ )
    else coerceN ( $\subseteq$ -# (seed#  $\Delta$ )) (trAtom  $\Delta$  xA)
```

The following function `conv` is then straightforward, calling `trAtom` on atoms at the variable case and using `extEnv` when crossing a binding. This function `conv` is simpler than the traversal function seen previously since we avoided the use of effectful computations, however the applicative version works as well. The function `conv` can then be used if we can associate all atoms to names and have the corresponding name supply.

```
conv :  $\forall \{\alpha\} \rightarrow$  Env  $\alpha$   $\rightarrow$  TmA  $\rightarrow$  Tm  $\alpha$ 
conv  $\Delta$  (V x)      = V (trAtom  $\Delta$  x)
conv  $\Delta$  ( $\lambda$  b t)  =  $\lambda$  _ (conv (extEnv b  $\Delta$ ) t)
conv  $\Delta$  (t · u)   = conv  $\Delta$  t · conv  $\Delta$  u
conv  $\Delta$  (Let b t u) = Let _ (conv  $\Delta$  t) (conv (extEnv b  $\Delta$ ) u)
```


To convert closed terms we need an empty environment. However the type Tm^A does not ensure closedness. Hence the conversion is partial. To build this partial conversion we define it in two steps. The first step uses the function `conv` with an empty environment. This empty environment maps any atom to the name 0^N , the name supply hence starts at 1^N . As the second step we use the function `closeTm?` which strengthens the type of the input term if it is closed and fails otherwise.

```
emptyEnv : Env (0B < ∅)
emptyEnv = const (0N) , sucs zeros

conv∅? : TmA →? Tm ∅
conv∅? = closeTm? ∘ conv emptyEnv
```

3.3 Towards elaborate uses of worlds

3.3.1 Data type of “one hole contexts”

The type Tm is just one basic example of an algebraic data type that involves names and binders. As a more challenging example, consider a type C of one-hole contexts associated with Tm . The type C is indexed with two worlds, which respectively play the role of an outer world and an inner world. The idea is, plugging a term of type $\text{Tm } \beta$ into the hole of a context of type $\text{C } \alpha \beta$ produces a term of type $\text{Tm } \alpha$. The definition of the type C is as follows:

```
data C α : World → Set where
  Hole   : C α α
  ·1-    : ∀ {β} → C α β → Tm α → C α β
  ·2-    : ∀ {β} → Tm α → C α β → C α β
  λ      : ∀ {β} b → C (b < α) β → C α β
  Let1  : ∀ {β} b → C α β
           → Tm (b < α) → C α β
  Let2  : ∀ {β} b → Tm α
           → C (b < α) β → C α β
```

Contexts bind names: the hole can appear under one or several binders. This is why, in general, a context has distinct outer and inner worlds. A context contains a list of binders that “connects” the outer and inner worlds: these binders are carried by the constructors λ .

A context and a term can be paired to produce a term-in-context:

```
CTm : World → Set
CTm α = ∃[ β ](C α β × Tm β)
```

It is straightforward to define a function `plug` from $\mathbf{CTm} \alpha$ to $\mathbf{Tm} \alpha$, which accepts a pair of a context and a term and plugs the latter into the former. Conversely, one can define a family of focusing functions ($\forall \{\alpha\} \rightarrow \mathbf{Tm} \alpha \rightarrow \mathbf{CTm} \alpha$), which split a term into a pair of a context and a term. There are several such functions, according to where one wishes to focus.

3.3.2 Patterns a la ML

As another instance of this idea, if one wished to extend our object language with ML-style patterns, one would index the type `Pat` of patterns with an outer world and an inner world, and one would use elaborate abstractions of the form $\exists [\beta] (\mathbf{Pat} \alpha \beta \times \mathbf{Tm} \beta)$.

3.3.3 Term and types: System F

We use Girard-Reynold's System F as an object language to illustrate how representations of object-level types can be indexed by worlds as well. Object-level terms are now indexed by two worlds: one for type names and one for term names.

```

module SysF where
  infixr 5 _=>_
  data Ty α : Set where
    V      : (x : Name α)           → Ty α
    _=>_   : (σ τ : Ty α)           → Ty α
    '∀'    : ∀ b (τ : Ty (b < α)) → Ty α

  data Tm α γ : Set where
    V      : ∀ (x : Name α)         → Tm α γ
    _·_    : ∀ (t u : Tm α γ)       → Tm α γ
    λ      : ∀ b (τ : Ty γ)
              (t : Tm (b < α) γ)    → Tm α γ
    _·τ_   : ∀ (t : Tm α γ) (τ : Ty γ) → Tm α γ
    Λ      : ∀ b (t : Tm α (b < γ)) → Tm α γ

```

3.4 Advanced example: normalization by evaluation

As an advanced example, we show how to express a normalization by evaluation algorithm in our system. This algorithm has been previously used as a benchmark by several researchers [Shinwell et al., 2003, Pitts, 2006, Licata and Harper, 2009, Cave and Pientka, 2012]. The challenge lies in the way

the algorithm mixes computational functions, name abstractions, and fresh name generation.

The object language of interest is again the pure λ -calculus. The algorithm exploits two different representations of object-level terms, which are respectively known as *syntactic* and *semantic* representations. Because these representations differ only in their treatment of name abstractions, they can be given a common definition, which is parameterized over the representation of abstractions:

```

module M (Abs : (World → Set) → World → Set) where
  data T α : Set where
    V      : Name α      → T α
    λ      : Abs T α     → T α
    _·_    : T α → T α → T α

```

The parameter `Abs` has kind $(\text{World} \rightarrow \text{Set}) \rightarrow (\text{World} \rightarrow \text{Set})$: it is an indexed-type transformer.

In order to obtain the syntactic representation, we instantiate `Abs` with the nominal abstractions that we have used everywhere so far: an abstraction is a package of a binder and of a term that inhabits the inner world. This yields the type `Term` of syntactic terms.

```

SynAbsN : (World → Set) → World → Set
SynAbsN F α = ∃[ b ] (F (b < α))

```

```

open M SynAbsN renaming (T to Term)

```

In order to obtain the semantic representation, we instantiate `Abs` with a different notion of abstraction, in the style of higher-order abstract syntax: an abstraction is a computational function, which substitutes a term for the bound name of the abstraction. This yields the type `Sem` of semantic terms. The type `Sem` is not an inductive data type; fortunately, with the `--no-positivity-check` flag, AGDA accepts this type definition, at the cost of breaking strong normalization (to minimize the risks we isolated this kind of example to separate modules where this flag is activated).

```

SemAbs : (World → Set) → World → Set
SemAbs F α = ∀ {β} → α ⊆ β → F β → F β

```

```

open M SemAbs renaming (T to Sem)

```

It is important to note that our semantic name abstractions involve bounded polymorphism in a world: we define `SemAbs F α` as $\forall\{\beta\} \rightarrow \alpha \subseteq \beta \rightarrow F \beta \rightarrow F \beta$, as opposed to the more naïve $F \alpha \rightarrow F \alpha$. This

3.4. ADVANCED EXAMPLE: NORMALIZATION BY EVALUATION 75

provides a more accurate and more flexible description of the behavior of substitution. Indeed, when instantiating an abstraction \mathbf{t} with some term \mathbf{u} , it makes perfect sense for \mathbf{u} to inhabit a larger world than \mathbf{t} , that is, for \mathbf{u} to refer to certain names that are fresh for \mathbf{t} . The result of the substitution then inhabits the same world as \mathbf{u} : that is, it potentially refers to these fresh names, in addition to all of the names that occurred free in the abstraction \mathbf{t} .

The types `SemAbs` (and `Sem`), are covariant with respect to the parameter α , which would not be the case with the naïve definition. In other words, it is possible to define a coerce operation for semantic terms:

```
coerceSem : ∀ {α β} → α ⊆ β → Sem α → Sem β
coerceSem pf (V a)    = V (coerceN pf a)
coerceSem pf (λ f)    = λ (λ pf' v → f (⊆-trans pf pf') v)
coerceSem pf (t · u)  = coerceSem pf t · coerceSem pf u
```

At a semantic abstraction, no recursive call is performed, because the body of the abstraction is opaque: it is a computational function. Instead, we exploit the transitivity of world inclusion and build a new semantic abstraction that inhabits the desired world.

The normalization by evaluation algorithm makes use of environments. Here, environments are functions from names to semantic terms. The function `_, ↦_` shows how to extend such an environment:

```
EvalEnv : (α β : World) → Set
EvalEnv α β = Name α → Sem β
-- α is the inner world
-- β is the outer world

_, ↦_ : ∀ {α β} (Γ : EvalEnv α β) b → Sem β → EvalEnv (b < α) β
_, ↦_ Γ b v = exportWith v Γ
-- meaning: b ↦ v
--           x ↦ Γ x
```

To a name, of type `Name α` an environment, of type `EvalEnv α β` associates a semantic term that lies outside the scope of the environment that is, a semantic term of type `Sem β`. The type `EvalEnv α β` is, again, covariant in its second world, as witnessed by the following coercion function:

```
coerceEnv : ∀ {α β1 β2} → β1 ⊆ β2 → EvalEnv α β1 → EvalEnv α β2
coerceEnv pf Γ = coerceSem pf ∘ Γ
```

The first part of the normalization by evaluation algorithm is a function `eval` that evaluates a syntactic term within an environment to produce a semantic term. When evaluating a λ -abstraction, we build a semantic ab-

straction, which encapsulates a recursive call to `eval`. The bounded polymorphism required by the definition of semantic abstractions forces us to coerce the environment Γ via `coerceEnv`.

```

eval :  $\forall \{ \alpha \beta \} \rightarrow \text{EvalEnv } \alpha \beta \rightarrow \text{Term } \alpha \rightarrow \text{Sem } \beta$ 
eval  $\Gamma$  ( $\lambda$  (a , t))
  =  $\lambda$  ( $\lambda$  pf v  $\rightarrow$  eval (coerceEnv pf  $\Gamma$  , a  $\mapsto$  v) t)
eval  $\Gamma$  (V x)          =  $\Gamma$  x
eval  $\Gamma$  (t . u)      = app (eval  $\Gamma$  t) (eval  $\Gamma$  u) where
  app :  $\forall \{ \alpha \} \rightarrow \text{Sem } \alpha \rightarrow \text{Sem } \alpha \rightarrow \text{Sem } \alpha$ 
  app ( $\lambda$  f) v = f  $\subseteq$ -refl v
  app n       v = n . v

```

The second part of the algorithm reifies a semantic term back into a syntactic term. When reifying a semantic abstraction, we build a syntactic abstraction. This requires generating a fresh name, and leads us to parameterizing `reify` with a supply of fresh names.

```

reify :  $\forall \{ \alpha \} \rightarrow \text{Supply } \alpha \rightarrow \text{Sem } \alpha \rightarrow \text{Term } \alpha$ 
reify s      (V a)    = V a
reify s      (n . v)  = reify s n . reify s v
reify (sB , s#) ( $\lambda$  f) =
   $\lambda$  (sB , reify (sucS (sB , s#)) (f ( $\subseteq$ -# s#) (V (nameB sB))))

```

The constructor `V` has type `Name $\alpha \rightarrow \text{Sem } \alpha$` hence is a valid initial environment of type `EvalEnv $\alpha \alpha$` . Evaluation under this initial environment, followed with reification, yields a normalization algorithm. This algorithm works with open terms: its argument, as well as its result, are terms in an arbitrary world α . However this function needs a name supply for the world α . The name supply can be incrementally constructed with `zeroS` and `sucS`. In particular we can normalize closed terms by giving `zeroS`:

```

nf :  $\forall \{ \alpha \} \rightarrow \text{Supply } \alpha \rightarrow \text{Term } \alpha \rightarrow \text{Term } \alpha$ 
nf supply = reify supply  $\circ$  eval V

```

Here is an example of the normalization of a term:

```

idT : Term  $\emptyset$ 
idT =  $\lambda$  (0B , V (0N))

test-nf : nf zeroS (((idT . (idT . idT)) . idT)  $\equiv$  idT
test-nf = refl

```

Chapter 4

Behind the scene of the NomPa library

4.1 Implementation of NomPa

The implementation of our interface is not surprising. Most of the code is types or proofs. Notice that this section details the internals of our library and that none of these definitions are meant to be known or used by the library client.

Worlds come first and are represented by lists of booleans. Such a list means that n is in the world if and only if the n^{th} value of the list is `true`. Formally the meaning of a world is defined by the membership predicate below. The choice of a list of booleans to represent a world was guided by two facts. First, the operations are structural: this eases the type-level computation and the proof goes through in particular in combination with de Bruijn indices as we details in section 5.2. Second, the elements are kept in order and (modulo trailing occurrences of `false`) two equivalent sets are represented the same way.

```
World : Set
World = List Bool

∅ : World
∅ = []

_∈_ : ℕ → World → Set
-    ∈ []           = ⊥
zero ∈ (false :: -) = ⊥
zero ∈ (true  :: -)  = ⊤
suc n ∈ (-       :: xs) = n ∈ xs
```

Worlds are meant to be computationally irrelevant. Meaning that none of the functions of this implementation make use of the worlds to compute the results. Hence, programs have an equivalent dynamic semantics when worlds are erased. While worlds are indeed used in membership, inclusion, and fresh-for proofs, these proofs are meant to be erased as well. Finally we want this erasing for three reasons: first for performances in the context of an optimizing compiler, second to have a separation between the “type system” and the running program, third to gain interesting theorems for free through the parametricity of world polymorphic functions (section 4.2).

Then binders are represented by natural numbers. The operator `_<_` defines how to extend a world with a binder. Given a binder `n` and a world `α`, `n < α` updates the world `α` with the value true at index `n`.

```

Binder : Set
Binder = ℕ

zeroB : Binder
zeroB = zero

sucB : Binder → Binder
sucB = suc

_<_ : Binder → World → World
zero < []      = true  :: []
suc n < []     = false :: n < []
zero < (- :: α) = true  :: α
suc n < (b :: α) = b    :: n < α

infixr 5 _<_

```

The semantics of `(y < α)` is to add `y` to the set `α`. Precisely the semantics is described by the following lemma:

```

<-sem : ∀ α x y → (x ∈ y < α) ≡ (if x ==ℕ y then ⊤ else x ∈ α)

```

Names are simply represented by a pair of a binder (a number) and a proof that this binder belongs to the indexing world. In AGDA we use a record:

```

record Name  $\alpha$  : Set where
  constructor -, -
  field
    binderN : Binder
    b $\in\alpha$  : binderN  $\in$   $\alpha$ 
infixr 4 -, -
open Name public

```

To produce a name, `nameB` simply packs the given binder and a proof which is always constructible in this simple case:

```

nameB :  $\forall$  { $\alpha$ } b  $\rightarrow$  Name (b  $\triangleleft$   $\alpha$ )
nameB b = b , {! proof omitted !}

```

The equality test `==N` and the exporting function `exportN?` both simply compare the internal representation of names and binders.

```

==N :  $\forall$  { $\alpha$ } (x y : Name  $\alpha$ )  $\rightarrow$  Bool
==N (x , -) (y , -) = x ==N y

exportN? :  $\forall$  {b  $\alpha$ }  $\rightarrow$  Name (b  $\triangleleft$   $\alpha$ )  $\rightarrow?$  Name  $\alpha$ 
exportN? {b} { $\alpha$ } (x , pf) =
  if x ==N b then nothing
  else just (x , {! proof omitted !})

```

World inclusion is modeled as set-theoretic inclusion (membership preservation). This membership preservation proof is used by `coerceN` to build the expected proof. All the inclusion rules (see figure 2.1) are computationally irrelevant and their implementations are omitted here.

```

infix 2  $\subseteq$ 
 $\subseteq$  : ( $\alpha$   $\beta$  : World)  $\rightarrow$  Set
 $\alpha \subseteq \beta$  =  $\forall$  x  $\rightarrow$  x  $\in$   $\alpha$   $\rightarrow$  x  $\in$   $\beta$ 

coerceN :  $\forall$  { $\alpha$   $\beta$ }  $\rightarrow$   $\alpha \subseteq \beta$   $\rightarrow$  Name  $\alpha$   $\rightarrow$  Name  $\beta$ 
coerceN  $\alpha \subseteq \beta$  (b , b $\in\alpha$ ) = b ,  $\alpha \subseteq \beta$  b b $\in\alpha$ 

```

The remaining part is the fresh-for relation (`#`). To cope with the proof of `suc#`, we give two characterizations of this relation. One is a set of syntactic rules (omitted here) and the other is semantic. These presentations are equivalent [Pouillard, 2011a]. For pragmatic reasons the fresh-for relation is stronger than the non-membership relation (that is, `x # α` implies `x \notin α` but the converse is false). The definition requires binder the be strictly greater than any binder in the world. This allows to build the `suc#`

primitive rule (omitted here) without inspecting the world. Not inspecting the worlds at run-time is necessary since we want to erase worlds.

```

_#_ : Binder → World → Set
x # α = ∀ y → y ∈ α → x > y

```

4.2 Soundness: logical relations and parametricity

Since our library is written in a type-safe language, one may wonder what soundness properties are to be proved. Moreover our names are indexed by worlds and hold membership proofs. The functions over names are given precise types and have been shown to be type-safe. The property that “well-typed programs do not go wrong” comes for free, but does not satisfy us. Indeed we have explained that certain operations (such as comparing two binders) must not be provided to programmers; the logical relation explains why this is the case.

To show that our library respects a model of good behavior with respect to names and binders, we build a model based on logical relations indexed/directed by types. This technique [Bernardy et al., 2010] is independent of this work and enables to define a notion of program equivalence. We use this technique to capture good behavior of functions involving names and binders. Using this technique, the set of specific definitions is kept to a minimum of one per introduced type (`World`, `Name`, and `_⊆_`). One proof per value introduced has to be done, which keeps the development modular and forward-compatible to the addition of more features.

This section is organized as follows. First we recall the basics of logical relations and parametricity. Then we give a toy example to practice a bit. Then relations for worlds, names, and world inclusions are given. Finally we make use of the construction to obtain free theorems [Wadler, 1989] about world-polymorphic functions over the type `Tm`.

4.2.1 Recap of the framework

A relation is said to be type-directed when it is recursively defined on the structure of types. Let \mathcal{R} be such a type-directed relation, and let τ be a type. Then, \mathcal{R}_τ is a relation on values of type τ , namely $\mathcal{R}_\tau : \tau \rightarrow \tau \rightarrow \text{Set}$. Recall that `Set` is also the type of propositions in AGDA.

A type-directed relation is called “logical” when the case for functions is defined extensionally. Here “extensionally” means that two functions are related when they produce related results out of related arguments. Let A_τ be a relation for the arguments and B_τ a relation for results. Two functions f_1 and f_2 are related if and only if for every pair of arguments (x_1, x_2) related

by A_r , the results $f_1 \ x_1$ and $f_2 \ x_2$ are related by B_r . This definition can be given in AGDA as well:

```
RelatedFunctions A_r B_r f_1 f_2 =
  ∀ {x_1 x_2} → A_r x_1 x_2
    → B_r (f_1 x_1) (f_2 x_2)
```

We say that a program or a value fits a logical relation when it is related to itself by the relation indexed by its type. We say that a logical relation is universal if every well-typed program fits the logical relation. This notion of universality was originally coined by John Reynolds as the “Abstraction Theorem” [Reynolds, 1983]. We call the “AGDA logical relation” the one defined by Bernardy et al. [2010] for a PTS (pure type system) and naturally extended as suggested to other features of AGDA. While no complete mechanized proof has been done for this we consider the AGDA logical relation as universal.

To simplify matters, the definitions shown here are not universe-polymorphic. The reader can find universe-polymorphic definitions in our complete AGDA development [Pouillard, 2011a].

To formally define a logical relation indexed by types, a common technique is to first inductively define the structure of types. This is known as a “universe of codes” \mathbf{U} . Then one defines a function called `E1` from codes to types. Finally one defines by induction a function called `[[_]]` from codes to relations on elements of type described by the given code. In AGDA the function `[[_]]` has the following type: $(\tau : \mathbf{U}) \rightarrow \mathbf{E1} \ \tau \rightarrow \mathbf{E1} \ \tau \rightarrow \mathbf{Set}$. Because of quantification and dependent-types, types contain variables, and a good deal of complexity is added to this scheme, so we opt for a lighter scheme. We do not define \mathbf{U} , `E1`, and `[[_]]`.

Instead, for each type constructor κ , we define a relation `[[κ]]` (this is the name of the relation, since there are no spaces). For the function type constructor the `RelatedFunctions` definition above is a good start. Actually this is a fine definition for non-dependent functions. The dependent version of `RelatedFunctions`, called `[[II]]` here, passes the relation argument $A_r \ x_1 \ x_2$ called x_r to the relation for results B_r . In short the relation for results now depends on the relation for arguments. Here is the definition in AGDA:

```
[[II]] A_r B_r f_1 f_2 = ∀ {x_1 x_2} (x_r : A_r x_1 x_2)
  → B_r x_r (f_1 x_1) (f_2 x_2)
```

Note that this definition generalizes the case of non-dependent functions and universal quantifications as well. For non-dependent functions we simply provide a relation for results which ignores its first argument (equivalent to `RelatedFunctions` and noted `[[→]]` from now on). For universal quan-

tifications, since the arguments are types, all we need is a relation for types (members of Set_0) themselves. Following our convention we call this relation $\llbracket \text{Set}_0 \rrbracket$. The relation $\llbracket \text{Set}_0 \rrbracket A_1 A_2$ is the set of relations between A_1 and A_2 :

$$\begin{aligned} \llbracket \text{Set}_0 \rrbracket & : \text{Set}_0 \rightarrow \text{Set}_0 \rightarrow \text{Set}_1 \\ \llbracket \text{Set}_0 \rrbracket A_1 A_2 & = A_1 \rightarrow A_2 \rightarrow \text{Set}_0 \end{aligned}$$

For reference, full definitions for core type theory are given in figure 4.1. As we said, we do not define the function $\llbracket - \rrbracket$. However we want a notation close to what this function would provide. To this end, instead of writing $\llbracket \tau \rrbracket$ for a given type expression τ , we write τ where we replaced each constructor κ by $\llbracket \kappa \rrbracket$, each non dependent arrow by $\llbracket \rightarrow \rrbracket$, each dependent arrow $(x : A) \rightarrow B$ by $\langle x_r : \llbracket A \rrbracket \rangle \llbracket \rightarrow \rrbracket \llbracket B \rrbracket$. By convention we subscript the names by r . Applications are translated to applications. Because of dependent types, this translation has to be extended to all terms but we do not do it here. Finally here are a few examples of the manual use of the function $\llbracket - \rrbracket$:

-- What we would like to write but cannot:

$\llbracket \mathbb{N} \rightarrow \mathbb{N} \rightarrow \text{Bool} \rrbracket =$

-- What we write instead:

$\llbracket \mathbb{N} \rrbracket \llbracket \rightarrow \rrbracket \llbracket \mathbb{N} \rrbracket \llbracket \rightarrow \rrbracket \llbracket \text{Bool} \rrbracket =$

-- What this means:

$\lambda f_1 f_2 \rightarrow$
 $\quad \forall \{x_1 x_2\} (x_r : \llbracket \mathbb{N} \rrbracket x_1 x_2)$
 $\quad \quad \{y_1 y_2\} (y_r : \llbracket \mathbb{N} \rrbracket y_1 y_2)$
 $\quad \rightarrow \llbracket \text{Bool} \rrbracket (f_1 x_1 y_1) (f_2 x_2 y_2)$

$\llbracket (A : \text{Set}_0) \rightarrow A \rightarrow A \rrbracket =$
 $\llbracket \text{II} \rrbracket \llbracket \text{Set}_0 \rrbracket (\lambda A_r \rightarrow A_r \llbracket \rightarrow \rrbracket A_r) =$

$\lambda f_1 f_2 \rightarrow$
 $\quad \forall \{A_1 A_2\} (A_r : A_1 \rightarrow A_2 \rightarrow \text{Set}_0)$
 $\quad \quad \{x_1 x_2\} (x_r : A_r x_1 x_2)$
 $\quad \rightarrow A_r (f_1 A_1 x_1) (f_2 A_2 x_2)$

-- Using the notation instead of $\llbracket \text{II} \rrbracket$:

$\llbracket (A : \text{Set}_0) \rightarrow \text{List } A \rrbracket =$
 $\langle A_r : \llbracket \text{Set}_0 \rrbracket \rangle \llbracket \rightarrow \rrbracket \llbracket \text{List} \rrbracket A_r =$

$\lambda l_1 l_2 \rightarrow$
 $\quad \forall \{A_1 A_2\} (A_r : A_1 \rightarrow A_2 \rightarrow \text{Set}_0)$
 $\quad \rightarrow \llbracket \text{List} \rrbracket A_r (l_1 A_1) (l_2 A_2)$

We now have the definition of the AGDA logical relation for the core

$$\begin{aligned}
\llbracket \text{Set}_0 \rrbracket & : \forall (A_1 A_2 : \text{Set}_0) \rightarrow \text{Set}_1 \\
\llbracket \text{Set}_0 \rrbracket A_1 A_2 & = A_1 \rightarrow A_2 \rightarrow \text{Set}_0 \\
\llbracket \text{Set}_1 \rrbracket & : \forall (A_1 A_2 : \text{Set}_1) \rightarrow \text{Set}_2 \\
\llbracket \text{Set}_1 \rrbracket A_1 A_2 & = A_1 \rightarrow A_2 \rightarrow \text{Set}_1 \\
-\llbracket \rightarrow \rrbracket - & : \forall \{A_1 A_2 B_1 B_2\} \rightarrow \llbracket \text{Set}_0 \rrbracket A_1 A_2 \rightarrow \llbracket \text{Set}_0 \rrbracket B_1 B_2 \\
& \quad \rightarrow \llbracket \text{Set}_0 \rrbracket (A_1 \rightarrow B_1) (A_2 \rightarrow B_2) \\
A_r \llbracket \rightarrow \rrbracket B_r & = \lambda f_1 f_2 \rightarrow \forall \{x_1 x_2\} \rightarrow A_r x_1 x_2 \rightarrow B_r (f_1 x_1) (f_2 x_2) \\
\text{infixr } 0 & -\llbracket \rightarrow \rrbracket - \\
\llbracket \text{II} \rrbracket & : \forall \{A_1 A_2\} (A_r : \llbracket \text{Set}_0 \rrbracket A_1 A_2) \\
& \quad \{B_1 B_2\} (B_r : (A_r \llbracket \rightarrow \rrbracket \llbracket \text{Set}_0 \rrbracket) B_1 B_2) \\
& \quad \rightarrow ((x : A_1) \rightarrow B_1 x) \rightarrow ((x : A_2) \rightarrow B_2 x) \rightarrow \text{Set}_1 \\
\llbracket \text{II} \rrbracket A_r B_r & = \lambda f_1 f_2 \rightarrow \forall \{x_1 x_2\} (x_r : A_r x_1 x_2) \rightarrow B_r x_r (f_1 x_1) (f_2 x_2) \\
\text{syntax } \llbracket \text{II} \rrbracket A_r (\lambda x_r \rightarrow f) & = \langle x_r : A_r \rangle \llbracket \rightarrow \rrbracket f \\
\llbracket \forall \rrbracket & : \forall \{A_1 A_2\} (A_r : \llbracket \text{Set}_0 \rrbracket A_1 A_2) \\
& \quad \{B_1 B_2\} (B_r : (\llbracket \text{Set}_0 \rrbracket \llbracket \rightarrow \rrbracket \llbracket \text{Set}_0 \rrbracket) B_1 B_2) \\
& \quad \rightarrow \llbracket \text{Set}_1 \rrbracket (\{x : A_1\} \rightarrow B_1 x) (\{x : A_2\} \rightarrow B_2 x) \\
\llbracket \forall \rrbracket A_r B_r & = \lambda f_1 f_2 \rightarrow \forall \{x_1 x_2\} (x_r : A_r x_1 x_2) \rightarrow B_r x_r (f_1 \{x_1\}) (f_2 \{x_2\}) \\
\text{syntax } \llbracket \forall \rrbracket A_r (\lambda x_r \rightarrow f) & = \forall \langle x_r : A_r \rangle \llbracket \rightarrow \rrbracket f
\end{aligned}$$

Figure 4.1: Logical relations for core types

```

data [[⊥]] : [[Set0]] ⊥ ⊥ where -- no constructors

data [[Bool]] : [[Set0]] Bool Bool where
  [[true]]  : [[Bool]] true  true
  [[false]] : [[Bool]] false false

data [[N]] : [[Set0]] N N where
  [[zero]]  : [[N]]          zero  zero
  [[suc]]   : ([[N]] [[→]] [[N]]) suc  suc

data -[[⊔]]- {A1 A2 B1 B2} (Ar : [[Set0]] A1 A2)
  (Br : [[Set0]] B1 B2) :
  A1 ⊔ B1 → A2 ⊔ B2 → Set0 where
  [[inj1]] : (Ar [[→]] Ar [[⊔]] Br) inj1 inj1
  [[inj2]] : (Br [[→]] Ar [[⊔]] Br) inj2 inj2

```

Figure 4.2: Logical relations for data types

type theory part. It extends nicely to inductive data types and records. The process is as follows: for each constructor κ of type τ , declare a new constructor $[[\kappa]]$ whose type is $[[\tau]] \kappa \kappa$. This process applies to type constructors and data constructors of data types, and type constructors and fields of record types. For reference, the logical relations for the data types we use in this development are in figure 4.2.

Parametricity results: “Theorems For Free” We have explained how AGDA’s logical relation is defined ($[[_]]$) on various types. We can now more precisely explain how to use the “Abstraction theorem” associated to AGDA terms.

For every well-typed term M at type τ , the theorem $[[\tau]] M M$ is provable.

In Type Theory (such as AGDA and COQ) the specification language is the same as the programming language, which helps a lot. The theorem statement: $[[\tau]] M M$ is a valid type in AGDA and a proof of it is a term of this type.

We warmly recommend the reader to study [Bernardy et al., 2010] to get a more in-depth understanding on the subject. We still try to give some insight on the “Abstraction Theorem”. We write $\Gamma \vdash M : \tau$ for a typing judgement on AGDA terms. In such a judgement, Γ is a typing environment holding previously defined terms and data types. While we explained how $[[_]]$

is defined on types, it extends nicely on terms which is actually necessary in AGDA since terms and types cohabit. The meta-function $\llbracket _ \rrbracket$ is also defined on environments by translating each binding of the environment.

$$\Gamma \vdash M : \tau \Rightarrow \llbracket \Gamma \rrbracket \vdash \llbracket M \rrbracket : \llbracket \tau \rrbracket M M$$

This theorem reads like, for all term M , well-typed under environment Γ at type τ , the term $\llbracket M \rrbracket$ is well-typed under environment $\llbracket \Gamma \rrbracket$ at type $\llbracket \tau \rrbracket M M$. If we read the type $\llbracket \tau \rrbracket M M$ as a proposition then the term $\llbracket M \rrbracket$ is its proof.

Hence, every well-typed term give rise to a free-theorem thanks to parametricity. However interesting results arose only when using polymorphism. In our setting polymorphism comes at two places. First, by making the types of our interface abstract we force the client to be polymorphic with respect to these types. Second, when the client write a world-polymorphic function this give rise to even more powerful parametricity results. In section 5.4.3 we give a more detailed account to various function types and their respective strength.

4.2.2 An example: Boolean values represented by numbers

We wish to explain how logical relations can help explain in what sense the interface to an abstract type is safe. To do so we introduce a tiny example about booleans represented using natural numbers. We want 0 to represent `false` and any other number to represent `true`. Therefore the boolean disjunction can be implemented using addition. We show that logical relations help build a model, ensure that a given implementation respects this model, and finally show that a client that uses only the interface also respect the model.

Note however that this is a toy example in several ways. There are no polymorphic functions in the interface, so no interesting free theorems are to be expected. While we could prove safety by defining a representation predicate in unary style, the logical relations approach is different. It relies on comparing concrete data as opposed to mapping to abstract data. The unary style would allow for a simpler construction, however this oversimplifies the problem here and would be no longer useful for proving our library.

Our tiny implementation of booleans using natural numbers is given below. It contains a type `B` that we want to keep abstract. It contains obvious definitions for `true`, `false`, and the disjunction `_∨_`. It intentionally has a dubious function `is42?`. Given that 42 and 41 both represent a truth value, the function `is42?` returns different results. This dubious function breaks our still informal expectations from such a module.

```

B : Set
B = ℕ

false : B
false = 0

true : B
true = 1

_∨_ : B → B → B
m ∨ n = m + n

is42? : B → B
is42? 42 = true
is42? _ = false

```

The next step is to define our expectations. To do so, we give a binary relation which tells when two `B` values have the same meaning. We do so with an (inductive) data type named `[[B]]` which states that 0 is related only with itself, and that any two non zero numbers are related, formalizing the intuition that two non zero numbers mean the same thing:

```

data [[B]] : B → B → Set where
  [[false]] : [[B]] 0 0
  [[true]]  : ∀ {m n} → [[B]] (suc m) (suc n)

```

When plugged into the machinery of logical relations this single definition suffices to define a complete model of well-typed programs. However, the plumbing requires some care. While the AGDA logical relation is universal, we have no such guarantee about the AGDA logical relation *where* the relation for `B` is no longer `[[N]]` but `[[B]]`. Fortunately changing the relation at a given type (`B` here) can be done safely. All we have to do is to consider programs abstracted away from `B` and its operations: `true`, `false` and `_∨_`. This can be done either through a mechanism for abstract types, or by requiring the client to be a function taking the implementation for `B` and its operations as argument.

However, to use `[[B]]` as the relation for `B`, we have to show that the definitions which make use of the representation of `B` actually fit the relation. Since `[[true]]` and `[[false]]` are obvious witnesses for `true` and `false`, only `_∨_` and `is42?` are left to be proved. Each time the goal to prove is systematic: wrap the type with `[[·]]` on each constructor and put the name of the function twice to state we want it to be related to itself. Here is the definition for `_[[∨]]_`:

$\llbracket \vee \rrbracket_- : ((\llbracket B \rrbracket \llbracket \rightarrow \rrbracket \llbracket B \rrbracket \llbracket \rightarrow \rrbracket \llbracket B \rrbracket)) _ \vee _ _ \vee _$

The type of $\llbracket \vee \rrbracket_-$ means that given inputs related in the model, the results are related in the model as well. Once unfolded the type looks like:

$\llbracket \vee \rrbracket_- : \forall \{x_1 \ x_2\} (x_r : \llbracket B \rrbracket \ x_1 \ x_2)$
 $\quad \{y_1 \ y_2\} (y_r : \llbracket B \rrbracket \ y_1 \ y_2)$
 $\quad \rightarrow \llbracket B \rrbracket (x_1 \vee y_1) (x_2 \vee y_2)$

The fact that input arguments come as implicit arguments greatly shortens definitions. Now, thanks to the inductive definition of $_+_-$, pattern-matching on the first relation suffices to reduce the goal, and allows this nice-looking definition of $\llbracket \vee \rrbracket_-$ where we see the usual lazy definition of the left biased disjunction:

$\llbracket \text{false} \rrbracket \llbracket \vee \rrbracket \ x = x$
 $\llbracket \text{true} \rrbracket \llbracket \vee \rrbracket \ _ = \llbracket \text{true} \rrbracket$

Let us now consider a proof for the function `is42?`. Fortunately there is no such proof since this function gladly breaks the intended abstraction. Instead we simply prove its negation by exhibiting that given two related inputs (42 and 27) we get non related outputs (since `is42? 42 = 1` and `is42? 27 = 0`).

$\neg \llbracket \text{is42?} \rrbracket : \neg((\llbracket B \rrbracket \llbracket \rightarrow \rrbracket \llbracket B \rrbracket) \text{is42?} \text{is42?})$
 $\neg \llbracket \text{is42?} \rrbracket \llbracket \text{is42?} \rrbracket \text{ with } \llbracket \text{is42?} \rrbracket \{42\} \{27\} \llbracket \text{true} \rrbracket$
 $\dots \quad \quad \quad | () \text{ -- absurd}$

Note that `is42?` is rejected by our model with no considerations about the other exported functions. Here we have a means to build 42 by using $_ \vee _$ and `true`. However with another implementation of $_ \vee _$ there would be no way to produce 42 and so no way to expose the wrong behavior of `is42?` using the interface. This would enable some proof techniques to actually show that no harm can be done given the limited interface and implementation. Such a proof would break down when extending the interface. In the end, using a model provides a better forward compatibility than syntactic proof techniques and enables proofs to be done in a modular way.

4.3 Relations for NomPa

For NOMPA, we apply the same process as with booleans. We define our expectations, by defining relations for introduced abstract types (worlds, binders, names, fresh-for and inclusions). Finally we prove that each value

and/or function exported fits the relation.

4.3.1 Relations for NomPa types

For reference the definitions are given in figure 4.3. We now describe them in turn.

While the types and the worlds define the programs we accept as valid, logical relations define when two programs have the same “meaning”. For instance, valid names are those which belong to their worlds, names with the same meaning are those related by the relation between their worlds. What matters is not just the *fact* that two worlds α_1 and α_2 are related, what matters is *how* they are related, because this dictates when a name in α_1 and a name in α_2 are related. Thus we define $\llbracket \text{World} \rrbracket \alpha_1 \alpha_2$ to be a set of relations between $\text{Name } \alpha_1$ and $\text{Name } \alpha_2$ (we note α_r a relation of type $\llbracket \text{World} \rrbracket \alpha_1 \alpha_2$).

Which relations should be part of this set? The more relations are part of it, the more power this gives to parametricity results. . . However, in order for the equality test on names to inhabit the logical relation, we must restrict $\llbracket \text{World} \rrbracket \alpha_1 \alpha_2$ to contain only relations that preserve equality in both directions, that is, only functional and injective relations:

$$\begin{aligned} \text{Preserve-}\equiv \mathcal{R} = \\ \forall x_1 y_1 x_2 y_2 \rightarrow \mathcal{R} x_1 x_2 \rightarrow \mathcal{R} y_1 y_2 \\ \rightarrow x_1 \equiv y_1 \leftrightarrow x_2 \equiv y_2 \end{aligned}$$

The relation $\llbracket \text{Name} \rrbracket$ is parameterized by a relation over worlds. Given two worlds α_1 and α_2 , and a relation α_r between them (the world α_r has type $\llbracket \text{World} \rrbracket \alpha_1 \alpha_2$), then a name x_1 in α_1 is related to a name x_2 in α_2 if and only if they are related by α_r : $\alpha_r x_1 x_2$. In short, the relation between names is completely defined by the relation between their worlds.

The relation $\llbracket \text{Binder} \rrbracket$ between two binders is both simple and meaningful. It is simple since it is the full relation, meaning that any two binders are related by $\llbracket \text{Binder} \rrbracket$. It is meaningful since it recalls us that one binder is not better than another one. For instance when comparing object-level terms as we do in section 3.1.3, the functions `cmpTm` and `_==Tm_` never compare binders.

This definition is meaningful for a second reason that appears when one looks at the free-theorems associated with functions that expect binders. For instance, a well-typed function `f` of type `Binder → Binder → Bool` can only be a constant function!

```

-- [[f]] is the parametricity theorem of f
[[f]] : ([[Binder]] [[→]] [[Binder]] [[→]] [[Bool]]) f f

-- f-const is a corollary of [[f]].
-- f-const shows that f is a constant function.
f-const : ∀ x1 x2 y1 y2 → f x1 y1 ≡ f x2 y2
f-const x1 x2 y1 y2 with [[f]] {x1} {x2} - {y1} {y2} -
... | [[true]] = refl
... | [[false]] = refl

```

This proves that there must be no way to distinguish two binders!

For the $\llbracket \subseteq \rrbracket$ relation, we exploit the fact there is only one way to use an inclusion witness, namely `coerceN`. Thus, for the purpose of building the “model”, we identify inclusions with their use in `coerceN`. That is, two witnesses $\alpha_1 \subseteq \beta_1$ and $\alpha_2 \subseteq \beta_2$ are related if and only if `coerceN α1 ⊆ β1` and `coerceN α2 ⊆ β2` are related (see figure 4.3). Another way to look at it is from the perspective of relation inclusion. A relation \mathcal{R}_1 is included in a relation \mathcal{R}_2 if and only if all pairs related by \mathcal{R}_1 are related by \mathcal{R}_2 as well. Because the function `coerceN` behaves such as the identity function, these two definitions coincide. If we expand the definitions for `[[Name]]` and $\llbracket \rightarrow \rrbracket$, we find that $\alpha_1 \subseteq \beta_1$ and $\alpha_2 \subseteq \beta_2$ are related if and only if the relation α_r is included in the relation β_r :

$$\begin{aligned} \llbracket \subseteq \rrbracket \alpha_r \beta_r \alpha_1 \subseteq \beta_1 \alpha_2 \subseteq \beta_2 \\ &= \forall \{x_1 \ x_2\} \rightarrow (x_1, x_2) \in \alpha_r \\ &\quad \rightarrow (\text{coerce}^N \alpha_1 \subseteq \beta_1 \ x_1, \text{coerce}^N \alpha_2 \subseteq \beta_2 \ x_2) \in \beta_r \end{aligned}$$

We now have to define operations on worlds that fit the logical relation.

The case for `[[∅]]` is trivial and uniquely defined. There is only one relation between empty sets: the empty relation.

```
[[∅]] : [[World]] ∅ ∅
```

We give below a set-theoretic definition of $\llbracket \triangleleft \rrbracket$. Our definition is not based on sets but equivalent to this. This new relation is shown to preserve the equalities.

```

-[[<]]- : ([[Binder]] [[→]] [[World]] [[→]] [[World]]) -<- -<-
-- not proper Agda
br [[<]] αr def { (b1, b2) } ∪ { (x, y) | (x, y) ∈ αr ∧ x ≠ b1 ∧ y ≠ b2 }

```

```

Preserve-≡ : {A B : Set0} (R : A → B → Set0) → Set0
Preserve-≡ R =
  ∀ x1 y1 x2 y2 → R x1 x2 → R y1 y2
  → x1 ≡ y1 ↔ x2 ≡ y2

--      [[World]] : [[Set1]] World World
record [[World]] (α1 α2 : World) : Set1 where
  constructor -,-
  field
    R          : Name α1 → Name α2 → Set
    R-pres-≡  : Preserve-≡ R

[[Name]] : ([[World]] [[→]] [[Set0]]) Name Name
-- : ∀ {α1 α2} → [[World]] α1 α2 → Name α1 → Name α2 → Set
[[Name]] (R , -) x1 x2 = R x1 x2

[[Binder]] : [[Set0]] Binder Binder
-- : Binder → Binder → Set
[[Binder]] - - = ⊤


-[[#]]- : ([[Binder]] [[→]] [[World]] [[→]] [[Set0]]) -#- -#-
-- : ∀ {b1 b2} → [[Binder]] b1 b2 → ∀ {α1 α2} → [[World]] α1 α2
--   → b1 # α1 → b2 # α2 → Set
-[[#]]- - - - - = ⊤

-[[⊆]]- : ([[World]] [[→]] [[World]] [[→]] [[Set0]])
  -⊆- -⊆-
-- : ∀ {α1 α2} → [[World]] α1 α2 →
--   ∀ {β1 β2} → [[World]] β1 β2 →
--     α1 ⊆ β1 → α2 ⊆ β2 → Set
-[[⊆]]- αr βr α1⊆β1 α2⊆β2
= ([[Name]] αr [[→]] [[Name]] βr) (coerceN α1⊆β1) (coerceN α2⊆β2)

```

Figure 4.3: Relations for NOMPA types

Here is a small drawing of the effect of $\llbracket \langle _ \rangle \rrbracket$ on relations between names:



4.3.2 NomPa values fit the relation

We now give a short overview of the proofs needed to show that our functions fit the relation. Formally, for each primitive operation \mathbf{p} which has type τ in our interface, we have to exhibit a definition $\llbracket \mathbf{p} \rrbracket$ which has type $\llbracket \tau \rrbracket \mathbf{p} \mathbf{p}$. All the proof can be found in our AGDA development [Pouillard, 2011a]. Thanks to the definition of $\llbracket \text{Binder} \rrbracket$, zero^B and suc^B , any function that returns a binder immediately fit the relation. The same reasoning holds for $_ \# \emptyset$ and $\text{suc} \#$ that immediately fit the relation. Thanks to the definition of $\llbracket \langle _ \rangle \rrbracket$, $\llbracket \text{name}^B \rrbracket$ holds by definition. Since $\llbracket \neg \text{Name} \emptyset \rrbracket$ receives names in the empty world, it trivially holds. For instance here is the type signature for $\llbracket \text{name}^B \rrbracket$:

```

[[nameB]] : (∀ (αr : [[World]]) (→)
  (br : [[Binder]]) (→)
  [[Name]] (br [[<_>]] αr)
) nameB nameB
--      : ∀ {α1 α2} (αr : [[World]] α1 α2)
--      {b1 b2} (br : [[Binder]] b1 b2)
--      → [[Name]] (br [[<_>]] αr) (nameB {α1} b1) (nameB {α2} b2)

```

For $\llbracket \text{==}^N \rrbracket$, once unfolded, the statement tells that the equality test commutes with a renaming. This means that the result of the equality test does not change when its inputs are consistently renamed. The proof for $\llbracket \text{==}^N \rrbracket$ is done in two parts. First, we have to relate the boolean valued function $_ \text{==}^N _$ to the fact it decides equality on names. Second, we make full use of the equality preservation property.

```

-[[==N]]_ : (∀⟨ αr : [[World]] ⟩[→])
             [[Name]] αr [→]
             [[Name]] αr [→]
             [[Bool]]
             ) _==N_ _==N_
--      : ∀ {α1 α2} (αr : [[World]] α1 α2)
--          {x1 x2} (xr : [[Name]] αr x1 x2)
--          {y1 y2} (yr : [[Name]] αr y1 y2)
--          → [[Bool]] (x1 ==N y1) (x2 ==N y2)

```

The proof that `exportN?` is in the relation relies on two points. First, the success (returning `just` or `nothing`) of `(exportN? {b} x)` depends only on the equality between `(nameB b)` and `x`. Second, the worlds related by `(br [⟨] αr)` can always be split into two groups, those being related by `(br [⟨] [∅])` and those related by `αr`. Said differently there is no link between the two parts.

Thanks to the definition of `-[[⊆]]-`, the proof that `coerceN` fits the relation is a simple application of the hypotheses. Then the real job is to show that all the inclusion rules fit the relation. This means that they all behave as the identity function. In the end all the logical relation proofs associated with all our primitive operations and rules are formally proved in our AGDA development [Pouillard, 2011a].

4.3.3 An example not fitting the relation

For instance consider the introduction of this function to compare two names:

```

_<=N_ : ∀ {α} → Name α → Name α → Bool
(m , _) <=N (n , _) = N._<=_ m n

```

This function is accepted by AGDA type-checking rules but does not meet our property that well-behaved functions should map α -equivalent inputs to α -equivalent outputs.

```

-[[<=N]] : ¬((∀⟨ αr : [[World]] ⟩[→])
             [[Name]] αr [→]
             [[Name]] αr [→] [[Bool]]) _<=N_ _<=N_
-[[<=N]] [[<=]] = ¬[[Bool]]-true-false ([[<=]] ? {0 , _} {1 , _} ?
                                         {1 , _} {0 , _} ?)
-- parts ('?') of the proofs are omitted for conciseness

```

Chapter 5

The de Bruijn approach

So far, we only considered a nominal representation. This is not the only way to deal with names and binders. We now focus on a different kind of representation, namely de Bruijn indices [de Bruijn, 1972]. This representation is said to be nameless because binders are no longer identified by an atom but by a notion of “distance” to the binding point.

This nameless approach solves part of the problem of representing binders by providing a canonical representation. More precisely binding occurrences are no longer named (we now use $\lambda.$ instead of $\lambda x.$). Bound variables are represented by the “distance” to the binding λ . This distance is the number (starting from 0) of enclosing λ s to cross in order to reach the binding λ .

This approach to a safer way to program with de Bruijn indices is described in our previous paper called “Nameless, Painless” [Pouillard, 2011b]. The library specialized to de Bruijn indices is called NAPA for **N**ameless **P**ainless. However we here stick to the general approach.

This chapter is organized as follows. The first section introduces de Bruijn indices through successive variations found in the literature. The second section 5.2, *extends* our programming interface in order to support de Bruijn indices. Also, the implementation is *extended* and not replaced. Section 5.3 explains how to use the interface to write functions on a de Bruijn representation of programs. In section 5.4 we extend the logical relation to support the de Bruijn style and we show how to exploit these properties on concrete examples.

5.1 Introduction to de Bruijn indices

Various techniques have been discovered to build a nameless representation. We have chosen a few of them which gradually set up the stage.

5.1.1 bare: The original approach

We call this one **bare**, because it relies solely on natural numbers. To make things more concrete here is an example of its use when defining our usual type for terms:

```
data TmB : Set where
  V      : (x      : ℕ)  → TmB
  _ · _  : (t u    : TmB) → TmB
  λ      : (t      : TmB) → TmB
  Let    : (t u    : TmB) → TmB
```

From the point of view of the binding structure, it is striking that no difference appears between the constructors of the data type. It is completely up to the programmer to manage the scoping difference introduced by λ and **Let**. This is even more worrying in the **Let** case since we have no clue of difference between the arguments.

Here is an example using this representation to build the λ -term for function application, namely $\lambda f. \lambda x. f\ x$.

```
appTmB : TmB
appTmB = λ (λ (V 1 · V 0))
```

The main advantages of this approach are its simplicity and its expressiveness. The expressiveness is maximal since no restriction is put on the usage of variables.

5.1.2 Maybe: The nested data type approach

The nested data type approach [Bellegarde and Hook, 1994, Bird and Paterson, 1999, Altenkirch and Reus, 1999] is a first step towards better properties about the binding structure of terms. Let us start with the definition of our type for terms with this approach:

```
data TmM (A : Set) : Set where
  V      : (x : A) → TmM A
  _ · _  : (t u : TmM A) → TmM A
  λ      : (t : TmM (Maybe A)) → TmM A
  Let    : (t : TmM A) (u : TmM (Maybe A)) → TmM A
```

There are three points to look at. The type Tm^M is parameterized by another type called A , so we can look at it as a kind of container. Note also that the variable case **V** does not hold a value of type \mathbb{N} but a value of type A . Last but not least the λ case holds a term whose parameter is not

simply a value of type `A` but a value of type `Maybe A`.

This last point makes the type `TmM` a nested data type, also called a non-regular type. This has the consequence of requiring polymorphic recursion to write recursive functions on such a type.

To understand why this is an adequate representation of λ -terms one has to look a bit more at the meaning of `Maybe`. If types are seen as sets, then `Maybe` takes a set and returns a set with one extra element. So each time we cross a `λ`, there is one extra element in the set of allowed variables, exactly capturing the fact that we are introducing a variable.

To see the difference with the previous approach, here is the λ -term for function application again:

```
appTmM : TmM ⊥
appTmM = λ (λ (V (just nothing) · V nothing))
```

Note the use of the empty type `⊥` to state the closedness of the `appTmM` term. Stating this kind of properties was impossible to do with the previous approach, without resorting to logical properties on the side.

5.1.3 The `Fin` approach

Another approach already described and used in [Altenkirch, 1993, McBride and McKinna, 2004] is to index everything (terms for example) by a bound. This bound is the maximum number of distinct free variables allowed in the value. This rule is enforced in two parts: variables have to be strictly lower than their bound, and the bound is incremented by one when crossing a name abstraction (a λ -abstraction for instance, called `λ` here).

The type `Fin n` is used for variables and represents natural numbers strictly lower than `n`. The name `Fin n` comes from the fact that it defines finite sets of size `n`. We call this approach `Fin` for its use of this type. The definition found in AGDA's standard library is the following:

```
data Fin : ℕ → Set where
  zero : {n : ℕ} → Fin (suc n)
  suc  : {n : ℕ} (i : Fin n) → Fin (suc n)
```

Given this type `Fin`, one can define the term data type using this approach:

```
data Tmf n : Set where
  V      : (x : Fin n) → Tmf n
  _·_    : (t u : Tmf n) → Tmf n
  λ      : (t : Tmf (suc n)) → Tmf n
  Let    : (t : Tmf n) (u : Tmf (suc n)) → Tmf n
```


As the previous approach, this representation helps enforce some well-formedness properties, for instance $\mathbf{Tm}^f \ 0$ is the type of closed λ -terms.

Here is the λ -term for application in this approach:

```
appTmf : Tmf 0
appTmf = λ (λ (V (suc zero) · V zero))
```

We can easily draw a link with the **Maybe** approach. Indeed, the type $\mathbf{Fin} \ (\mathbf{suc} \ n)$ has exactly one more element than the type $\mathbf{Fin} \ n$. However, these approaches are not equivalent for at least two reasons. The **Maybe** approach can accept any type to represent variables. This makes the structure more like a container and this can be particularly helpful to define the substitution as the composition of $\mathbf{mapTm} : \forall\{A \ B\} \rightarrow (A \rightarrow B) \rightarrow \mathbf{Tm} \ A \rightarrow \mathbf{Tm} \ B$ and $\mathbf{joinTm} : \forall\{A\} \rightarrow \mathbf{Tm} \ (\mathbf{Tm} \ A) \rightarrow \mathbf{Tm} \ A$ as in [Bellegarde and Hook, 1994, Bird and Paterson, 1999, Altenkirch and Reus, 1999]. The **Fin** approach has advantages as well: the representation is concrete and simpler since closer to the **bare** approach. However this apparent simplicity comes at a cost: we will see that its concrete representation is one root of the problem.

5.2 An interface for de Bruijn indices

In order to effectively program with de Bruijn indices we wanted to have a dedicated programming interface. This is indeed in contrast with our previous work [Pouillard and Pottier, 2010] where we had two implementations – nominal and de Bruijn – of a common interface.

We now claim that we need specialized operations to better support nominal and de Bruijn styles. Happily the work done for nominal nicely scales to de Bruijn indices by *adding* new primitive operations on worlds, names, and world inclusions.

Following the development that we have done for the nominal representation, we index names and terms by worlds. Everything we have built so far remains valid (and useful). The types and definitions for **World**, \emptyset , **Name**, $_==^N _$, $_Name\emptyset$, $_ \subseteq _$, \mathbf{coerce}^N , $_ \subseteq \mathbf{-refl}$, $_ \subseteq \mathbf{-trans}$ and $_ \subseteq \mathbf{-}\emptyset$ are directly re-used.

While binders (**Binder**, $_ \triangleleft _$, \mathbf{zero}^B , $\mathbf{suc}^B \dots$) do not need to be exposed in order to follow a de Bruijn presentation, we use them in the implementation of de Bruijn primitive operations.

To understand what definitions we need to add to deal with de Bruijn indices, we identify two characteristics of a binding construct such as λ . With de Bruijn indices two things occur at λ -abstractions: 0 is introduced as a new binder, and all the previous binders are incremented by one.

The key feature of this approach is to change how binders are referenced in a particular subterm. Our explicit notion of world can be used to express this change. We extend our notion of worlds with an operator `_+1` (formally defined on page 98). Since worlds can be thought of as sets, the operator `_+1` has an informal set-theoretic definition of adding one to each element of the set. Postponing a formal definition, here is the declaration:

```
_+1 : World → World
```

With this single extension of world operations (`_+1`), we can build name abstractions in the style of de Bruijn indices. Here is our type `TmD` revised to follow de Bruijn indices:

```
data TmD α : Set where
  V    : Name α → TmD α
  _·_  : TmD α → TmD α → TmD α
  λ    : TmD (0B < (α + 1)) → TmD α
  Let  : TmD α → TmD (0B < (α + 1)) → TmD α
```

The only change between `Tm` and `TmD` is in the binding constructs. Instead of holding a binder `b` that is bound in the subterm, `0B` is bound in the subterm and all the previous binders are moved by one. Because this world operation is idiomatic, it deserves its own notation:

```
_↑1 : World → World
α ↑1 = 0B < (α + 1)
```

We call the operation `_↑1` shifting by one. Note that `_↑1` conveniently hides the underlying usage of binders. The type `TmD` now reads like:

```
data TmD α : Set where
  V    : Name α → TmD α
  _·_  : TmD α → TmD α → TmD α
  λ    : TmD (α ↑1) → TmD α
  Let  : TmD α → TmD (α ↑1) → TmD α
```

We get back the nested data type approach described in section 5.1.2 if one replaces `World` by `Set`, `Name` by the identity, and `_↑1` by `Maybe`. The same connection can be done with the `Fin` approach described in section 5.1.3 if one replaces `World` by `N`, `Name` by `Fin`, and `_↑1` by `suc`.

Our internal representation for worlds (section 4.1) using lists of boolean values is chosen to make `_+1` easy to define and to work with:

```
_+1 : World → World
α +1 = false :: α
```

We defined `_↑1` on top of the primitive operators `_+1` and `_◁_`. Unfolding `_↑1` and `_◁_` shows the `true` nature of `_↑1`:

```
↑1-is-true : : ∀ {α} → α ↑1 ≡ true :: α
↑1-is-true : = refl
-- Since α ↑1 equals zero ◁ false :: α
-- which reduces to true :: α
```

These one-step definitions `_+1` and `_↑1` are extended to any number to produce `_+W_` and `_↑k_` of type `World → ℕ → World`. The world `α +W k` is `α +1...+1` (`_+1` iterated `k` times) and the world `α ↑ k` is `α ↑1...↑1` (`_↑1` iterated `k` times).

To the best of our knowledge, making the distinction between two forms of shifting operations for worlds (namely `_+1` and `_↑1`) has never been investigated in the context of representing names and binders. As we see later in detail in section 5.4.3, this distinction is important and omitting it leads to less accurate types. Less accurate types led us to a broken definition of shifting in our first work on the subject [Pouillard and Pottier, 2010] as we detail in section 5.4.3.

ℓ -bound and ℓ -free names: In a context where we are working under ℓ binders, we call ℓ -bound, a name bound somewhere in the scope of ℓ binders. We call ℓ -free, a name that is free for all ℓ binders. In other words, a de Bruijn index is ℓ -bound if it is strictly less than ℓ ; it is ℓ -free otherwise.

`zeroN`: Since binders are not needed in the de Bruijn presentation, names are constructed directly. To start with, the simplest name is `zeroN`. The name `zeroN` represents zero and inhabits any world shifted by one. While this is a primitive operation from a de Bruijn point of view, it can be defined in terms of our nominal interface:

```
zeroN : ∀ {α} → Name (α ↑1)
zeroN = nameB (0B)
```

The signatures of the core primitive operations on names are given in figure 5.1. One can add any constant to a name in any world with `addN`. As-

$$\begin{aligned}
\text{add}^N & : \forall \{\alpha\} \ k \rightarrow \text{Name } \alpha \\
& \qquad \qquad \qquad \rightarrow \text{Name } (\alpha +^W k) \\
\text{subtract}^N & : \forall \{\alpha\} \ k \rightarrow \text{Name } (\alpha +^W k) \\
& \qquad \qquad \qquad \rightarrow \text{Name } \alpha \\
\text{cmp}^N & : \forall \{\alpha\} \ \ell \rightarrow \text{Name } (\alpha \uparrow \ell) \\
& \qquad \qquad \qquad \rightarrow \text{Name } (\emptyset \uparrow \ell) \uplus \text{Name } (\alpha +^W \ell)
\end{aligned}$$

$$\begin{aligned}
\text{syntax } \text{add}^N & \quad k \ x = x +^N k \\
\text{syntax } \text{subtract}^N & \quad k \ x = x \dot{-}^N k \\
\text{syntax } \text{cmp}^N & \quad \ell \ x = x <^N \ell
\end{aligned}$$

Figure 5.1: Core operations on names




Figure 5.2: Operations on names

$$\begin{aligned}
\subseteq\text{-}\emptyset\text{+}1 & : \emptyset + 1 \subseteq \emptyset \\
\subseteq\text{-}\uparrow 1\text{-}\uparrow 1 & : \forall\{\alpha \beta\} \rightarrow \alpha \subseteq \beta \leftrightarrow \alpha \uparrow 1 \subseteq \beta \uparrow 1 \\
\subseteq\text{-}\text{+}1\text{-}\text{+}1 & : \forall\{\alpha \beta\} \rightarrow \alpha \subseteq \beta \leftrightarrow \alpha + 1 \subseteq \beta + 1 \\
\subseteq\text{-}\text{+}1\text{-}\uparrow 1 & : \forall\{\alpha\} \rightarrow \alpha + 1 \subseteq \alpha \uparrow 1
\end{aligned}$$

Figure 5.3: New rules for world inclusion

suming worlds are erased at run-time, by parametricity the resulting world of add^N clearly shows that this function does exactly its addition job. One can do the opposite operation with subtract^N . Thanks to its precise type this function is total and the inverse of add^N .

Given any world α , a name in the world $\alpha \uparrow \ell$ is either strictly lower than ℓ (and so also lives in $\emptyset \uparrow \ell$), or greater or equal to ℓ (thus also lives in $\alpha +^W \ell$). This is exactly what the function $\text{cmp}^N \ell$ is about. Given a name x , $\text{cmp}^N \ell x$ returns a disjoint sum of names which can be read in two parts. It first gives which side of the disjoint sum it stands, and second it gives a refined version of the input name x .

Figure 5.2 depicts the different types of names we have and how our operations relates them. Starting from the bottom, names with worlds of the form $\emptyset \uparrow k$ are definitely k -bound. Their value is completely known. The two arrows state the type isomorphism between $\text{Name } (\emptyset \uparrow k)$ and $\text{Fin } k$. Above we have names that may be bound or free ($\text{Name } (\alpha \uparrow k)$). A dynamic test (cmp^N) can tell whether such a name is k -bound or not. Above we have names that are known to be greater than k ($\text{Name } (\alpha +^W k)$). They are k -free names. Above we have free names ($\text{Name } \alpha$). On the top we have impossible names since they are said to belong to an empty world. From them we can derive everything.

These primitive operations are enough to show an isomorphism between $\text{Fin } n$ and $\text{Name } (\emptyset \uparrow n)$. From this, every program involving Fin can be translated into our system. This means that our system does not restrict the programmer more than the Fin approach. However as soon as one uses finer types than $\text{Name } (\emptyset \uparrow n)$, then fewer “wrong programs” type-check and more properties hold as we see in section 5.4.3.

We extend the inclusion relation with a set of new rules given in figure 5.3. The $\subseteq\text{-}\emptyset\text{+}1$ rule states that $\emptyset + 1$ is empty. Inclusion is preserved both ways by the contexts $\text{-}\uparrow 1$ and $\text{-}\text{+}1$. Finally $\text{-}\text{+}1$ can be weakened to $\text{-}\uparrow 1$. This accounts for the fact that $\alpha \uparrow 1$ means $\{0\} \cup (\alpha + 1)$ and so is a proper superset of $\alpha + 1$. This set of rules is implemented according to the definition of inclusion, namely $\forall x \rightarrow x \in \alpha \rightarrow x \in \beta$. On top of these base

rules, we derive others that we omit here (some are used in the code).

```
--  $\subseteq$ - $\uparrow$ 1 $\rightarrow$  $\uparrow$ 1 is derivable
 $\subseteq$ - $\uparrow$ 1 $\rightarrow$  $\uparrow$ 1 :  $\forall\{\alpha\ \beta\}\rightarrow \alpha \subseteq \beta \rightarrow \alpha \uparrow 1 \subseteq \beta \uparrow 1$ 
 $\subseteq$ - $\uparrow$ 1 $\rightarrow$  $\uparrow$ 1 =  $\subseteq$ - $\triangleleft$  0
```

5.2.1 Singleton worlds!

We said that our worlds denote finite subsets of \mathbb{N} and are more precise than in the `Fin` approach. Actually they can be as precise as we wish, since any subset of \mathbb{N} can be described by our operations on worlds (`\emptyset` , `$+1$` , and `$\uparrow 1$`). In particular they can be singleton worlds. From singleton worlds we build singleton types for names:

```
Worlds :  $\mathbb{N} \rightarrow$  World
Worlds n =  $\emptyset \uparrow 1 +^W n$ 

Names :  $\mathbb{N} \rightarrow$  Set
Names = Name  $\circ$  Worlds

_ns :  $\forall n \rightarrow$  Names n
_ns n = zeroN +N n
```

Singleton worlds not only exist, they are also preserved by our two updating operations, namely `adds` and `subtracts`.

```
adds :  $\forall \{n\}\ k \rightarrow$  Names n  $\rightarrow$  Names (k + $\mathbb{N}$  n)
adds {n} k x = addN k x
  <-because  $\subseteq$ -assoc-+  $\subseteq$ -refl n k ->

subtracts :  $\forall \{n\}\ k \rightarrow$  Names (k + $\mathbb{N}$  n)  $\rightarrow$  Names n
subtracts {n} k x = subtractN k x
  <-because  $\subseteq$ -assoc-+'  $\subseteq$ -refl n k ->
```

5.3 Examples and advanced operations

5.3.1 Some convenience functions

Here are a few functions built on top of the interface (without using the concrete representation of names). `sucN` is `addN 1` and `sucN \uparrow` is a variant that includes a coercion from `$\alpha + 1$` to `$\alpha \uparrow 1$` . The function `_N` turns a number `n`

into a name that inhabits any world with at least $n + 1$ consecutive binders.

```
sucN : ∀ {α} → Name α → Name (α + 1)
sucN = addN 1
```

```
sucN↑ : ∀ {α} → Name α → Name (α ↑ 1)
sucN↑ = coerceN ⊆-+1↑1 ∘ sucN
```

```
_N : ∀ {α} n → Name (α ↑ suc n)
_N {α} n = zeroN +N n
  <-because α ↑ 1 +W n ⊆⟨ ⊆-+-↑ n ⟩
           α ↑ 1 ↑ n ⊆⟨ ⊆-exch-↑-↑ 1 n ⟩
           α ↑ suc n ■ ->
where open ⊆-Reasoning
```

The function `subtractN?`, similarly to the function `cmp` tells whether a given name is l -bound or l -free. In case the name is free, an “exported” version of it is returned. This function forms the base case of “exporting” functions such as `subtractTmD?` explained later on page 107. The function `predN?` is a simple specialization of `subtractN?`.

```
subtractN? : ∀ {α} ℓ → Name (α ↑ ℓ)
           →? Name α

subtractN? ℓ x
  with x <N ℓ
... | inj1 - = nothing
... | inj2 x' = just (x' -N ℓ)

predN? : ∀ {α} → Name (α ↑ 1) →? Name α
predN? = subtractN? 1
```

On top of `subtractN?` we also build a convenient eliminator for names which is simply the elimination of the result of `subtractN?`.

```
subtractWithN : ∀ {α A} ℓ → A → (Name α → A)
              → Name (α ↑ ℓ) → A
subtractWithN v f = maybe f v ∘' subtractN?
```

The function `shiftName ℓ k pf` shifts its argument by `k` if this name is l -free, otherwise it leaves the l -bound name untouched. This function makes use of `cmpN` and coerces the outputs to the required type. It also perform a

coercion on the fly, giving extra flexibility for free.

```

shiftName : ∀ {α β} ℓ k → (α +W k) ⊆ β
           → Name (α ↑ ℓ)
           → Name (β ↑ ℓ)

shiftName ℓ k pf x
  with x <N ℓ
... | inj1 x' = x'      ⟨-because pf1 -⟩ -- ℓ-bound
... | inj2 x' = x' +N k ⟨-because pf2 -⟩ -- ℓ-free
  where
  pf1 = ⊆-cong-↑ ⊆-∅ ℓ
  pf2 = ⊆-trans (⊆-exch-+++ ⊆-refl ℓ k)
           (⊆-ctx-↑ pf ℓ)

```

The function `protect↑` shifts a name transformer. Let `f` be a function from names to names. The function `protect↑ ℓ f` is a version of `f` that is applicable under `ℓ` binders. That is, if `f` has type `Name α → Name β` then `protect↑ ℓ f` has type `Name (α ↑ ℓ) → Name (β ↑ ℓ)`. Let `x` be a name under `ℓ` binders. When `x` is `ℓ`-bound, it is left untouched by `protect↑`. When `x` is `ℓ`-free, we can subtract `ℓ` to `x`, give it to `f`, and then add `ℓ` to get the result.

```

protect↑ : ∀ {α β} ℓ
         → (Name α → Name β)
         → (Name (α ↑ ℓ) → Name (β ↑ ℓ))

protect↑ ℓ f x
  with x <N ℓ
... | inj1 x' = x'
           ⟨-because ⊆-cong-↑ ⊆-∅ ℓ -⟩
... | inj2 x' = f (x' -N ℓ) +N ℓ
           ⟨-because ⊆-+-↑ ℓ -⟩

```

By combining `protect↑` and `addN` one obtains an alternative implementation of `shiftName` called `shiftName'`. However `shiftName` is more efficient since it avoids to subtracting `ℓ` to add it back after adding `k`.

```

shiftName' : ∀ {α β} ℓ k → (α +W k) ⊆ β
           → Name (α ↑ ℓ) → Name (β ↑ ℓ)
shiftName' ℓ k pf = protect↑ (coerceN pf ∘ addN k) ℓ

```


5.3.2 Building terms

Building terms with our de Bruijn style library is as easy as building them in a plain de Bruijn index representation. The structure is exactly the same, and the variables are made of numbers (of type \mathbb{N}) using $_N$. Below we define representations of the identity function as `idTmD`, the application operator as `appTmD`, and the composition function as `compTmD`:

```
idTmD : ∀ {α} → TmD α
idTmD = λ(V (0 N))
```

```
appTmD : ∀ {α} → TmD α
appTmD = λ(λ(V (1 N) · V (0 N)))
```

```
compTmD : ∀ {α} → TmD α
compTmD = λ(λ(λ(V (2 N)) · (V (1 N) · V (0 N))))
```

5.3.3 Computing free variables

Our first example of functions over terms computes a list of the free variables in the input term. The function `fvD` below, while straightforward, has the subtle cases of binding constructs (`λ` here). In these cases we have to remove the bound name from the list of free variables produced by the recursive call. In this nameless representation this amounts to removing occurrences of 0 and subtracting 1 to other name occurrences. This is done by the function `rm0` which calls `predN?` on each element of the list and merges the results. Note that forgetting to remove the bound variable results in a type error. In the same vein the typing of `fvD` ensures that all returned variables do appear free in the given term.

```
rm0 : ∀ {α} → List (Name (α ↑1))
      → List (Name α)
rm0 [] = []
rm0 (x :: xs) with predN? x
...         | just x' = x' :: rm0 xs
...         | nothing = rm0 xs
```

```

fvD : ∀ {α} → TmD α → List (Name α)
fvD (V x)      = [ x ]
fvD (fct · arg) = fvD fct ++ fvD arg
fvD (λ t)      = rm0 (fvD t)
fvD (Let t u)  = fvD t ++ rm0 (fvD u)

```

5.3.4 Generic traversal

As with the nominal approach (3.2) we build a generic traversal function to avoid code duplication among various operation on terms. However in this nameless style the matter is simplified. The information to carry down during the traversal can be summarized by a single natural number, namely the ℓ parameter that we have already used. This parameter records how many binders have been entered.

```

module TraverseTmD
  {E} (E-app : Applicative E)
  {α β} (trName : ∀ ℓ → Name (α ↑ ℓ)
          → E (TmD (β ↑ ℓ)))

  where
  open Applicative E-app

  tr : ∀ ℓ → TmD (α ↑ ℓ) → E (TmD (β ↑ ℓ))
  tr ℓ (V x)      = trName ℓ x
  tr ℓ (t · u)    = pure _·_ ⊗ tr ℓ t ⊗ tr ℓ u
  tr ℓ (λ t)      = pure λ ⊗ tr (suc ℓ) t
  tr ℓ (Let t u) = pure Let ⊗ tr ℓ t ⊗ tr (suc ℓ) u

  trTmD : TmD α → E (TmD β)
  trTmD = tr 0

```

Like in section 3.2 we made a special case for variable occurrences: `trName` is not wrapped with the constructor `V`. Mapping names to terms yields capture-avoiding substitution almost for free. However, in the meantime we build `trTmD` which maps names to names. It does so by applying `pure V` to the name-to-name function:

```

open TraverseTmD

trTmD :
  ∀ {E} (E-app : Applicative E) {α β}
    (trName : ∀ ℓ → Name (α ↑ ℓ)
              → E (Name (β ↑ ℓ)))
  → TmD α → E (TmD β)
trTmD E-app trName
= trTmD E-app (λ ℓ x → pure V ⊗ trName ℓ x)
  where open Applicative E-app

```

Renaming functions: In many functions over terms the handling of variables shares a common part. Given a variable under ℓ binders, we test if the variable is ℓ -bound. If so we leave it untouched, otherwise we subtract ℓ and go on a specific processing after which we add ℓ again to the free variables of the result. The traversal function is augmented by this processing of bound variables to build `renameTmDA`:

```

-- Like protect↑ but generalized to Applicative functors
protect↑A = {! omitted !}

renameTmDA : ∀ {E} (E-app : Applicative E)
              {α β} (θ : Name α → E (Name β))
              → TmD α → E (TmD β)
renameTmDA E-app θ
= trTmD E-app (protect↑A E-app θ)

```

Then by picking either the identity functor (`id-app`), or the `Maybe` one we build two functions (a total and a partial one) to perform any renaming, namely `renameTmD` and `renameTmD?`:

```

renameTmD : ∀ {α β} → (Name α → Name β)
              → TmD α → TmD β
renameTmD = trTmD id-app ∘ protect↑
-- or
-- renameTmD = renameTmDA id-app

renameTmD? : ∀ {α β} → (Name α →? Name β)
              → TmD α →? TmD β
renameTmD? = renameTmDA Maybe.applicative

```

Lifting name functions: Any operation on names can now be lifted to terms. The function `addTmD` lifts the function `addN` from names to terms. It takes a number `k` and a term and adds `k` to all the free variables:

```
addTmD : ∀ {α} k → TmD α → TmD (α +W k)
addTmD = renameTmD ∘ addN
```

The functions `subtractTmD` and `subtractTmD?` respectively lift the functions `subtractN` and `subtractN?` from names to terms. They enable to subtract a number `k` to all the free variables of a term:

```
subtractTmD : ∀ {α} k → TmD (α +W k) → TmD α
subtractTmD = renameTmD ∘ subtractN
```

```
subtractTmD? : ∀ {α} ℓ → TmD (α ↑ ℓ) →? TmD α
subtractTmD? = renameTmD? ∘ subtractN?
```

While `coerceN` can be lifted to terms in the same way, we can by-pass the `protect↑` dynamic tests and directly “protect” the inclusion witness with an appropriate inclusion rule. Put differently there are two ways to go from a proof `pf` of $\alpha \subseteq \beta$ to $(\text{Name } (\alpha \uparrow \ell) \rightarrow \text{Name } (\beta \uparrow \ell))$: one is to protect the result of `coerce` (`protect↑ ∘ coerceN pf`) and the other is to `coerce` with protected inclusion witness (`coerceN ∘ ⊆-cong-↑ pf`):

```
coerceTmD : ∀ {α β} → α ⊆ β → TmD α → TmD β
coerceTmD pf = trTmD id-app (coerceN ∘ ⊆-cong-↑ pf)
-- or less efficiently:
-- coerceTmD = renameTmD ∘ coerceN
```

Lifting `addN` to terms can be done more efficiently as well. Indeed `addTmD` internally uses `renameTmD` which uses `protect↑`. Here the dynamic test performed by `protect↑` is necessary. However when the name is ℓ -free we subtract ℓ and add it back after adding `k`. The function `shiftName` avoids this extra computation, hence the following `shiftTmD`:

```
shiftTmD : ∀ {α β} k → (α +W k) ⊆ β → TmD α → TmD β
shiftTmD k p = trTmD id-app (λ ℓ → shiftName ℓ k p)
-- or less efficiently:
-- shiftTmD k pf = renameTmD (coerceN pf ∘ addN k)
```

The function `closeTmD?` is built like in nominal style (section 3.2):

```
closeTmD? : ∀ {α} → TmD α →? TmD ∅
closeTmD? = renameTmD? (const nothing) -- any free var leads to a failure
```

Capture avoiding substitution To implement capture avoiding substitution for the type Tm^D , all we need is a specific function `trName` for `trTmD`. Here substitutions are represented as functions from names to terms. The function `substVarTmD` handles the case for variables. This function is very close to `protect↑` but extended to functions returning terms.

```

substVarTmD : ∀ {α β} → (Name α → TmD β) →
                ∀ ℓ → Name (α ↑ ℓ) → TmD (β ↑ ℓ)
substVarTmD f ℓ x
  with x <N ℓ
... | inj1 x' = V (x' (-because ⊆-cong-↑ ⊆-∅ ℓ -))
... | inj2 x' = shiftTmD ℓ (⊆-+-↑ ℓ) (f (x' ⋅N ℓ))

```

The main function `substTmD` instantiates `trTmD` with the identity applicative functor and `substVarTmD`:

```

substTmD : ∀ {α β} → (Name α → TmD β)
                → (TmD α → TmD β)
substTmD = trTmD id-app ∘ substVarTmD

```

As an illustration, the function `β-redD` performs a β -reduction at the root of the term using the function `substTmD`:

```

β-redD : ∀ {α} → TmD α → TmD α
β-redD (λ fct · arg) = substTmD (subtractWithN 1 arg V) fct
β-redD t              = t

```

5.3.5 Nameless term comparison

We show that all the subtle work is done at the level of names in a separate and reusable function called `cmpName↑`. This function takes a function that compares two free names and builds one that compares two names under ℓ bindings. It does so by comparing both of them to ℓ . If they are both bound they can be safely compared using `_==N_` since they now are of the same type. If they are both free, they can be compared using the function received as argument. Otherwise they are different.

```

cmpName↑ : ∀ {α β} ℓ → |Cmp| Name α β
                → |Cmp| Name (α ↑ ℓ) (β ↑ ℓ)
cmpName↑ ℓ Γ x y with x <N ℓ | y <N ℓ
... | inj1 x' | inj1 y' = x' ==N y'
... | inj2 x' | inj2 y' = Γ (x' ⋅N ℓ) (y' ⋅N ℓ)
... | -       | -       = false

```

The function `cmpTmD` structurally compares two terms in a simple way, only keeping track of the number of traversed binders and calling `cmpName↑` at variables:

```

cmpTmD : ∀ {α β} → |Cmp| Name α β → |Cmp| TmD α β
cmpTmD {α} {β} Γ = go 0 where
  go : ∀ ℓ → |Cmp| TmD (α ↑ ℓ) (β ↑ ℓ)
  go ℓ (V x)      (V y)      = cmpName↑ ℓ Γ x y
  go ℓ (t · u)    (v · w)    = go ℓ t v ∧ go ℓ u w
  go ℓ (λ t)      (λ u)      = go (suc ℓ) t u
  go ℓ (Let t u) (Let v w) = go ℓ t v ∧ go (suc ℓ) u w
  go - - -      = false

```

By instantiating the name comparator Γ by the homogeneous name comparison function `_==N_` we get an equality test on nameless terms. While this function works on open terms, they are requested to inhabit a common world.

```

_==TmD_ : ∀ {α} → |Cmp| TmD α α
_==TmD_ = cmpTmD _==N_

```

However in the particular case of comparing terms in a common world, we can define a simpler and a little faster equality test on terms. To do so we exploit the fact that the name abstraction is canonical and thus the term comparison is kept homogeneous during the whole comparison. In particular the function `cmpName↑` is not used.

```

eqTmD : ∀ {α} → |Cmp| TmD α α
eqTmD (V x)      (V y)      = x ==N y
eqTmD (t · u)    (v · w)    = eqTmD t v ∧ eqTmD u w
eqTmD (λ t)      (λ u)      = eqTmD t u
eqTmD (Let t u) (Let v w) = eqTmD t v ∧ eqTmD u w
eqTmD - - -      = false

```

5.4 Logical Relations for de Bruijn indices

5.4.1 Relations for the de Bruijn world operations

We now have to define how our operation on worlds fits the logical relation. Figure 5.4 depicts how the operators `[[+1]]` and `[[↑1]]` apply to the graph of a relation.




Figure 5.4: Shifting versus adding example

$_ [+1]$: ($\llbracket \text{World} \rrbracket \llbracket \rightarrow \rrbracket \llbracket \text{World} \rrbracket$) $_ +1 _ +1$
 $\alpha_r \llbracket +1 \rrbracket \stackrel{\text{def}}{=} \{ (x+1, y+1) \mid (x, y) \in \alpha_r \}$ -- not proper Agda

$_ \llbracket \uparrow 1 \rrbracket$: ($\llbracket \text{World} \rrbracket \llbracket \rightarrow \rrbracket \llbracket \text{World} \rrbracket$) $_ \uparrow 1 _ \uparrow 1$
 $_ \llbracket \uparrow 1 \rrbracket \alpha_r = \llbracket \text{zero}^B \rrbracket \llbracket < \rrbracket (\alpha_r \llbracket +1 \rrbracket)$

5.4.2 De Bruijn functions fit the relation

We now give a short overview of the proofs needed to show that our functions fit the relation. Thanks to the definition of $_ [+1]$, $\text{add}^N 1$ and $\text{subtract}^N 1$ fit the relation. These two are later extended to $\text{add}^N k$ and $\text{subtract}^N k$ by an induction on k . Here are for instance the type signatures for $\llbracket \text{add}^N \rrbracket$ and $\llbracket \text{cmp}^N \rrbracket$:

$\llbracket \text{add}^N \rrbracket$: ($\forall \langle \alpha_r : \llbracket \text{World} \rrbracket \quad \rangle \llbracket \rightarrow \rrbracket$
 $\langle k_r : \llbracket \mathbb{N} \rrbracket \quad \rangle \llbracket \rightarrow \rrbracket$
 $\llbracket \text{Name} \rrbracket \alpha_r \llbracket \rightarrow \rrbracket$
 $\llbracket \text{Name} \rrbracket (\alpha_r \llbracket +^W \rrbracket k_r)$) $\text{add}^N \text{add}^N$

$\llbracket \text{cmp}^N \rrbracket$: ($\forall \langle \alpha_r : \llbracket \text{World} \rrbracket \quad \rangle \llbracket \rightarrow \rrbracket$
 $\langle k_r : \llbracket \mathbb{N} \rrbracket \quad \rangle \llbracket \rightarrow \rrbracket$
 $\llbracket \text{Name} \rrbracket (\alpha_r \llbracket \uparrow \rrbracket k_r) \llbracket \rightarrow \rrbracket$
 $\llbracket \text{Name} \rrbracket (\llbracket \emptyset \rrbracket \llbracket \uparrow \rrbracket k_r) \llbracket \uplus \rrbracket$
 $\llbracket \text{Name} \rrbracket (\alpha_r \llbracket +^W \rrbracket k_r)$
 $\rangle \text{cmp}^N \text{cmp}^N$

The definition of cmp^N is not a simple induction on its first argument. It calls $_ \leq _$ (which does an induction) and returns a Boolean value. Based on this Boolean value, the function cmp^N returns either inj_1 or inj_2 (the constructors of $\llbracket \uplus \rrbracket$) with the same name (with a different proof). To prove $\llbracket \text{cmp}^N \rrbracket$ we show the equivalence with a simpler inductive function and show that this simpler function is in the relation. Thanks to the ex-




Figure 5.5: Various term types: from concrete to abstract

tensionality of the logical relation, no additional axiom is required to show that $\text{cmp}^{\mathbb{N}}$ fits the relation.

We have shown that the new inclusion rules (Figure 5.3) fit the relation.

5.4.3 On the strength of free theorems

Every well-typed term comes with a free theorem (section 4.2.1). However depending on the type of the term the strength of the theorem varies a lot. For instance at type \mathbb{N} the theorem only tell that the term is equal to itself. At type $\mathbb{N} \rightarrow \text{Bool}$ the theorem says no more than the function is deterministic. At type $\forall\{A : \text{Set}\} \rightarrow A \rightarrow A$ it says that the function behaves as the identity function, which is much stronger. We now give a few elements of what can affect the strength of free theorems in our context. We continue to use our type $\text{Tm}^{\mathbb{D}}$ to represent some data structures with names and binders but it could be any other. Moreover we focus on functions from terms to terms, where the input and the output worlds are the same. We discuss the strength of the abstraction level of the input or output term of such a function. From this perspective figure 5.5 depicts the relation corresponding to these types.

Various term types The weakest type we can give to a function taking a term with names and binders is $\forall\{\ell\} \rightarrow \text{Tm}^{\mathbb{D}} (\emptyset \uparrow \ell) \rightarrow \dots$. A term re-

ceived by such a function can have a statically unknown number of distinct free variables, but we know that these variables are comprised in the interval $[0 .. \ell-1]$. The free theorem of $\forall\{\ell\} \rightarrow \mathbf{Tm}^D (\emptyset \uparrow \ell) \rightarrow \mathbf{Tm}^D (\emptyset \uparrow \ell)$ says no more than “the function is deterministic” because the logical relation $\llbracket \mathbf{Tm}^D \rrbracket (\emptyset \uparrow \ell)$ is the identity. We also know because of typing that the resulting term cannot have ℓ -free occurrences.

The type $\mathbf{Tm} (\alpha \uparrow \ell)$ is more abstract than $\mathbf{Tm}^D (\emptyset \uparrow \ell)$, an unknown world is used instead of the empty world. However this type still shows a known amount of bindings. Then, a function of type $\forall\{\alpha \ell\} \rightarrow \mathbf{Tm}^D (\alpha \uparrow \ell) \rightarrow \mathbf{Tm}^D (\alpha \uparrow \ell)$ has a stronger free theorem than the previous one. It says that the function commutes with a renaming of the variables in the world α (section 5.4.4). This is a common type to deal with open terms under a partial environment:

```
some-op :  $\forall\{\alpha \ell\} \rightarrow (\mathbf{Fin} \ell \rightarrow \mathbf{Info}) \rightarrow \mathbf{Tm}^D (\alpha \uparrow \ell) \rightarrow \mathbf{Tm}^D (\alpha \uparrow \ell)$ 
some-op  $\Gamma$  t = {! omitted !}
```


The most abstract type for terms makes no assumptions on their free variables. This can be done by quantifying by an arbitrary world. The free theorem associated with the type $\forall \{\alpha\} \rightarrow \mathbf{Tm}^D \alpha \rightarrow \mathbf{Tm}^D \alpha$ is then stronger and says that the function commutes with any renaming of the free variables. This particular type is studied in detail in section 5.4.4.

Extra arguments Note that adding extra arguments to a function can drastically affect the strength of its free theorem. An extreme example is the type $\forall\{\alpha\} \rightarrow (\alpha \equiv \emptyset) \rightarrow \mathbf{Tm}^D \alpha \rightarrow \mathbf{Tm}^D \alpha$, which is ruined by its $\alpha \equiv \emptyset$ argument. This argument makes the type equivalent to $\mathbf{Tm}^D \emptyset \rightarrow \mathbf{Tm}^D \emptyset$ which has a trivial free-theorem. While the above example is extreme, this is an important point to watch out for when adding extra arguments to a function. Another example is the type $\forall\{\alpha \beta\} \rightarrow \mathbf{Tm}^D \alpha \rightarrow \mathbf{Tm}^D \beta \rightarrow \mathbf{Tm}^D \beta$ versus $\forall\{\alpha \beta\} \rightarrow (\mathbf{Name} \alpha \rightarrow \mathbf{Name} \beta \rightarrow \mathbf{Bool}) \rightarrow \mathbf{Tm}^D \alpha \rightarrow \mathbf{Tm}^D \beta \rightarrow \mathbf{Tm}^D \beta$. The first one cannot compare the free variables of the two given terms while the second can apply a user-supplied function to do so.

Shifting versus adding Last but not least using $_{+^W}$ instead of $_{\uparrow}$ significantly improves the strength of the associated free theorem. Consider the function `protectedAdd`:

```
protectedAdd :  $\forall \{\alpha\} \ell k \rightarrow \mathbf{Name} (\alpha \uparrow \ell)$ 
                $\rightarrow \mathbf{Name} (\alpha +^W k \uparrow \ell)$ 
protectedAdd  $\ell$  k = protect $\uparrow$   $\ell$  (add $^N$  k)
```


The following diagram depicts the graph behind the function `protectedAdd ℓ k`:



Consider now a weaker type, namely:

$$\begin{aligned} \text{protectedAdd}\uparrow &: \forall \{\alpha\} \ell \mathbf{k} \rightarrow \text{Name} (\alpha \uparrow \ell) \\ &\quad \rightarrow \text{Name} (\alpha \uparrow \mathbf{k} \uparrow \ell) \\ \text{protectedAdd}\uparrow \ell \mathbf{k} &= \text{protect}\uparrow \ell (\text{add}^{\mathbf{M}} \uparrow \mathbf{k}) \end{aligned}$$

We depict the graph behind the function `protectedAdd` \uparrow . The graph is actually the same than the graph for the function `protectedAdd`. We extend the set of nodes to match the codomain stated by the type.




We simply replaced the occurrence of `_W_` by `_ \uparrow _`. The consequences of this change are disastrous: this type allows more behaviors for its functions. Indeed the `addN k` function can be given the latter type (but not the former) when using the appropriate inclusion witness to exploit the commutativity of `_ \uparrow _`:

```

unprotectedAdd : ∀ {α} k l → Name (α ↑ l)
                → Name (α ↑ k ↑ l)

unprotectedAdd k l
  = coerceN (⊆-exch-↑-↑' l k) ∘ addN↑ k
  -- this is ok since (α ↑ k) ↑ l ≡ (α ↑ l) ↑ k
  -- whereas          (α +W k) ↑ l ≠ (α ↑ l) +W k

```



In our previous work [Pouillard and Pottier, 2010] we had a function called `shiftName` which was incorrectly defined with the `unprotectedAdd` behavior. Indeed the correct expected behavior is the one of `protectedAdd↑`. In this paper we had only the `↑` operation, and thus we missed the finer type we can give to `protectedAdd`. It took us a long time to discover this mistake, since we thought that our logical relation argument was enough. The logical relations proofs were indeed correct, the free-theorem were also correct but were not always “saying” what we expected. In summary we want to highlight that all the logical relations results have to be taken with great care. A weaker function type than expected can ruin the intended informal properties.

5.4.4 Using logical relations and parametricity

To formally show that a world-polymorphic function `f` of type $\forall \{\alpha\} \rightarrow \text{Tm}^D \alpha \rightarrow \text{Tm}^D \alpha$ commutes with a renaming of the free variables, we proceed as follows. First we recall the natural definition of logical relation on the type `TmD`. Second we present the type `Ren` of renamings as injective functions. Third renaming is shown equivalent to being related at the type `TmD`. Finally we prove our commutation lemma by using the free-theorem associated to the function `f`.

The logical relation for the type `TmD` is mechanical. It states that two terms are related if they have the same structure and related free variables:

```

data [[TmD]] {α1 α2} (αr : [[World]] α1 α2) :
  TmD α1 → TmD α2 → Set where
  [[V]]      : ∀ {x1 x2} (xr : [[Name]] αr x1 x2)
              → [[TmD]] αr (V x1) (V x2)
  -[[·]]-    : ∀ {t1 t2 u1 u2}
              (tr : [[TmD]] αr t1 t2)
              (ur : [[TmD]] αr u1 u2)
              → [[TmD]] αr (t1 · u1) (t2 · u2)
  [[λ]]     : ∀ {t1 t2} (tr : [[TmD]] (αr [[↑1]]) t1 t2)
              → [[TmD]] αr (λ t1) (λ t2)
  [[Let]]   : ∀ {t1 t2 u1 u2}
              (tr : [[TmD]] αr t1 t2)
              (ur : [[TmD]] (αr [[↑1]]) u1 u2)
              → [[TmD]] αr (Let t1 u1) (Let t2 u2)

```

In the following, we choose a “renaming” to be an injective function Φ of names to names. The type of a renaming that maps names in the world α to names in the world β is $\text{Ren } \alpha \beta$. Note that there is a slight difference between a renaming from α to β ($\text{Ren } \alpha \beta$) and the worlds α and β being “related” ($[[\text{World}]] \alpha \beta$). On renamings, we define three functions $\langle _ \rangle^N$, $\langle _ \rangle^W$ and $\langle _ \rangle^{\text{Tm}}$ which respectively turn a renaming into a function of names to names, into a relation between worlds, and into a function of terms to terms.

```

Ren : (α1 α2 : World) → Set
Ren α1 α2 = Injection (Nm α1) (Nm α2)
  -- * Injection From To is the set of the injective functions
  --   from the setoid From to the setoid To.
  -- * Nm α is the setoid on Name α

⟨ \_ ⟩N : ∀ {α β} → Ren α β → (α →N β)
⟨ Φ ⟩N = {! omitted !} -- Projects out the function over names

⟨ \_ ⟩W : ∀ {α β} → Ren α β → [[World]] α β
⟨ \_ ⟩W {α} {β} Φ =  $\mathcal{R}$  ,  $\mathcal{R}$ -pres-≡
  where  $\mathcal{R} : \text{Name } \alpha \rightarrow \text{Name } \beta \rightarrow \text{Set}$ 
         $\mathcal{R} \ x \ y = \langle \Phi \rangle^N \ x \equiv y$ 
         $\mathcal{R}$ -pres-≡ : Preserve-≡  $\mathcal{R}$ 
         $\mathcal{R}$ -pres-≡ = {! omitted !}

⟨ \_ ⟩Tm : ∀ {α β} → Ren α β → (TmD α → TmD β)
⟨ Φ ⟩Tm = renameTmD ⟨ Φ ⟩N

```

We now observe that given a renaming Φ , it is equivalent for two terms t_1 and t_2 to be related by $[[\text{Tm}^D]] \langle \Phi \rangle^W$ and for t_2 to be equal to t_1 renamed

with Φ . Put differently at the type \mathbf{Tm} the relation coincides with the α -equivalence (using the identity renaming for Φ).

$\mathbf{Rename} : \forall \{\alpha \beta\} \rightarrow \mathbf{Ren} \alpha \beta \rightarrow \mathbf{Tm}^D \alpha \rightarrow \mathbf{Tm}^D \beta \rightarrow \mathbf{Set}$
 $\mathbf{Rename} \Phi \ t_1 \ t_2 = \langle \Phi \rangle^{\mathbf{Tm}} t_1 \equiv t_2$

$\llbracket \mathbf{Tm}^D \rrbracket \Leftrightarrow \mathbf{Rename} : \forall \{\alpha \beta\} (\Phi : \mathbf{Ren} \alpha \beta) \rightarrow \llbracket \mathbf{Tm}^D \rrbracket \langle \Phi \rangle^W \Leftrightarrow \mathbf{Rename} \Phi$

Finally given a function f and a proof f_r that f is in the logical relation, we can show that any renaming Φ commutes with the function f . To prove this we apply our $\llbracket \text{Tm}^D \rrbracket \Leftrightarrow \text{rename}$ lemma in both directions and use f_r with the renaming Φ lifted as a $\llbracket \text{World} \rrbracket$.

```
-- Pointwise equality
f ≐ g = ∀ x → f x ≡ g x

ren-comm : (f  : ∀ {α} → TmD α → TmD α)
           (fr : (∀ (αr :  $\llbracket \text{World} \rrbracket \rrbracket \rightarrow \llbracket \llbracket \text{Tm}^D \rrbracket \rrbracket \alpha_r \llbracket \rightarrow \rrbracket \llbracket \llbracket \text{Tm}^D \rrbracket \rrbracket \alpha_r \rrbracket$ ) f f)
           {α β : World} (Φ : Ren α β)
           → ⟨ Φ ⟩Tm ∘ f ≐ f ∘ ⟨ Φ ⟩Tm
ren-comm f fr Φ t
  =  $\llbracket \llbracket \text{Tm}^D \rrbracket \rrbracket \Rightarrow \text{Rename } \Phi (f_r \langle \Phi \rangle^W (\text{Rename} \Rightarrow \llbracket \llbracket \text{Tm}^D \rrbracket \rrbracket \Phi \equiv .\text{refl}))$ 
```


Chapter 6

Variations and Related Work

This chapter covers various experimentations and connects our approach to more existing work. While some of the connections are not fully developed, they show the path to a complete connection and still give some insight.

The chapter is structured as follows. The first section explores the use of *atoms* as a separate kind of *global* names. In this section we explain the construction of Locally Nameless and Locally Named libraries where the type for global names is `Atom` and the type for local names is our type `Name`. We discuss the limitations of these approaches as programming interfaces, but underline the complementary role they can have when combined with our library.

In section 6.2 we describe another nameless approach from N.G. de Bruijn [de Bruijn, 1972] that is now called “de Bruijn levels”. We show how our existing library accommodates this style with only the addition of two inclusion rules.

In section 6.3 we describe a common interface for binders called links. This interface revives the presentation of our previous work [Pouillard and Pottier, 2010]. We show how the “unifying” part of this link-based approach can be integrated in our library. This approach enables greater code re-use without sacrificing much of the flexibility of our approach.

In section 6.4 we briefly explore how different approaches can be combined in order to be used at the same time. To illustrate this we combine nominal and de Bruijn indices in the same data type. This could allow for instance to convert a subterm to de Bruijn indices in order to apply a given algorithm.

In section 6.5 we cover a series of description languages used to describe “nominal types”. We show how the use of our explicit worlds can cope with pretty much everything these languages can describe. To do so we explain how each construct of these description languages can be expressed with worlds.

Section 6.6 discusses some more related work. We focus on relatively

recent work with which we have not yet drawn a concrete connection so far.

6.1 Dynamically stratified representations

Well known techniques such as *Locally Nameless* [Charguéraud, 2011, Aydemir et al., 2008], or *Locally Named* [Pollack et al., 2011, Sato and Pollack, 2010] can be qualified as “dynamically stratified representations”. Indeed these techniques use a double representation of variables, one for *local* variables, and one for *global* variables. The idea is to use local variables as the internal representation of the binding structure and to provide operations to switch back and forth from local to global variables. Name abstractions are introduced with the *close* operation. Given a subterm t and a global name X , a name abstraction is built by picking a local name x and replacing X by x in t , to form the body of the newly created name abstraction. Conversely, to *open* a name abstraction, one gives a global name X and an abstraction to be opened. The abstraction is internally made of a local name x and a subterm t . The opening operation returns t where X replaces x .

The main purpose of this double representation is to hide the handling of local names to the client and to let the client use only global names. To be sure to never encounter a local name, we want our terms to be *locally closed*, meaning without any free local variables. This property guarantees that starting from a locally closed term and properly opening abstractions, we never reach a local variable. Hence we only deal with global variables and avoid the hard work implied by local names and their abstractions.

While the representation for local names distinguishes Locally Nameless and Locally Named techniques, both use a nominal approach to represent global names. Hence we reuse *atoms* from section 2.1.1 and package them up as a little module. These atoms form our set of global names. This tiny module contains a type for atoms `Atom`, an equality test `_==A_`, and an injective function from natural numbers to atoms:

```
module Data.Atom where
  abstract
    Atom : Set
    Atom = ℕ

    _==A_ : (x y : Atom) → Bool
    _==A_ = _==ℕ_

    _A : ℕ → Atom
    _A = id
```

Since our library supports both de Bruijn indices and the nominal style,

the type `Name` can be used to represent local names. Our type for names is indexed by a world, hence internally, terms are indexed as well. However the manipulated terms being locally closed, the world is always the empty world in the end. Hence we happen to reach a local name, this name inhabits the empty world. Thanks to the function `Name \emptyset -elim` (based on `¬Name \emptyset`), we know that such a case is impossible.

6.1.1 Locally closed terms

To illustrate these techniques and how we can use our library to develop these techniques, we use a type for terms like we have done so far. However we do not define the type `TmA` yet (which is not the type `TmA` from section 2.1.1): instead, we assume the following types and data constructors.

```
-- The type for terms
TmA : Set

-- The type for name-abstractions,
-- abstracting one name over one term
-- example: [x]t
AbsTm : Set

-- Constructor for ‘global’ variables
VA : (x : Atom) → TmA

-- Constructor for ‘local’ variables
VN : (x : Name  $\emptyset$ ) → TmA

-- Constructor for  $\lambda$ -abstractions
 $\lambda$  : (abs : AbsTm) → TmA

-- Function application
_·_ : (t u : TmA) → TmA

-- Local definition
-- example: ‘let x = t in u’  $\Rightarrow$  Let t [x]u
Let : (t : TmA) → (abs : AbsTm) → TmA
```

In order to deal with the type `AbsTm`, the two functions `openTm` and `closeTm` are given to eliminate and introduce values of type `AbsTm`.

```
closeTm : Atom → TmA → AbsTm
openTm : Atom → AbsTm → TmA
```

Generally when writing the function `openTm`, it is no more difficult to plug a term than it is to plug an atom. This is especially true since the plugged term is locally closed, since no capture can happen. Hence we prefer to define the function `openSubstTm` for each implementation and define `openTm` generically.

```
openSubstTm : TmA → AbsTm → TmA
```

```
openTm : Atom → AbsTm → TmA
openTm = openSubstTm ∘ VA
```

6.1.2 Building terms

To build terms easily one can build smart constructors hiding the underlying use of `closeTm`. Hence we introduce two functions `λ'` and `let'` to build λ -abstractions and local definitions in a more natural way. In particular in the case of `let'` this is the occasion to put the arguments in the expected order:

```
λ' : Atom → TmA → TmA
λ' x t = λ (closeTm x t)
```

```
let' : Atom → TmA → TmA → TmA
let' x t u = Let t (closeTm x u)
```

We can now build various terms quite easily:

```
idT : TmA
idT = λ' x (VA x)
  where x = 0A
```

```
falseT : TmA
falseT = λ' x (λ' x (VA x))
  where x = 0A
```

```
trueT : TmA
trueT = λ' x (λ' y (VA x))
  where x = 0A
       y = 1A
```

```
apT : TmA
apT = λ' x (λ' y (VA x · VA y))
  where x = 0A
       y = 1A
```

```

 $\Omega^T : \text{Tm}^A$ 
 $\Omega^T = \text{let}' \delta \ ( \lambda' x \ (V^A x \cdot V^A x) )$ 
            $(V^A \delta \cdot V^A \delta)$ 
  where  $\delta = 0^A$ 
          $x = 1^A$ 

```

Opening and closing operations need to deal with local variables. However any other operation could be built with only `openTm` and `closeTm`, hence without any knowledge of the local variables. Using `openTm` and `closeTm` also has three main limitations: convenience, performance and non-structural recursion. Since all of them apply on most functions over terms we take a concrete example with a function `size` which tries to count the number of constructors in a term:

```

size : TmA → N
size (VA _) = 1
size (fct · arg) = 1 + size fct + size arg
size (λ abs) = 1 + size (openTm (0A) abs)
size (Let t abs) = 1 + size t + size (openTm (0A) abs)
size (VN x) = Name()-elim x

```

Convenience While `closeTm` can be hidden with smart constructors, `openTm` has to be explicitly used, which hinders pattern-matching. Indeed, without view-patterns (see [Wadler, 1987] for the seminal work on view patterns), pattern-matching has to be stopped at abstractions, in order to call `openTm`. The deeper the pattern, the worse the trouble. Moreover in this presentation we have to provide a global name to open the abstraction, which often requires an atom supply to give different atoms each time. In the `size` function we constantly open abstractions with the atom 0^A which is acceptable only because we do not care about the identity of variables in this example.

6.1.3 Performances

From a performance point of view `openTm` and `closeTm` require to traverse the whole subterm. They thus increase the asymptotic complexity of most algorithms. For a term of size n the asymptotic complexity of `openTm` and `closeTm` is naturally between $O(n)$ and $O(n \cdot \log(n))$. To our knowledge the complexity cannot be reduced much without losing properties such as persistence. In Caml [Pottier, 2006] the opening operation lazily performs the required substitution. We describe Caml and connect it to our work in section 6.5.2. For traditional algorithms the cost of opening (and

closing) is amortized but the presence of delayed substitutions makes the worst-case complexity exponential.

Let us look at how performance degrades on some algorithms. For instance an algorithm on a term that is simple enough to not involve the binding structure and that usually has a linear complexity, has a quadratic complexity when using `openTm` and `closeTm`. Algorithms to extract some information such as the size of a term (our function `size`) or the presence of some constructs in a term are of this kind. Algorithms to transform some non-binding constructs such as constant-folding, control-flow rewriting, assertions insertion or suppression, etc. are of this kind as well. Even some algorithms operating on the binding structure would benefit from avoiding `openTm` and `closeTm`. Here are for instance two functions in nominal style to transform back and forth, local `let` bindings into applied λ -abstractions. With our library no run-time cost is incurred and these two operations work in linear time:

```

module  $\lambda \Rightarrow \text{Let}$  where
   $\langle\langle - \rangle\rangle : \forall \{\alpha\} \rightarrow \text{Tm } \alpha \rightarrow \text{Tm } \alpha$ 
   $\langle\langle \lambda b t \cdot u \rangle\rangle = \text{Let } b \langle\langle u \rangle\rangle \langle\langle t \rangle\rangle$ 
   $\langle\langle V x \quad \quad \rangle\rangle = V x$ 
   $\langle\langle t \cdot u \quad \quad \rangle\rangle = \langle\langle t \rangle\rangle \cdot \langle\langle u \rangle\rangle$ 
   $\langle\langle \lambda b t \quad \quad \rangle\rangle = \lambda b \langle\langle t \rangle\rangle$ 
   $\langle\langle \text{Let } b t u \rangle\rangle = \text{Let } b \langle\langle t \rangle\rangle \langle\langle u \rangle\rangle$ 

```

```

module  $\text{Let} \Rightarrow \lambda$  where
   $\langle\langle - \rangle\rangle : \forall \{\alpha\} \rightarrow \text{Tm } \alpha \rightarrow \text{Tm } \alpha$ 
   $\langle\langle \text{Let } b t u \rangle\rangle = \lambda b \langle\langle u \rangle\rangle \cdot \langle\langle t \rangle\rangle$ 
   $\langle\langle V x \quad \quad \rangle\rangle = V x$ 
   $\langle\langle t \cdot u \quad \quad \rangle\rangle = \langle\langle t \rangle\rangle \cdot \langle\langle u \rangle\rangle$ 
   $\langle\langle \lambda b t \quad \quad \rangle\rangle = \lambda b \langle\langle t \rangle\rangle$ 

```

6.1.4 Non-structural recursion

The third main limitation of these approaches – at least from the point of view of a total language such as AGDA – is that `openTm` hinders structural recursion. Structural recursion is a syntactic criterion requiring to make recursive calls only on syntactic subterms of the input. Indeed if `t` is a subterm of a term `u`, `openTm s t` is not a subterm of `u`, preventing any recursive call. Of course we can claim for a better support for this kind of recursion, since `openTm` preserves the *size* of its input.

Our `size` function does not escape the rule and AGDA’s termination checker complains by coloring in red two offending recursive calls. In order to workaroud this limitation, we give a new parameter to `size` called `fuel`

of type \mathbb{N} , on which occurs the structural recursion. When we hit 0 we have not enough gas to proceed and return a dummy value. If given enough fuel the function behaves correctly. Generally it suffices to pick an upper bound for the fuel parameter. Here a good upper bound would be the size of the term...

```
-- ‘‘fuel-extended’’ to pass the termination checker
size : (fuel :  $\mathbb{N}$ )  $\rightarrow$   $\text{Tm}^A \rightarrow \mathbb{N}$ 
size (suc n) (VA _) = 1
size (suc n) (fct · arg) = 1 + size n fct + size n arg
size (suc n) ( $\lambda$  abs) = 1 + size n (openTm (0A) abs)
size (suc n) (Let t abs) = 1 + size n t + size n (openTm (0A) abs)
size (suc n) (VN x) = Name $\emptyset$ -elim x
size 0      _ = 0 -- dummy
```

6.1.5 Free atoms using openTm

A more interesting example is to compute the list of free variables, more precisely of free atoms. In order to define it independently of the representation of local name abstractions, we use `openTm`. Since the identity of atoms matters here, we take an atom supply as an argument to open abstractions with atoms that are fresh for the current context. Indeed there is no need to thread the atom supply here since we remove the atom afterward. A `fuel` argument is added to pass the termination checker. These two parameters would benefit from being hidden as a monad, at least for conveniently using and combining these functions. The function `fa` uses the function `rmA` which was defined in section 2.1.1.

```
module FreeAtoms where
  fa : (fuel atomSupply :  $\mathbb{N}$ )  $\rightarrow$   $\text{Tm}^A \rightarrow \text{List Atom}$ 
  fa (suc n) _ (VN x)      = Name $\emptyset$ -elim x
  fa (suc n) _ (VA x)      = [ x ]
  fa (suc n) s (fct · arg) = fa n s fct ++ fa n s arg
  fa (suc n) s ( $\lambda$  abs)
    = rmA (sA) (fa n (1 + s) (openTm (sA) abs))
  fa (suc n) s (Let t abs)
    = fa n s t ++
      rmA (sA) (fa n (1 + s) (openTm (sA) abs))
  fa 0      _ _ = [] -- dummy
```

6.1.6 Common parts of the internal representation

The type Tm^A is for locally closed terms. This type is based on a more general type for terms with potential free local names. This type is called PreTm and is indexed by a world. The only implementation specific part is the representation of name abstractions. The type AbsPreTm represents a name abstraction and can be implemented in various ways.

```
AbsPreTm : World → Set
```

```
data PreTm α : Set where
  VN   : (x : Name α)           → PreTm α
  VA   : (x : Atom)            → PreTm α
  _ · _ : (t u : PreTm α)       → PreTm α
  λ     : (t : AbsPreTm α)      → PreTm α
  Let   : (t : PreTm α) (u : AbsPreTm α) → PreTm α
```

```
AbsPreTm = {! implementation specific !}
```

Once the types PreTm and AbsPreTm are defined, defining the locally closed variants amounts to choose the empty world.

```
-- Locally closed terms
TmA : Set
TmA = PreTm ∅

-- Locally closed abstracted terms
AbsTm : Set
AbsTm = AbsPreTm ∅
```

6.1.7 Locally Nameless

We can now focus on specific details of local variables. We start with *Locally Nameless* [Charguéraud, 2011, Aydemir et al., 2008]. To do so we only have to choose the representation of name abstractions. We make use of the type SynAbs^D defined in Figure 6.1:

```
AbsPreTm = SynAbsD PreTm
```

Substituting definitions for AbsPreTm and SynAbs^D in the definition of PreTm yields the type Tm^D with an extra constructor for atoms:

$$\begin{aligned}
\text{SynAbs}^N \quad F \alpha &= \exists [\mathbf{b}] (F (\mathbf{b} \triangleleft \alpha)) \\
\text{SynAbs}^D \quad F \alpha &= F (\alpha \uparrow 1) \\
&\cong \exists [\beta] (\alpha \hookrightarrow^D \beta \times F \beta) \\
\text{SynAbs}^L \quad F \mathbf{b} \alpha &= F (\text{suc}^B \mathbf{b}) (\mathbf{b} \triangleleft \alpha) \\
\text{SemAbs} \quad F \alpha &= \forall \{ \beta \} \rightarrow \alpha \subseteq \beta \rightarrow F \beta \rightarrow F \beta \\
\text{SemAbs}^{+W} \quad F \alpha &= \forall \{ \beta \} \mathbf{k} \rightarrow (\alpha +^W \mathbf{k}) \subseteq \beta \rightarrow F \beta \rightarrow F \beta \\
\text{SynAbs}^{N\#} \quad F \alpha &= \exists [\mathbf{b}] (\mathbf{b} \# \alpha \times F (\mathbf{b} \triangleleft \alpha)) \\
&\cong \exists [\beta] (\alpha \hookrightarrow^N \beta \times F \beta)
\end{aligned}$$

Figure 6.1: Various forms of name abstraction

```

-- Once AbsPreTm and SynAbsD are expanded:
data PreTm α : Set where
  VN   : (x : Name α)           → PreTm α
  VA   : (x : Atom)            → PreTm α
  _ · _ : (t u : PreTm α)       → PreTm α
  λ     : (t : PreTm (α ↑ 1))   → PreTm α
  Let   : (t : PreTm α) (u : PreTm (α ↑ 1)) → PreTm α

```

We then define a module `TraversePreTm` (code omitted) like in section 5.3.4 to generically derive various operations including the coercion operation, and the substitution operation:

```

-- These two functions are derived from TraversePreTm
coercePreTm : ∀ {α β} → α ⊆ β → PreTm α → PreTm β
substPreTm  : ∀ {α β} → (Name α → PreTm β)
              → PreTm α → PreTm β

```

We can finally define `openSubstTm` and `closeTm`. The function `openSubstTm` is a simple use of substitution on terms using the function `substPreTm`. The argument to `substPreTm` replaces `0` by `t` using the function `exportWith`:

```

openSubstTm : TmA → AbsTm → TmA
openSubstTm t = substPreTm (exportWith t (Name-elim 0+1 ⊆ 0))

```

The function `closeTm` is essentially the combination of a coercion from the empty world to the world `0 ↑ 1` and a replace operation of an atom by a name referring to the newly introduced binder at the root. To make this replacing operation we simply carry on a name which we initialize with `0N` and increment when going down under the abstractions with `sucN↑`. While we could have fused these operations we did not for the clarity of the code.


```

closeTm : Atom → TmA → AbsTm
closeTm a t = ! (0N) (coerce∅PreTm t) where
  -- ‘‘! y t’’ replaces ‘‘a’’ by ‘‘y’’ in ‘‘t’’
  ! : ∀ {α} → Name α → PreTm α → PreTm α
  ! y (VA x) = if x ==A a then VN y else VA x
  ! y (VN x) = VN x
  ! y (t · u) = ! y t · ! y u
  ! y (λ t) = λ (! (sucN↑ y) t)
  ! y (Let t u) = Let (! y t) (! (sucN↑ y) u)

```

6.1.8 Locally Named

Local names bindings can be changed from de Bruijn indices to a nominal style by updating `AbsPreTm`. We use `SynAbsN` from section 3.4 and Figure 6.1:

```
AbsPreTm = SynAbsN PreTm
```

The resulting type `PreTm` is similar to the type `Term` from section 3.4 and to the type `TmA` from section 2.1.1. However, name-abstractions are packed into pairs and there is an extra constructor for atoms.

As usual we define a traversal function along the lines of section 3.2 with an extra argument to deal with the constructor `VA`:

```

TraversePreTm.tr :
  ∀ {E} (E-app : Applicative E)
  {Env} (trKit : TrKit Env (E ∘ PreTm))
  (trAtm : ∀ {α β} → Env α β → Atom → E (PreTm β))
  {α β} → Env α β → PreTm α → E (PreTm β)

```

From this traversal we obtain various operations including substitution:

```

substPreTm : ∀ {α β} → Supply β → (Name α → PreTm β)
           → PreTm α → PreTm β

```

From the function `substPreTm` we quickly get the function `openSubstTm`:

```

openSubstTm : TmA → AbsTm → TmA
openSubstTm t (-, u)
  = substPreTm (0S) (exportWith t Name∅-elim) u

```

The function `closeTm` is a bit more challenging. We define a new traversal kit called `renameKit+⊆` on a new environment type `SubstEnv+⊆`. This

environment type packages a renaming environment and an inclusion witness to maintain the fact that we freshened the output term and hence can import anything from the root output world.

```
record SubstEnv+⊆ Res β₀ α β : Set where
  constructor -,-
  field
    ren : SubstEnv Res α β
    pf  : β₀ ⊆ β
  open SubstEnv ren public
```

This traversal kit behaves such as the renaming kit (`renameKit`) on page 65:

```
renameKit+⊆ : ∀ β₀ → TrKit (SubstEnv+⊆ Name β₀) Name
renameKit+⊆ β₀ = mk trName trBinder extEnv
  where
    Env = SubstEnv+⊆ Name β₀
    open SubstEnv+⊆
    extEnv : ∀ {α β} b (Δ : Env α β) → Env (b < α) (seedB Δ < β)
    extEnv b (ren , pf) =
      (TrKit.extEnv renameKit b ren ,
       ⊆-trans pf (SubstEnv.seed⊆ ren))
```

Finally `closeTm` is defined using the traversal function `TraversePreTm.tr`, the kit `renameKit+⊆`, and the special behavior on atoms defined in the local function `trAtm`:

```
closeTm : Atom → TmA → AbsTm
closeTm a t = 0B , tr (renameEnvS 1 , ⊆-refl) t where
  β₀ = -

trAtm : ∀ {α β} → SubstEnv+⊆ Name β₀ α β → Atom → PreTm β
trAtm Δ a' = if a ==A a'
  then VN (coerceN (SubstEnv+⊆.pf Δ) (0N))
  else VA a'

kit = mapKit id VN (renameKit+⊆ β₀)
tr  = TraversePreTm.tr id-app kit trAtm
```

While we have a construction for Locally Named terms, this is not yet exactly the approach described by Sato and Pollak [Pollack et al., 2011, Sato and Pollack, 2010]. Indeed we have not yet integrated the height function which yields a different function to close terms. This means that we do not

get the nice properties either, for instance our Locally Named representation is not canonical. However the canonical representation is less crucial in our setting since our interface prevents from observing these representation details.

6.1.9 When does this stratified technique pay off?

So far we described these “dynamically stratified representations”, explained how we build them on top of our library, and explained the limitations of an exclusive usage of open/close functions. We claim that this technique can be of great effect when the static discipline of our library becomes too heavy. Contrary to most of the Locally Nameless implementations we argue that the name-abstraction should not be an abstract type. In particular with our library, no safety is lost by keeping name-abstraction concrete since we have precise worlds. Moreover, as we have shown, adding atoms to represent external names is the only change we have to make beforehand to our data type to support this technique.

Generalizing from atoms to arbitrary constants We advise users of our library to have a case for constants in their representation of terms for instance. These “constants” could be numbers, atoms, or a type parameter to be chosen afterward. We illustrate with two situations how this can be used effectively.

In the first situation, the client builds a function `f` which requires an extra environment argument to work on open terms. In the second situation the client is only able to define a function `f-c` for closed terms. One can often reduce the first situation to the second if we have an environment for the empty world. In these situations we end up with a function working only on closed terms. However if we pick the type of constants to be atoms, we have a function `f-lc` working on locally closed terms.

```
-- Terms parameterized by constants
Tm : Set → World → Set

-- We suppose that producing an initial environment
-- for arbitrary worlds is difficult.
Env : World → World → Set
```

```

-- We suppose the environment for empty worlds
-- to be easy to produce.
ε : Env ∅ ∅

-- Situation 1: A function on open terms
f : ∀ {α β Cst} → Env α β → Tm Cst α → Tm Cst β

-- Situation 2: Like f but works on closed terms
f-c : ∀ {Cst} → Tm Cst ∅ → Tm Cst ∅
f-c = f ε

-- Locally closed terms
LCTm = Tm Atom ∅

-- Like f but on locally closed terms
f-lc : LCTm → LCTm
f-lc = f-c

```

As we can see with the definitions of `f-c` and `f-lc` not much happened. In this example we are indeed supposing that nothing is done on constants and that this is correct to treat them parametrically. In this case what has been done is simply to tell the user of the function to convert all the free variables into atoms before calling the function.

We can spare the conversion to the user and do it ourselves. To do so, we show that given enough traversal functions on terms the two following types are equivalent:

$$\begin{aligned} & \forall \{Cst\} \rightarrow Tm Cst \emptyset \rightarrow Tm Cst \emptyset \\ & \cong \\ & \forall \{\alpha Cst\} \rightarrow Tm Cst \alpha \rightarrow Tm Cst \alpha \end{aligned}$$

We focus on the non-immediate direction, that is going from the first type to the second. In this situation we can even define this transformation as a higher-order function. Here we do not use atoms but keep names, and we preserve the original constants as well. We call the function `f` with the type `Name α ⊔ Cst` for constants, where `Cst` is the type of original constants and `α` is the world of free variables. Once the constants are preserved using the `bindTm` function, the free variables are turned into constants as well using the function `substTm`. Then the function `f` operates on the locally closed term. This yields a locally closed term as well that we coerce to the world `α`. Finally we restore the original constants and free variables using the function `bindTm` and the eliminator of sums: `[-, _]`. The following function `transform` is in a parameterized module `M` such that it only depends

on our library:

```

-- Hides the environment argument of a given function
-- over terms.
module M
  (Tm : Set → World → Set)
  (cst : ∀ {α Cst} → Cst → Tm Cst α)
  (var : ∀ {α Cst} → Name α → Tm Cst α)
  (substTm : ∀ {α Cst} → (Name α → Tm Cst ∅)
                → Tm Cst α → Tm Cst ∅)
  (bindTm : ∀ {α A B} → (A → Tm B α) → Tm A α → Tm B α)
  (coerce∅Tm : ∀ {α Cst} → Tm Cst ∅ → Tm Cst α)
  (f : ∀ {Cst} → Tm Cst ∅ → Tm Cst ∅)
where

transform : ∀ {α Cst} → Tm Cst α → Tm Cst α
transform
  = bindTm [ cst , var ] -- untag constants and variables
  o coerce∅Tm           -- closed terms inhabit any world
  o f                   -- operate on locally closed terms
  o substTm (cst o inj₂) -- tag variables as constants by 2
  o bindTm (cst o inj₁)  -- tag original constants by 1

```

We can see open and close functions like a way to move the border between the bound names and free names. When we do not call the open function and directly go under a name abstraction, it is as if we implicitly moved the border. This implicit move is one of the main sources of bugs in programs dealing with names and binders. In our system the border is reflected in the types using the worlds. Hence this distinction incurs no cost at run-time and still prevents mistakes caused by confusing bound and free names.

We have shown how to build dynamically stratified representations on top of our library. We have reused our technique to represent binders and local names, but used an atom type for global names. We have built both a Locally Nameless implementation and a basic Locally Named implementation including the traditional open and close functions. These techniques have shown their advantages as a tool for meta-theory, to reason about programming languages. We have shown the limitations that we think these techniques have for programming. We think these techniques are akin to dynamic typing. Indeed the tagging of names as local or global helps write functions without considering the free variables of the input. We have then shown how we can recover this benefit by having terms parameterized by constants. As soon as constants are handled parametrically we can transform any function over closed terms to a function on open terms. However

we recognize these dynamically stratified approaches may be simpler to use than our approach and still provide more safety than the bare approaches.

6.2 de Bruijn levels

Let us briefly mention another representation called “de Bruijn levels”. This is another nameless representation different from de Bruijn indices. λ -abstractions are implicitly labeled starting at 0 from the root, and variable occurrences simply refer to these labels. De Bruijn levels are actually very close to a nominal presentation since every reference to a given binder is represented identically. However, binders are not freely chosen, as in nominal style, but dictated by the position of the term starting from the root. We can accommodate this style by indexing terms by the next available binder:

```
data TmL s α : Set where
  V      : Name α → TmL s α
  λ     : (t u : TmL s α) → TmL s α
  λ     : TmL (sucB s) (s < α) → TmL s α
  Let   : TmL s α → TmL (sucB s) (s < α) → TmL s α
```

The type Tm^L is a variant of our type Tm which is not only indexed by a world but by a binder s as well. This binder s tells how far we are from “the root”. This binder is carried down by a constructor such as λ . At a λ -abstraction, the binder index is used and becomes bound in the subterm; moreover, the index is incremented in the subterm. Most functions work as in nominal style: for instance, the function fv requires no fundamental changes. The extra cost is this new index appearing in the types of terms. The benefit is that given an initial binder index, the representation is canonical.

6.2.1 Term comparison

Comparing terms in de Bruijn level style is no different than in nominal style. A name-comparator is carried down and extended with the same function extendNameCmp . In the end, the resulting function makes no restriction on the seeds of the input terms. This emphasises again that our comparison function makes no attempt to compare binders between the two arguments. Bound variables are compared positionally. Here is the signature for the function cmpTm^L :

$$\text{cmpTm}^L : \forall \{ \alpha \beta s_1 s_2 \} \rightarrow |\text{Cmp}| \text{ Name } \alpha \beta \\ \rightarrow (\text{Tm}^L s_1 \alpha \rightarrow \text{Tm}^L s_2 \beta \rightarrow \text{Bool})$$

However one may want to exploit the canonical structure of our representation for terms. We have already done so for de Bruijn indices with the function `eqTmD` in section 5.3.5. When both terms starts with the same binder index and world, then their types stays in synch. When reaching variable occurrences the equality test on names can be applied [Pouillard, 2011a].

```

_==TmL_ : ∀ {α s} → TmL s α → TmL s α → Bool
V x1      ==TmL V x2      = x1 ==N x2
(t1 · u1) ==TmL (t2 · u2) = t1 ==TmL t2 ∧ u1 ==TmL u2
λ t1      ==TmL λ t2      = t1 ==TmL t2
Let t1 u1 ==TmL Let t2 u2 = t1 ==TmL t2 ∧ u1 ==TmL u2
_          ==TmL _          = false

```

De Bruijn level operations are similar to nominal operations. A first difference is that binding constructs do not have to hold a binder since the binder is fixed by the index. This gives a nameless flavor to this presentation but our typed presentation shows where the binders have gone: in the types.

To illustrate how binders moved out of the terms we can look at snippets of functions with all the arguments made explicit:

```

fN : ∀ {α} → Tm α → ...
fN {α} (λ b t) = ...

fL : ∀ {b α} → TmL b α → ...
fL {b} {α} (λ t) = ...

```

The first snippet is in nominal style and the binder `b` is found in the term. The second snippet uses de Bruijn levels and the binder `b` comes as a separate (implicit) argument. In the end we have the same data in both snippets. We can see this as a good news, as it helps recover the information lost by this nameless representation.

For instance functions such as `fv` and `coerceTm` require almost no change to be adapted from nominal to de Bruijn levels.

6.2.2 Kits and traversals

We develop kits and traversals in a way similar to section 3.2. Since our terms are now indexed by a binder as well, environments are indexed by two worlds and two binders.

```

EnvTypeL : Set1
EnvTypeL = (s1 s2 : Binder) (α β : World) → Set

```

Kits are made of only two functions `trNameL` and `extEnvL`. There is no need for an equivalent of the function `trBinder` since binders now appear in indices and can thus be controlled through the environment type.

```
TrNameL : (Env : EnvTypeL) (Res : Binder → World → Set) → Set
TrNameL Env Res = ∀ {s1 s2 α β} → Env s1 s2 α β
                → Name α → Res s2 β
```

```
ExtEnvL : (Env : EnvTypeL) → Set
ExtEnvL Env =
  ∀ {s1 s2 α β} → Env s1 s2 α β
                → Env (sucB s1) (sucB s2) (s1 < α) (s2 < β)
```

```
record TrKitL (Env : EnvTypeL)
              (Res : Binder → World → Set) : Set where
  constructor mk
  field
    trNameL : TrNameL Env Res
    extEnvL : ExtEnvL Env
```

We have several traversal functions making use of a kit. We give only the most general of them. As in section 3.2 we abstract over an applicative functor, and the kit's result type is a term. This one suffices to derive all the functions presented in the following:

```
module TraverseATmL
  {E} (E-app : Applicative E)
  {EnvL} (kitL : TrKitL EnvL (λ s α → E (TmL s α)))
  where
    open Applicative E-app
    open TrKitL kitL

    trTmL : ∀ {α β s1 s2} → EnvL s1 s2 α β → TmL s1 α → E (TmL s2 β)
    trTmL Δ (V x)      = trNameL Δ x
    trTmL Δ (t · u)    = pure _ · _ * trTmL Δ t * trTmL Δ u
    trTmL Δ (λ t)      = pure λ _ * trTmL (extEnvL Δ) t
    trTmL Δ (Let t u) = pure Let * trTmL Δ t * trTmL (extEnvL Δ) u
```

The equivalent of `SubstEnv` is called `SubstEnvL`. This one does not require a `Supply` field, a fresh-for proof suffices. The `RenameEnvL` is a specialization to functions returning names.


```

record SubstEnvL Res (s1 s2 : Binder) (α β : World) : Set where
  constructor -, -
  field
    trNameL : Name α → Res s2 β
    s2#β     : s2 # β

```

```
RenameEnvL = SubstEnvL (const Name)
```

```
renameKitL : TrKitL RenameEnvL (const Name)
```

6.2.3 Derived functions

Along the lines of section 3.2, the kit `renameKitL` is extended to a kit `renameAKitL` which supports an effectful application. Like before, from these kits we derive several functions. Functions like `renameTmL`, `renameTmLA`, `renameTmL?`, `closeTmL`, and `substTmL` are built like their analog from section 3.2. The name supply argument is now just a fresh-for proof since the seed appears in the index. Here are the type signatures for these functions:

```

renameTmL : ∀ {α β s1 s2} → s2 # β → (Name α → Name β)
           → TmL s1 α → TmL s2 β

```

```

renameTmLA : ∀ {E} (E-app : Applicative E)
             {α β s1 s2} → s2 # β → (Name α → E (Name β))
             → TmL s1 α → E (TmL s2 β)

```

```

renameTmL? : ∀ {α β s1 s2} → s2 # β → (Name α →? Name β)
           → TmL s1 α →? (TmL s2 β)

```

```

closeTmL? : ∀ {s α} → TmL s α →? TmL s ∅

```

```

substTmL : ∀ {α β s s2} → s2 # β → (Name α → TmL s2 β)
           → (TmL s α → TmL s2 β)

```

We then lift addition and subtraction from names to terms using the function `renameTmL`. The resulting functions get interesting types. In particular the input seed is disconnected from the input seed and this these functions enables to re-level the terms:

$$\begin{aligned} \text{protectedAddTm}^L &: \forall \{\alpha \ s_1 \ s_2\} \ k \rightarrow s_2 \# (\alpha +^W k) \\ &\quad \rightarrow \text{Tm}^L \ s_1 \ \alpha \rightarrow \text{Tm}^L \ s_2 \ (\alpha +^W k) \\ \text{protectedAddTm}^L \ k \ \text{pf} &= \text{renameTm}^L \ \text{pf} \ (\text{add}^N \ k) \\ \text{protectedSubtractTm}^L &: \forall \{\alpha \ s_1 \ s_2\} \ k \rightarrow s_2 \# \alpha \\ &\quad \rightarrow \text{Tm}^L \ s_1 \ (\alpha +^W k) \rightarrow \text{Tm}^L \ s_2 \ \alpha \\ \text{protectedSubtractTm}^L \ k \ \text{pf} &= \text{renameTm}^L \ \text{pf} \ (\text{subtract}^N \ k) \end{aligned}$$

6.2.4 New inclusion rules

In order to tightly integrate de Bruijn levels, and more generally to mix nominal binders and `_+1` on worlds, we introduce two new inclusion rules. These two rules state that `_+1` distributes over `_<`:

$$\begin{aligned} \underline{\subseteq}\text{-dist-+1-<} &: \forall \{\alpha \ b\} \rightarrow (b \ < \ \alpha) \ +1 \ \subseteq \ (\text{suc}^B \ b) \ < \ (\alpha \ +1) \\ \underline{\subseteq}\text{-dist-<-+1} &: \forall \{\alpha \ b\} \rightarrow (\text{suc}^B \ b) \ < \ (\alpha \ +1) \ \subseteq \ (b \ < \ \alpha) \ +1 \end{aligned}$$

We generalize these two primitive rules to not only `1` but any constant:

$$\begin{aligned} \underline{\subseteq}\text{-dist-+-<} &: \forall \{\alpha \ b\} \ k \rightarrow (b \ < \ \alpha) \ +^W k \ \subseteq \ (b \ +^B k) \ < \ (\alpha \ +^W k) \\ \underline{\subseteq}\text{-dist-+-<} \ \text{zero} &= \underline{\subseteq}\text{-refl} \\ \underline{\subseteq}\text{-dist-+-<} \ (\text{suc} \ n) &= \underline{\subseteq}\text{-trans} \ (\underline{\subseteq}\text{-cong-+1} \ (\underline{\subseteq}\text{-dist-+-<} \ n)) \\ &\quad (\underline{\subseteq}\text{-dist-+1-<} \ _) \\ \underline{\subseteq}\text{-dist-<-+} &: \forall \{\alpha \ b\} \ k \rightarrow (b \ +^B k) \ < \ (\alpha \ +^W k) \ \subseteq \ (b \ < \ \alpha) \ +^W k \\ \underline{\subseteq}\text{-dist-<-+} \ \text{zero} &= \underline{\subseteq}\text{-refl} \\ \underline{\subseteq}\text{-dist-<-+} \ (\text{suc} \ n) &= \underline{\subseteq}\text{-trans} \ (\underline{\subseteq}\text{-dist-<-+1} \ _) \\ &\quad (\underline{\subseteq}\text{-cong-+1} \ (\underline{\subseteq}\text{-dist-<-+} \ n)) \end{aligned}$$

6.2.5 Addition and subtraction kits

The add kit (`AddL.kit`) enables to add a constant to all the names and binders of a given traversable term. The environment of this kit is made of two proofs, one on worlds the other on binders. This environment could hence be removed by an optimizing compiler. Once compiled the resulting function is not hindered by the safe handling of names and binders.

```
module AddL k where
  EnvL : EnvTypeL
  EnvL b1 b2 α β = (α +W k) ⊆ β
                    × (b1 +B k) ≡ b2
```

```

trNameL : TrNameL EnvL (const Name)
trNameL (pf , -) = coerceN pf ∘ addN k

kitL : TrKitL EnvL (const Name)
kitL = mk trNameL {! proof omitted !}

```

We then briefly mention a generalization of the previous kit which allows not only addition but subtraction as well. This kit works similarly and we just show the environment type and the signature of the function on terms:

```

AddSubtractEnvL s1 s2 α β = (s1 +B k1) ≡ (s2 +B k2)
                        × (α +W k1) ⊆ (β +W k2)

addSubtractTmL : ∀ k1 k2
  {α β s1 s2} → (s1 +B k1) ≡ (s2 +B k2)
                → (α +W k1) ⊆ (β +W k2)
                → TmL s1 α → TmL s2 β

```

From this general operation `addSubtractTmL` we derive several operations. We can coerce terms, add a constant, subtract a constant, change the seed index of a closed term, and combinations of those.

```

coerceTmL   : ∀ {s α β} → α ⊆ β → TmL s α → TmL s β
addTmL     : ∀ {α s} k → TmL s α → TmL (s +B k) (α +W k)
subtractTmL : ∀ {α s} k → TmL (s +B k) (α +W k) → TmL s α
rerootTmL  : ∀ {α} k1 k2 → TmL (k1B) ∅ → TmL (k2B) α

```

6.2.6 Conversion with nominal terms

Converting de Bruijn level terms to nominal could not be simpler. The code [Pouillard, 2011a] is a simple recursive copy where AGDA can even infer the binders to choose to build the nominal terms. When converting, we are simply moving a static information to a dynamic place. When reaching a binding construct the binder is known to be the index of the term. This binder also a valid binder to be used to build the resulting binding construct in nominal style. Here is the conversion function we obtain:

```

module TmL⇒TmN where
  ⟨⟨ _ ⟩⟩ : ∀ {α s} → TmL s α → TmN α
  ⟨⟨ V x      ⟩⟩ = V x
  ⟨⟨ λ t      ⟩⟩ = λ _ ⟨⟨ t ⟩⟩
  ⟨⟨ t · u    ⟩⟩ = ⟨⟨ t ⟩⟩ · ⟨⟨ u ⟩⟩
  ⟨⟨ Let t u  ⟩⟩ = λ _ ⟨⟨ u ⟩⟩ · ⟨⟨ t ⟩⟩

```

The conversion from nominal to de Bruijn levels is a bit more complicated. Indeed nominal bound variables have to be renamed to follow the discipline of levels. We implement this function by first writing a traverse function from nominal terms to de Bruijn level terms. Then we make use of the renaming kit with little changes. Here are the functions we obtain:

```

renameTmL⇒TmN : ∀ {α β s} → s # β → (Name α → Name β)
                    → (TmN α → TmL s β)
renameTmL⇒TmN = {! code omitted !}

convTmN⇒TmL : ∀ {s β} → s # β → TmN () → TmL s β
convTmN⇒TmL s#β = renameTmL⇒TmN s#β Name()-elim

convTmN⇒TmL() : ∀ s → TmN () → TmL s ()
convTmN⇒TmL() s = convTmN⇒TmL (s #())

```

Conclusion We have shown how de Bruijn levels can be implemented using our library. Our types are now indexed by a world and a binder. This binder is a seed for a binder supply. This seed is to be used for the binder of the next binding construct. For the subterms of binding constructs we increment the seed and extend the world in a nominal style using `<_<`. We have shown how close this style was from the nominal style and have derived operations on terms by applying techniques developed for the nominal style. This style also reuses tools we built for de Bruijn indices such as additions on worlds and names (`+W_` and `+N_`). The extra index can be cumbersome to track, but can also help, as in the equality test on terms. We took advantage of the representation to build a faster equality test and a faster addition function.

6.3 Links: Binders as World Relations

So far with our library, name-abstraction induced a particular form of world for the subterm. In nominal style, we use `<_<` and in de Bruijn indices style we use `<_↑1`. These forms of worlds induce the use of dependent types in nominal style and computation on types in de Bruijn indices style. In nominal style, dependent types occur since binders appear in worlds. In de Bruijn computation on the level of worlds (hence types) occur with the use of arbitrary shifts and additions on worlds $\alpha +^W k$ and $\alpha \uparrow \ell$. In our work of 2010 [Pouillard and Pottier, 2010], the interface relies only on world-polymorphism and since world are erased it avoids dependent types. After our work of 2010 we put aside the idea of links and have developed the

interface presented so far using dependent types leading to a more precise and concrete interface.

In this section we come back to a link based interface. We start by describing an interface for *binding-links* or simply *links*. A link is a type parameterized by two worlds which supports the following three operations: `weaken`, `strengthen?`, and `nameOf`. The first parameter of the link is called the outer world and the second the inner world. A link can be seen as a binder `b` plus a relation between the inner and the outer world. The outer world describes the binders before the introduction of `b` and the inner world describes the binders after its introduction. The first operation moves a name from the outer world to the inner world, it is a weakening. The second operation tries to move a name from the inner world to the outer world but may fail if we try to move the name being bound by the binding-link. The third operation reveals the name being bound, which belongs to the inner world. In AGDA we have the following record:

```
record Link ( _↔_ : (α β : World) → Set ) : Set where
  constructor mk
  field
    weaken      : ∀ {α β} → α ↔ β → Name α → Name β
    strengthen? : ∀ {α β} → α ↔ β → Name β →? Name α
    nameOf     : ∀ {α β} → α ↔ β → Name β
```

6.3.1 Terms and examples

Using links, one can define data types, like our type `Tm` for terms. As usual the interesting cases are the binding constructs. In particular in the construct `Let` we see how the use of inner and outer worlds control the defined scopes.

```
data Tm α : Set where
  V      : (x : Name α) → Tm α
  _·_    : (t u : Tm α) → Tm α
  λ      : ∀ {β} (l : α ↔ β) (t : Tm β) → Tm α
  Let    : ∀ {β} (l : α ↔ β) (t : Tm α) (u : Tm β) → Tm α
```

Using only the function `strengthen?` one can build our usual functions `fv` and `rm` to compute the list of free variables and remove a name from a list of names:

```

rm : ∀ {α β} → α ↔ β → List (Name β) → List (Name α)
rm - []          = []
rm ℓ (x :: xs) with strengthen? ℓ x
...             | nothing = rm ℓ xs
...             | just x'  = x' :: rm ℓ xs

fv : ∀ {α} → Tm α → List (Name α)
fv (V x)      = [ x ]
fv (t · u)    = fv t ++ fv u
fv (λ ℓ t)    = rm ℓ (fv t)
fv (Let ℓ t u) = fv t ++ rm ℓ (fv u)

```

We now define a comparator for terms. It follows precisely what we have done before and here again only the function `strengthen?` is necessary:

```

extCmpName : ∀ {α β α' β'} → |Cmp| Name α β → α ↔ α' → β ↔ β'
                    → |Cmp| Name α' β'

extCmpName f ℓ1 ℓ2 x1 x2
  with strengthen? ℓ1 x1 | strengthen? ℓ2 x2
... | just x1'           | just x2' = f x1' x2'
... | nothing            | nothing  = true
... | -                  | -       = false

cmpTm : ∀ {α β} → |Cmp| Name α β → |Cmp| Tm α β
cmpTm f (V x)      (V x') = f x x'
cmpTm f (t · u)    (t' · u') = cmpTm f t t' ∧ cmpTm f u u'
cmpTm f (λ ℓ t)    (λ ℓ' t') = cmpTm (extCmpName f ℓ ℓ') t t'
cmpTm f (Let ℓ t u) (Let ℓ' t' u') = cmpTm f t t'
                                        ∧ cmpTm (extCmpName f ℓ ℓ') u u'

cmpTm _ _ _ _ = false

_==Tm_ : ∀ {α} → |Cmp| Tm α α
_==Tm_ = cmpTm _==N_

```

6.3.2 Implementations

An appealing definition for nominal links would to pair a binder \mathbf{b} and a proof that β is equal to $\mathbf{b} \triangleleft \alpha$ this is called a weak link in [Pouillard and Pottier, 2010]. However to support the weakening operation a fresh-for proof is required. Hence a nominal link is made of a binder \mathbf{b} , a proof that $\mathbf{b} \# \alpha$, and a proof that β is equal to $\mathbf{b} \triangleleft \alpha$, this is called a strong link in [Pouillard and Pottier, 2010]. Our implementation reuses the type `Supply` which already packages a binder plus a fresh-for proof. Our nominal implementation

is straightforward. The weakening operation uses a coercion. The strengthening operation boils down to our function `exportN?`. The operation `nameOf` returns the underlying binder converted to a name. In AGDA we have the following module:

```

module Nom where
  record  $\_ \hookrightarrow \_$   $\alpha \beta$  : Set where
    constructor  $\_ , \_$ 
    field
      supply : Supply  $\alpha$ 
    open Supply supply
    field
       $\beta$ -def :  $\beta \equiv \text{seed}^B \triangleleft \alpha$ 

  weaken :  $\forall \{ \alpha \beta \} \rightarrow \alpha \hookrightarrow \beta \rightarrow \text{Name } \alpha \rightarrow \text{Name } \beta$ 
  weaken (supply ,  $\equiv$ .refl) = coerceN (Supply.seed $\subseteq$  supply)

  strengthen? :  $\forall \{ \alpha \beta \} \rightarrow \alpha \hookrightarrow \beta \rightarrow \text{Name } \beta \rightarrow ? \text{Name } \alpha$ 
  strengthen? ( ,  $\equiv$ .refl) = exportN?

  nameOf :  $\forall \{ \alpha \beta \} \rightarrow \alpha \hookrightarrow \beta \rightarrow \text{Name } \beta$ 
  nameOf ((seedB , -) ,  $\equiv$ .refl) = nameB seedB

  nomLinks : Link  $\_ \hookrightarrow \_$ 
  nomLinks = mk weaken strengthen? nameOf

```

We now do the same with de Bruijn indices. A link in de Bruijn style is just a proof that the inner world β is equal to the outer world α shifted by one. The weakening is the successor function and the strengthening the predecessor function. The function `nameOf` returns `0`:

```

module DeBruijn where
  _↔_ : (α β : World) → Set
  α ↔ β = β ≡ (α ↑1)

  weaken : ∀ {α β} → α ↔ β → Name α → Name β
  weaken ≡ .refl = sucM↑

  strengthen? : ∀ {α β} → α ↔ β → Name β →? Name α
  strengthen? ≡ .refl = predM?

  nameOf : ∀ {α β} → α ↔ β → Name β
  nameOf ≡ .refl = zeroN

  deBruijnLinks : Link _↔_
  deBruijnLinks = mk weaken strengthen? nameOf

```

6.3.3 Building links and terms

In order to build terms we need to introduce more operations, though. Indeed we have no initial links to use. To this end we introduce a simple interface to build links. This interface provides two operations: `init` and `next`. More precisely these operations work on links where the inner world is made existential. Here is our interface in AGDA:

```

record FreshLink {_↔_} (link : Link _↔_) : Set where
  constructor mk
  open Link link

  Fresh : World → Set
  Fresh α = ∃[ β ](α ↔ β)

  field
    -- An infinite ‘‘chain’’ of binding-links
    init   : Fresh ∅
    next   : ∀ {α} (x : Fresh α) → Fresh (proj1 x)

```

As we see in the type of the operation `next`, the links are chained. Let `fresh` be a value of type `Fresh α`. The outer world of `next fresh` is the inner world of `fresh`. Here is our straightforward nominal implementation of the `FreshLink` interface based on operations on name supplies:


```

module NomFresh where
  init : Fresh  $\emptyset$ 
  init = _ , 0s ,  $\equiv$ .refl

  next :  $\forall \{\alpha\}$  (x : Fresh  $\alpha$ )  $\rightarrow$  Fresh (proj1 x)
  next (._ , s ,  $\equiv$ .refl) = _ , sucs s ,  $\equiv$ .refl

  nomFreshLinks : FreshLink nomLinks
  nomFreshLinks = mk init next

```

The de Bruijn implementation is so trivial that AGDA infers most of the definition by unification:

```

module DeBruijnFresh where
  init : Fresh  $\emptyset$ 
  init = _ ,  $\equiv$ .refl

  next :  $\forall \{\alpha\}$  (x : Fresh  $\alpha$ )  $\rightarrow$  Fresh (proj1 x)
  next _ = _ ,  $\equiv$ .refl

  deBruijnFreshLinks : FreshLink deBruijnLinks
  deBruijnFreshLinks = mk init next

```

Building terms is a matter of allocating binding-links and weakening names such that they are in the expected world. To illustrate the construction of terms here are two examples:

```

idTm : Tm  $\emptyset$ 
idTm =  $\lambda$  x xTm
  where x = proj2 init
        xTm = V (nameOf x)

apTm : Tm  $\emptyset$ 
apTm =  $\lambda$  x ( $\lambda$  y (xTm . yTm))
  where x = proj2 init
        y = proj2 (next init)
        xTm = V (weaken y (nameOf x))
        yTm = V (nameOf y)

```

6.3.4 Kits and traversals

We now leverage our technique to build kits and traversals. Adapting the work that we did in the nominal style is mostly straightforward. Kits are

made of a function to operate on names and a function to operate on links. To operate on a link the function `commEnv`, given a link and an environment, returns both a new link and a new environment. The type `Comm` is the most difficult part of these traversals and is already described in our 2010 work [Pouillard and Pottier, 2010]:

```
Comm : ( _↗_ : (α β : World) → Set ) → Set
Comm _↗_ = ∀ {α β γ} → (α ↔ β × α ↗ γ)
           → ∃ [ δ ] (γ ↔ δ × β ↗ δ)
```

A traversal kit is then made of a function `trName` and a function `commEnv`. As before the kit is parameterized over the type of environments `Env` and the type of results `Res`.

```
record TrKit (Env : (α β : World) → Set)
             (Res : World → Set) : Set where
  constructor mk
  field
    trName   : ∀ {α β} → Env α β → Name α → Res β
    commEnv  : Comm Env
```

We now build the renaming kit. Again the solution is close to the nominal one even if it works for de Bruijn style links as well. Here the supply is of type `Fresh`. The function `importFun` takes two links and lifts a function over names in outer worlds α and β to a function over names in inner worlds γ and δ . The function `importFun` then form the core part of the commutation function for renaming environments. As with nominal kits the environment is a mapping of names from the input world to the output world. This mapping freshens the bound variables and uses a name supply to do so.

```
importFun : ∀ {α β γ δ} → α ↔ γ
           → β ↔ δ
           → (Name α → Name β)
           → (Name γ → Name δ)

importFun l1 l2 f x
  with strengthen? l1 x
... | nothing = nameOf l2
... | just x' = weaken l2 (f x')
```

```
RenameEnv : (α β : World) → Set
```

```

RenameEnv  $\alpha \beta = \text{Fresh } \beta \times (\text{Name } \alpha \rightarrow \text{Name } \beta)$ 

renameKit : TrKit RenameEnv Name
renameKit = mk proj2 comm
  where
    comm : Comm RenameEnv
    comm (x , ( - , y ) , f)
      = - , y , next ( - , y ) , importFun x y f

```

We now define the traversal function for our particular type for terms. As usual the environment is carried down the recursive function. The function `trName` is called at variables. The function `commEnv` is used at name-abstractions. Given the link found in the name-abstraction and the current environment, the function `commEnv` returns both a new link to close the resulting name-abstraction and a new environment to continue the traversal on subterms in scope of the name-abstraction:

```

module Tr {Env : ( $\alpha \beta$  : World)  $\rightarrow$  Set}
  (trKit : TrKit Env Tm)

  where
    open TrKit trKit

    tr :  $\forall \{\alpha \beta\} \rightarrow \text{Env } \alpha \beta \rightarrow \text{Tm } \alpha \rightarrow \text{Tm } \beta$ 
    tr  $\Gamma$  (V x)      = trName  $\Gamma$  x
    tr  $\Gamma$  (t . u)   = tr  $\Gamma$  t . tr  $\Gamma$  u
    tr  $\Gamma$  ( $\lambda \ell$  t) with commEnv ( $\ell$  ,  $\Gamma$ )
    ...                | - ,  $\ell'$  ,  $\Gamma' = \lambda \ell'$  (tr  $\Gamma'$  t)
    tr  $\Gamma$  (Let  $\ell$  t u) with commEnv ( $\ell$  ,  $\Gamma$ )
    ...                | - ,  $\ell'$  ,  $\Gamma' = \text{Let } \ell'$  (tr  $\Gamma$  t) (tr  $\Gamma'$  u)

```

We put our traversal function at work by using the renaming kit to build a renaming function. Since our renaming kit is producing names and our traversal function expecting a kit which produces terms we have to lift it. The function `kitVar` is a kit transformer applying the constructor `V` to the result of the input function `trName`. Then the function `renameTm` gives the kit and the initial arguments to the traversal function.

```

kitVar :  $\forall \{\text{Env}\} \rightarrow \text{TrKit Env Name} \rightarrow \text{TrKit Env Tm}$ 
kitVar (mk trName commEnv) = mk ( $\lambda \Gamma \rightarrow \text{V} \circ \text{trName } \Gamma$ ) commEnv

renameTm :  $\forall \{\alpha \beta\} \rightarrow \text{Fresh } \beta \rightarrow (\text{Name } \alpha \rightarrow \text{Name } \beta)$ 
   $\rightarrow \text{Tm } \alpha \rightarrow \text{Tm } \beta$ 
renameTm fr f = Tr.tr (kitVar renameKit) (fr , f)

```

With the function `renameTm` we can lift the function `weaken` from names to terms. In this function the input link serves both as the initial weakening function for free names and as a name supply for the freshening purposes. Here is the function `weakenTm`:

```
weakenTm : ∀ {α β} → α ↦ β → Tm α → Tm β
weakenTm ℓ t = renameTm (next (- , ℓ)) (weaken ℓ) t
```

Our last application is a capture avoiding substitution function. To simplify matters we build a substitution kit specialized for our term type. The only important difference with the renaming kit is the use of the function `weakenTm` instead of `weaken`. The kit is then given to our traversal function to produce the function `substTm`. We then specialize our substitution function to replace a name/link by a closed term in a term with just this free variable. Finally we write the function `β-red` which reduces the β -redex on the root of the term if there is one.

```
SubstEnv : (α β : World) → Set
SubstEnv α β = Fresh β × (Name α → Tm β)
```

```
substKit : TrKit SubstEnv Tm
substKit = mk proj2 comm
  where
    importFunTm : ∀ {α β γ δ} → α ↦ γ
                  → β ↦ δ
                  → (Name α → Tm β)
                  → (Name γ → Tm δ)

    importFunTm ℓ1 ℓ2 f x
      with strengthen? ℓ1 x
    ... | nothing = V (nameOf ℓ2)
    ... | just x' = renameTm (next (- , ℓ2)) (weaken ℓ2) (f x')

    comm : Comm SubstEnv
    comm (x , (- , y) , f)
      = (- , y , (next (- , y) , importFunTm x y f))
```

```
substTm : ∀ {α β} → Fresh β → (Name α → Tm β) → Tm α → Tm β
substTm fr f = Tr.tr substKit (fr , f)
```

```
_-[:=-] : ∀ {α} → Tm α → ∅ ↦ α → Tm ∅ → Tm ∅
```

```
t [ x := u ] = substTm init (strengthenWith u V x) t
```

```
 $\beta$ -red : Tm  $\emptyset \rightarrow$  Tm  $\emptyset$ 
 $\beta$ -red ( $\lambda$   $\ell$  t . u) = t [  $\ell := u$  ]
 $\beta$ -red t = t
```

Conclusion By instantiating links to be either nominal links or de Bruijn links we obtain the corresponding implementation of the β -reduction of λ -terms. In contrast to our previous work [Pouillard and Pottier, 2010] we simplified the presentation to only retain one kind of binding-link. To reuse the terminology of [Pouillard and Pottier, 2010] we keep only the strong links. We think the purpose of this binding-link abstraction is to be more focused on a clean and simple interface than a comprehensive and precise interface. Indeed precision has been achieved in our library through binders, world operations, world inclusions, and operations on names. The binding-link abstraction hides details about the representation. While a nominal binder can shadow the previous one, a link, thanks to the fresh-for proof, cannot. While a de Bruijn binder is canonical, a link is not. While we can perform arithmetic with de Bruijn indices in our library, with links we cannot. Yet terms with binding-links can be equipped with kits and traversals. In short the paradigm is still very expressive but less flexible.

6.4 Combining nominal and nameless styles

The support for de Bruijn indices and de Bruijn levels, has been introduced by only adding new items to the interface. Neither the interface, the implementation nor the logical relation need to be altered, only extended.

Here we present one of the multiple combinations that could be made. We have two constructors for λ -abstraction, one in nominal style and one in de Bruijn indices style. Right after, we show the function `fv` which produces no surprises.

```
data Tm  $\alpha$  : Set where
  V      : (x : Name  $\alpha$ )  $\rightarrow$  Tm  $\alpha$ 
   $\lambda^N$  :  $\forall$  b (t : Tm (b  $\triangleleft$   $\alpha$ ))  $\rightarrow$  Tm  $\alpha$ 
   $\lambda^D$    : (t : Tm ( $\alpha \uparrow 1$ ))  $\rightarrow$  Tm  $\alpha$ 
  _·_    : (t u : Tm  $\alpha$ )  $\rightarrow$  Tm  $\alpha$ 
```

```

fv : ∀ {α} → Tm α → List (Name α)
fv (V x)      = [ x ]
fv (fct · arg) = fv fct ++ fv arg
fv (λN b t)   = rm b (fv t)
fv (λD t)     = rm0 (fv t)

```

We now show two variants of a term which combines those two styles of bindings. One thing to notice is that even if a name is introduced in a nominal style, one has to shift it to cross de Bruijn style binder. The first variant binds the first name in nominal style and the second name in the de Bruijn style. To reach the nominal binder from the inside `sucN↑` is used. The second variant binds the first name in de Bruijn style and the second in nominal style. To cross the nominal binder, the de Bruijn occurrence needs to be coerced showing that two worlds are included:

```

apTm1 : ∀ {α} → Tm α
apTm1 = λN x (λD (V (sucM↑ (nameB x)) · V (0N)))
  where x = 42B

apTm2 : Tm ∅
apTm2 = λD (λN x (V (coerceN pf (0N)) · V (nameB x)))
  where x = 42B
       pf = {! proof omitted !}

```

An attractive feature of this combined style would be to partially convert to de Bruijn indices. Once converted, the canonical structure can be used. Since the conversion to de Bruijn indices requires to maintain a mapping from nominal names to de Bruijn indices we build a more general renaming function. Our renaming function (called `ren`) is similar to `renameTmD` (section 5.3.4) but accepts nominal bindings as well. It is similar to `renameTm` (section 3.2) but does not require a name supply argument to build the result term. This function `ren` hence both converts to de Bruijn indices and applies the supplied function on free variables. To build the function `ren`, the user-supplied function on names must be adapted to cross name-abstractions. We reuse the function `protect↑1` from section 5.3.4 which is adapted when both the source and target abstractions are de Bruijn indices. We build a variant of the function `protect↑1` called `protect<↑1` which works when the source abstraction is nominal and the target abstraction is de Bruijn.

```

protect<↑1 : ∀ {α β b} → (Name α → Name β)
              → (Name (b < α) → Name (β ↑1))
protect<↑1 f = exportWith (0N) (sucN↑ ∘ f)
-- protect<↑1 f b = 0
-- protect<↑1 f x = 1 + f x

ren : ∀ {α β} → (Name α → Name β) → Tm α → Tm β
ren f (λN b t) = λD (ren (protect<↑1 f) t)
ren f (V x)     = V (f x)
ren f (λD t)    = λD (ren (protect↑1 f) t)
ren f (t · u)   = ren f t · ren f u

```

Using this renaming function we build a simple normalization function which just supplies the identity function as the renaming:

```

norm : ∀ {α} → Tm α → Tm α
norm = ren id

```

Using this combined approach one may lose the convenience of nominal bindings. To recover this convenience, we build a function called `openTm` such as in section 6.1. The function `openTm` takes a name supply `s` and an open de Bruijn term `t`. The function renames the de Bruijn index zero to the first name of the name supply. We then define a simple function called `nomview` which converts the toplevel λ -abstraction from de Bruijn to nominal. This function `nomview` enables to look at the terms as though there were only nominal λ -abstractions. However using `nomview` and `openTm` incurs the limitations we discussed in section 6.1, namely convenience, performance, and non-structural recursion.

```

openTm : ∀ {α} (s : Supply α) → Tm (α ↑1) → Tm (seedB s < α)
openTm {α} (b , b#) = ren f where
  f : Name (α ↑1) → Name (b < α)
  f = exportWith (nameB _) (coerceN (⊆-# b#) ∘ predN)

nomview : ∀ {α} → Supply α → Tm α → Tm α
nomview s (λD t) = λN _ (openTm s t)
nomview _ t      = t

```

A tantalizing change would be to include new constructors to represent the open and close operations internally. This would enable a lazy opening and closing for terms. For instance we could try to add these data constructors:

$$\begin{aligned} \text{Close} & : \forall \{\alpha\} \mathbf{b} \rightarrow \text{Tm} (\mathbf{b} \triangleleft \alpha) \rightarrow \text{Tm} (\alpha \uparrow 1) \\ \text{Open} & : \forall \{\alpha \mathbf{b}\} \rightarrow \mathbf{b} \# \alpha \rightarrow \text{Tm} (\alpha \uparrow 1) \rightarrow \text{Tm} (\mathbf{b} \triangleleft \alpha) \end{aligned}$$

However, we did not go very far with these constructors. Pattern-matching on these is already limited given the indices of the result type are not invertible. Using explicit type equalities helps to go a bit further but not much:

$$\begin{aligned} \text{Close} & : \forall \{\alpha \beta\} (\text{eq} : \beta \equiv \alpha \uparrow 1) \mathbf{b} \rightarrow \text{Tm} (\mathbf{b} \triangleleft \alpha) \rightarrow \text{Tm} \beta \\ \text{Open} & : \forall \{\alpha \beta \mathbf{b}\} (\text{eq} : \beta \equiv \mathbf{b} \triangleleft \alpha) \rightarrow \mathbf{b} \# \alpha \rightarrow \text{Tm} (\alpha \uparrow 1) \rightarrow \text{Tm} \beta \end{aligned}$$

We have explored one particular combination, namely combining nominal and de Bruijn indices styles. We have shown how to build terms, but also how to get a canonical form using de Bruijn indices. We have also shown how to convert one de Bruijn binding to a nominal one. Exploring other combinations is left as a topic for further research. This may include combinations with de Bruijn levels, higher-order representations (such as our normalization by evaluation function in section 3.4), nominal binders with fresh-for proofs to maintain distinct binders, etc.

6.5 Nominal types

In this section we successively describe various specification languages to describe languages. More concretely a specification language enables to write type definitions and their binding structure. In section 6.5.1, we describe nominal signatures as used in FRESHML-LITE [Shinwell et al., 2003]. We first describe signatures with a single sort of atoms and then generalize to multiple sorts. We then focus on languages with *pattern* types such as FRESHML and FRESH OCAML in section 6.5.1, Caml in section 6.5.2 and BINDERS UNBOUND in section 6.5.3.

6.5.1 FreshML & Fresh OCaml types

FRESHML [Shinwell et al., 2003] is a programming language, designed by Mark R. Shinwell, Andrew M. Pitts and Murdoch J. Gabbay. The language is designed to make programming with binders simple. The language is an ML-like language augmented with types and language constructs to deal with names and binders. The language has a conventional call-by-value operational semantics and a type system preventing to break name abstraction. This means that two α -equivalent terms cannot be distinguished within the language. The language offers an adequate encoding of nominal terms (to be described shortly) and guarantees that the identity of a bound name cannot be observed. However, it does not prevent a newly generated atom from escaping its scope.

FRESH OCAML [Shinwell and Pitts, 2005] is an extension of the OCAML [Leroy et al., 2005] language. This language extension enriches OCAML with a primitive type for names and one type construction for name abstractions. From a language design point of view, FRESH OCAML brings the features of FRESHML to OCAML.

FRESHML-LITE is a light version of FRESHML without pattern types and restricted to a single sort of atoms. Indeed FRESHML and FRESH OCAML have a simple notion of *pattern types* that we do not treat here. Pattern types are covered in section 6.5.2 and 6.5.3 about Caml and BINDERS UNBOUND. About multiple sorts we show how to handle those and then go back to a single sort to simplify matters.

Nominal terms are the values described by *nominal signatures*. Here is the definition of a nominal signature as of Definition 2.1 of “Nominal Unification” [Urban et al., 2004]:

“A nominal signature is specified by: a set of sorts of atoms (typical symbol ν); a disjoint set of sorts of data (typical symbol δ); and a set of function symbols (typical symbol \mathbf{f}), each of which has an arity of the form $\tau \rightarrow \delta$. Here τ ranges over (compound) sorts given by the grammar $\tau ::= \nu \mid \delta \mid 1 \mid \tau \times \tau \mid \langle \nu \rangle \tau$. Sorts of the form $\langle \nu \rangle \tau$ classify terms that are binding abstractions of atoms of sort ν over terms of sort τ . We will explain the syntax and properties of such terms in a moment.”

Our goal now is to show how we can describe any nominal signature using our library. We could expect sorts of atoms, sorts of data and compound sorts to be translated directly as types in AGDA, and thus to have type `Set` (no index). However this would be incompatible with our technique which indices names (atoms) and terms by worlds. We call *expression types* what corresponds to the compound sorts in nominal signatures. The term *expression* takes all its sense when combined with *pattern types* that we will see later. In AGDA we call the type of expression types `E`.

Nominal signatures on one atom sort We now describe the interface to build and combine expression types and nominal signatures. For simplicity we assume that only one sort of atom is used, that we call ν . The first component is called ν^e and corresponds to atoms of sort ν . The second component is called 1^e and corresponds to the unit expression type. We can take the product of two expression types with $_ \times^e _$. Finally the expression type $\langle \nu \rangle^e \tau$ classifies the terms that are name-abstractions of atoms of sort ν over terms of type expression τ . Note that since we constrained the interface to a single sort of atoms, name-abstraction expects only one argument, the expression type. To declare a function symbol we introduce $_ \rightarrow^e _$ which takes two expression types and returns an AGDA type (i.e. `Set`). In AGDA we have the following declarations for the interface:

```

module MonoSortedNominalSignature where
  E      : Set1
  νe    : E
  1e    : E
  _×e_  : E → E → E
  <ν>e_  : E → E
  ↪e_   : E → E → Set

  -- Non-essential but used in a later example
  Liste : E → E

```

We now give an example of the use of this interface to build a nominal signature. We use AGDA’s data type definitions to handle the introduction of data sorts, function symbols, and the handling of recursion. This example describes a fragment of ML. We recall as a comment how this nominal signature was originally written [Urban et al., 2004].

```

{- A Nominal Signature:
   Example 2.2 from ‘Nominal Unification’
-}
sort of atoms: vid
sort of data:  exp
function symbols:
  vr  : vid → exp
  app : exp × exp → exp
  lam : <vid>exp → exp
  lv  : exp × <vid>exp → exp
  lf  : <vid>((<vid>exp) × exp) → exp
-}

```

```

module NominalSignatureExample where
  open MonoSortedNominalSignature

  data Exp : E where
    vr  : νe ↪e Exp
    app : Exp ×e Exp ↪e Exp
    lam : <ν>e Exp ↪e Exp
    lv  : Exp ×e <ν>e Exp ↪e Exp
    lf  : <ν>e((<ν>e Exp) ×e Exp) ↪e Exp

```

As we can see, the definition of a nominal signature can be easily done within our system. The resulting `Exp` type is fully usable. As an example we give the definition of the function `fv`.

```

fv : Exp →e Liste νe
fv (vr x)                = [ x ]
fv (app (t , u))        = fv t ++ fv u
fv (lam (b , t))        = rm b (fv t)
fv (lv (t , b , u))     = fv t ++ rm b (fv u)
fv (lf (bf , (b , t) , u)) = rm bf (rm b (fv t) ++ fv u)

```

This function `fv` already reveals parts of the definitions for our nominal signature combinators that we now present in detail. We have not much other options than using `Name` as the definition for ν^e . This forces \mathbb{E} to be `World → Set`. This last definition can also be read as “expression types are the world-indexed types” which should come as no surprise as it captures the very intuition behind our system. The unit type expression (1^e) ignores the world index and yields AGDA’s unit type. The product of expression types ($-\times^e-$) distributes the world to the subcomponents. Name abstraction pairs a binder and a sub-expression type whose world is properly extended. Finally, the combinator to declare a function symbol ($\lrcorner\rightarrow^e-$) distributes to the underlying arrow type but also quantifies over the world. Here are the definitions in AGDA.

```

module MonoSortedNominalSignature where
   $\mathbb{E}$  : Set1
   $\mathbb{E}$  = World → Set

  νe :  $\mathbb{E}$ 
  νe = Name

  1e :  $\mathbb{E}$ 
  1e _ =  $\top$  --  $\top$  is Agda's unit type

   $-\times^e-$  :  $\mathbb{E} \rightarrow \mathbb{E} \rightarrow \mathbb{E}$ 
  ( $E_1 \times^e E_2$ ) α =  $E_1$  α ×e  $E_2$  α

   $\langle \nu \rangle^e$  :  $\mathbb{E} \rightarrow \mathbb{E}$ 
  ( $\langle \nu \rangle^e E$ ) α = ∃ [ b ] (E (b < α))

   $\lrcorner\rightarrow^e$  :  $\mathbb{E} \rightarrow \mathbb{E} \rightarrow \text{Set}$ 
   $E_1 \lrcorner\rightarrow^e E_2 = \forall \{ \alpha \} \rightarrow E_1 \alpha \rightarrow E_2 \alpha$ 

```

One might want to unfold the type of a few symbols we defined. Modulo currying, the symbols `Exp`, `vr`, `lam` and `fv` correspond to `Tm`, `V`, `λ` and `fv` from section 2.2:

```

Exp : World → Set
vr  : ∀ {α} → Name α → Exp α
lam : ∀ {α} → ∃[ b ] Exp (b < α) → Exp α
fv  : ∀ {α} → Exp α → List (Name α)

```

Multi-sorted nominal signatures We now enrich our expression types with multiple sorts of atoms. We call `Sort` the type of atom sorts. This type `Sort` is actually part of the definition of a nominal signature and thus is not part of the interface. Actually, the interface is parameterized over the type for atom sorts. We rename the ν^e as `Namee` in this version. `Namee` associates an \mathbb{E} to each value of `Sort`. The unit, list and product for expression types are now derived from the more general combinators `Neutrale`, `Neutral1e`, `Neutral2e`, etc. These combinators respectively lift any type `Set`, `Set → Set`, `Set → Set → Set` to \mathbb{E} , $\mathbb{E} \rightarrow \mathbb{E}$, $\mathbb{E} \rightarrow \mathbb{E} \rightarrow \mathbb{E}$, etc. The expression type $\langle \nu \rangle^e \tau$ classifies the terms that are name-abstractions of atoms of sort ν over terms of type expression τ . We keep the other combinators unchanged. In AGDA we have the following declarations for the interface:

```

module MultiSortedNominalSignature
  (Sort : Set) ( _==_ : (x y : Sort) → Bool ) where

  E      : Sort → Set1
  Namee : Sort → E
  <_>e  : Sort → E → E
  ↪e   : E → E → Set
  Neutrale : Set → E
  Neutral1e : (Set → Set) → (E → E)
  Neutral2e : (Set → Set → Set) → (E → E → E)

  1e : E
  1e = Neutrale ⊤

  Liste : E → E
  Liste = Neutral1e List

  _×e_ : E → E → E
  _×e_ = Neutral2e _×_

```

To illustrate this new interface and the use of multiple sorts we present a simple signature for System F where type variables are distinguished from term variables. Notice that apart from the boring definition for `_==_` our

embedded language for signatures is lightweight.

```

module NominalSignatureExample where
  data Sort : Set where
    vtm vty : Sort

  _==_ : (x y : Sort) → Bool
  vtm == vtm = true
  vty == vty = true
  vtm == vty = false
  vty == vtm = false

  open MultiSortedNominalSignature nomPa Sort _==_

  data Ty :  $\mathbb{E}$  where
    var : Namee vty  $\mapsto^e$  Ty
    arr : Ty  $\times^e$  Ty  $\mapsto^e$  Ty
    all : < vty >e Ty  $\mapsto^e$  Ty

  data Tm :  $\mathbb{E}$  where
    vr  : Namee vtm  $\mapsto^e$  Tm
    app : Tm  $\times^e$  Tm  $\mapsto^e$  Tm
    lam : Ty  $\times^e$  < vtm >e Tm  $\mapsto^e$  Tm
    App : Tm  $\times^e$  Ty  $\mapsto^e$  Tm
    Lam : < vty >e Tm  $\mapsto^e$  Tm

```

We show that computing the free type variables in both types and terms can be programmed without any extra difficulty. Here the type system not only prevents from leaking bound names but also prevents from mixing the two sorts of variables.

```

fvtyTy : Ty  $\mapsto^e$  Liste (Namee vty)
fvtyTy (var x) = [ x ]
fvtyTy (arr ( $\sigma$  ,  $\tau$ )) = fvtyTy  $\sigma$  ++ fvtyTy  $\tau$ 
fvtyTy (all (b ,  $\tau$ )) = rm b (fvtyTy  $\tau$ )

```

```

fvtyTm : Tm  $\mapsto^e$  Liste (Namee vty)
fvtyTm (vr x) = []
  -- [ x ] would not type-check
fvtyTm (app (t , u)) = fvtyTm t ++ fvtyTm u
fvtyTm (lam (τ , b , t)) = fvtyTy τ ++ fvtyTm t
  -- rm b ... would not type-check
fvtyTm (App (t , τ)) = fvtyTm t ++ fvtyTy τ
fvtyTm (Lam (b , t)) = rm b (fvtyTm t)

```

We can now turn to the implementation of the multi-sorted interface. Expression types are no longer indexed by a single world but by a collection of worlds. This collection of worlds is itself indexed by atom sorts, and represented by a function from sorts to worlds. The combinator `Namee` is the type `Name` in the world given by looking up the world associated to the sort ν . The combinator `<_>e_` needs to update the collection of worlds at the given atom sort. To do so we define the function `_[::=_]` to do this update. The combinators `Neutrale`, `Neutral1e`, and `Neutral2e` are distributing the index to the subcomponents. Since $\lambda A \rightarrow \text{Ix} \rightarrow A$ is an applicative functor (often called Reader) for all index types `Ix`, the neutral combinators correspond to `pure`, `liftA`, and `liftA2` of this applicative functor. Here is the AGDA implementation.

```

module MultiSortedNominalSignature
  (Sort : Set) ( _==_ : (x y : Sort) → Bool ) where

  E : Set1
  E = (Sort → World) → Set

  Namee : Sort → E
  Namee ν Γ = Name (Γ ν)

  _[::=_] : (Sort → World) → Sort → World → (Sort → World)
  (Γ [ ν := α ]) ν' = if ν == ν' then α else Γ ν'

  <_>e_ : Sort → E → E
  (< ν >e E) Γ = ∃[ b ](E (Γ [ ν := b < Γ ν ]))

```

```

Neutrale : Set →  $\mathbb{E}$ 
Neutrale A _ = A

Neutral1e : (Set → Set) → ( $\mathbb{E}$  →  $\mathbb{E}$ )
Neutral1e F E  $\Gamma$  = F (E  $\Gamma$ )

Neutral2e : (Set → Set → Set) → ( $\mathbb{E}$  →  $\mathbb{E}$  →  $\mathbb{E}$ )
Neutral2e Op E1 E2  $\Gamma$  = Op (E1  $\Gamma$ ) (E2  $\Gamma$ )

1e :  $\mathbb{E}$ 
1e = Neutrale  $\top$ 

Liste :  $\mathbb{E}$  →  $\mathbb{E}$ 
Liste = Neutral1e List

_×e_ :  $\mathbb{E}$  →  $\mathbb{E}$  →  $\mathbb{E}$ 
_×e_ = Neutral2e _×_

```

FreshML types In addition to nominal signatures, FRESHML and FRESH OCAML support function types in their signatures. This can be taken into account by adding an arrow combinator. We introduce it in two flavors:

```

_→e_ :  $\mathbb{E}$  →  $\mathbb{E}$  →  $\mathbb{E}$ 
(E1 →e E2)  $\alpha$  = E1  $\alpha$  → E2  $\alpha$ 

_⇒e_ :  $\mathbb{E}$  →  $\mathbb{E}$  →  $\mathbb{E}$ 
(E1 ⇒e E2)  $\alpha$  =  $\forall \{\beta\} \rightarrow \alpha \subseteq \beta \rightarrow (E_1 \rightarrow^e E_2) \beta$ 

⇒-to-→ :  $\forall \{E_1 E_2\} \rightarrow (E_1 \Rightarrow^e E_2) \mapsto^e (E_1 \rightarrow^e E_2)$ 
⇒-to-→ f = f  $\subseteq$ -refl

```

We introduced the latter when defining a normalization by evaluation algorithm in section 3.4. Moreover the latter version is generally more useful than the former since it can be “imported” by merging the given inclusion witness. Here is how to coerce such a function:

```

Coe :  $\mathbb{E}$  → Set
Coe E =  $\forall \{\alpha \beta\} \rightarrow \alpha \subseteq \beta \rightarrow E \alpha \rightarrow E \beta$ 

coerce-⇒e :  $\forall \{E_1 E_2\} \rightarrow Coe (E_1 \Rightarrow^e E_2)$ 
coerce-⇒e pf f = f  $\circ$   $\subseteq$ -trans pf

```

6.5.2 Caml types

Caml [Pottier, 2006] is both a tool and a language to describe data type with a binding structure. Caml has been developed by François Pottier. The tool takes data type definitions in a file and produces an OCAML [Leroy et al., 2005] module with OCAML data types and generated boiler plate code.

The data types language of Caml extends ML algebraic data types that we find in OCAML with annotations to describe the scoping rules. More precisely, Caml types extends nominal signatures with a notion of *pattern types*.

Caml data types are of two kinds: expression and pattern types. The expression kind is made to be used for, well, expression-like types. For instance our type `Tm` belongs to this kind, but type-expressions, statements, processes, etc. are expression types. We reuse the terminology of the previous section and call \mathbb{E} the type of expression types. As before we reuse the following primitive expression types: `Namee`, `_→e_`, `Neutrale`, `Neutral1e`, `Neutral2e`. We reuse also the following derived expression types: `1e`, `Liste`, `_×e_`. In order for Caml to derive functions about these data types, functions are not allowed to appear in signatures. Name-abstractions are made of a *pattern type* and surrounded by `<->`. Pattern types are called \mathbb{P} and also support various combinators to be described shortly.

```
-- A pattern type
P   : Set1

-- Name abstraction
<-> : P → E
```

Pattern types appear when we start a name-abstraction. Like expression types, pattern types are closed under products, and have the arrow `_→P_` to declare a data constructor (function symbol). In addition, patterns can be binders, or expression types with *inner* and *outer* annotations. Since pattern types embed expression types, base types can be lifted to patterns as well. Here are the signatures of the various forms of pattern types:


```

module CamlTypes where
  open MonoSortedNominalSignature public

  BinderP  :  $\mathbb{P}$ 
   $\_ \times^{\mathbb{P}} \_$  :  $\mathbb{P} \rightarrow \mathbb{P} \rightarrow \mathbb{P}$ 
   $\_ \mapsto^{\mathbb{P}} \_$  :  $\mathbb{P} \rightarrow \mathbb{P} \rightarrow \text{Set}$ 

  OuterP  :  $\mathbb{E} \rightarrow \mathbb{P}$ 
  InnerP  :  $\mathbb{E} \rightarrow \mathbb{P}$ 

  NeutralP :  $\text{Set} \rightarrow \mathbb{P}$ 
  NeutralP = OuterP  $\circ$  Neutrale
             -- which is identical to
             --   InnerP  $\circ$  NeutralP

  -- List of patterns as special data type
  -- since we could not reuse List here.
  data ListP (P :  $\mathbb{P}$ ) :  $\mathbb{P}$  where
    [] :  $1^{\mathbb{P}} \mapsto^{\mathbb{P}} \text{List}^{\mathbb{P}} \text{P}$ 
    :: :  $\text{P} \times^{\mathbb{P}} \text{List}^{\mathbb{P}} \text{P} \mapsto^{\mathbb{P}} \text{List}^{\mathbb{P}} \text{P}$ 

```

To appear in pattern types, expression types are qualified as either *outer* or *inner*, and thus are respectively in-scope and out-of-scope of the surrounding abstraction. The disposition of binders and embedded expressions does not matter for Caml. The *outer* expressions are not affected by any binder of the pattern. Similarly all the *inner* expressions of a pattern are put in scope of all the binders introduced by the pattern.

Example To illustrate the use of expression and pattern types, we define our term type in the Caml style and call it `Exp`. The `lam` case is made of a name-abstraction holding a pattern product with one binder and one inner term. The `let` case is more interesting and makes use of the `OuterP` combinator to hold the first subterm out of the scope of the binder. As with nominal signatures we are using a `data` declaration to take care of introducing the constructors and dealing with the recursive nature of the type:

```

data Exp :  $\mathbb{E}$  where
  var   :  $\text{Name}^e \mapsto^e \text{Exp}$ 
  app   :  $\text{Exp} \times^e \text{Exp} \mapsto^e \text{Exp}$ 
  lam   :  $\langle \text{Binder}^{\mathbb{P}} \times^{\mathbb{P}} \text{Inner}^{\mathbb{P}} \text{Exp} \rangle \mapsto^e \text{Exp}$ 
  let'  :  $\langle \text{Binder}^{\mathbb{P}} \times^{\mathbb{P}} \text{Outer}^{\mathbb{P}} \text{Exp} \times^{\mathbb{P}} \text{Inner}^{\mathbb{P}} \text{Exp} \rangle \mapsto^e \text{Exp}$ 

```

We continue our illustration of `Caml`'s expressiveness with local blocks of local definitions. To achieve this we simply reuse the declaration for `let'` and wrap the binding with `ListP`:

```
let★ : < ListP (BinderP ×P OuterP Exp) ×P InnerP Exp > ↦e Exp
```

To go from non-recursive blocks to recursive ones, a single step is required. Replacing `OuterP` by the `InnerP` annotation suffices to make all the bound names accessible to the right hand part of this construct.

```
rec : < ListP (BinderP ×P InnerP Exp) ×P InnerP Exp > ↦e Exp
```

We finally add some pattern-matching facilities to our small language. To do so, we introduce a pattern type named `Pat`. This type supports wildcard patterns, single binders, and pairs of patterns:

```
data Pat : P where
  wildcard : 1P ↦P Pat
  binder   : BinderP ↦P Pat
  pair     : Pat ×P Pat ↦P Pat
```

Once we have patterns, we package them as branches. A branch holds a pattern, an optional guard, and a term associated with the pattern:

```
Branch : E
Branch = < Pat ×P InnerP (Maybee Exp) ×P InnerP Exp >
```

The type `Exp` is then extended with a new `match` construct, which holds a scrutiny term and a list of branches:

```
match : Exp ×e Liste Branch ↦e Exp
```

In our AGDA development we show how to modularly define a free-variables function. To do so we define a combinator to compute the free variables for each form of expression types and pattern types. For pattern types we need no less than three actions to collect the inner free variables, the outer free variables, and to remove pattern binders from a list of free variables. Then at name-abstractions, free variables is the union of outer free variables and inner ones minus those bound by the pattern. In the end, one can build an `fv` function on terms whose body is made of combinators and exactly follow the definition of terms.

Behind the scenes We now turn to the definitions needed to make these combinators work with our library. We focus on a single sort of atoms since we have shown how we can support multiple sorts. Expression types are not changed and directly reused from section 6.5.1.

The C_{aml} change to expression types is the name-abstraction. However, before discussing the name-abstraction combinator, we need to cover pattern types. Pattern types are indexed by no less than two worlds, and an operator on worlds. Let \mathbb{P} be a pattern type, let α and β be worlds, let Op be an operator on worlds ($\text{World} \rightarrow \text{World}$), then $\mathbb{P} \alpha \beta \text{Op}$ is a fully applied pattern type of type Set . The world α plays the role of the outer world, meaning before binders introduced by the pattern. The world β is the inner world, meaning after all the binders introduced by the pattern. Thus, the worlds α and β are outer and inner worlds with respect to the complete pattern found between name-abstraction brackets ($\langle _ \rangle$). These two worlds α and β are respectively used by $\text{Outer}^{\mathbb{P}}$ and $\text{Inner}^{\mathbb{P}}$ to pick the world of the expression type. Then the operator explains what the given pattern type changes to a world. If the pattern does not bind anything (as is the case of $\text{Outer}^{\mathbb{P}}$ and $\text{Inner}^{\mathbb{P}}$), then the operator is equal to the identity. Since a binder pattern binds one binder, the operator is equal to the corresponding binder insertion: $\mathbf{b} \triangleleft \dots$. Pattern products distribute α and β to subcomponents. Subcomponents world operation get composed into a single operation to index the product. Finally the abstraction brackets package an operation and a fully applied pattern type whose outer world is the current world, the inner world is computed by applying Op to the current world. Here are the definitions in AGDA:

```

 $\mathbb{P} : \text{Set}_1$ 
 $\mathbb{P} = (\alpha \beta : \text{World}) (\text{Op} : \text{World} \rightarrow \text{World}) \rightarrow \text{Set}$ 

 $\text{Binder}^{\mathbb{P}} : \mathbb{P}$ 
 $\text{Binder}^{\mathbb{P}} \_ \_ \text{Op} = \exists [ \mathbf{b} ] (\text{Op} \equiv \_ \triangleleft \_ \mathbf{b})$ 

 $\text{Outer}^{\mathbb{P}} : \mathbb{E} \rightarrow \mathbb{P}$ 
 $\text{Outer}^{\mathbb{P}} \mathbb{E} \alpha \_ \text{Op} = \mathbb{E} \alpha \times \text{Op} \equiv \text{id}$ 

 $\text{Inner}^{\mathbb{P}} : \mathbb{E} \rightarrow \mathbb{P}$ 
 $\text{Inner}^{\mathbb{P}} \mathbb{E} \_ \beta \text{Op} = \mathbb{E} \beta \times \text{Op} \equiv \text{id}$ 

data  $\_ \times^{\mathbb{P}} \_ (P_1 P_2 : \mathbb{P}) : \mathbb{P}$  where
   $\_ \_ : \forall \{ \text{Op}_1 \text{Op}_2 \alpha \beta \}$ 
     $\rightarrow P_1 \alpha \beta \text{Op}_1$ 
     $\rightarrow P_2 \alpha \beta \text{Op}_2$ 
     $\rightarrow (P_1 \times^{\mathbb{P}} P_2) \alpha \beta (\text{Op}_2 \circ \text{Op}_1)$ 

```

$$\langle _ \rangle : \mathbb{P} \rightarrow \mathbb{E}$$

$$\langle P \rangle \alpha = \exists [Op] (P \alpha (Op \alpha) Op)$$

$$\mapsto^P _ : \mathbb{P} \rightarrow \mathbb{P} \rightarrow \text{Set}$$

$$\mapsto^P P_1 P_2 = \forall \{ \alpha \beta Op \} \rightarrow P_1 \alpha \beta Op \rightarrow P_2 \alpha \beta Op$$

In the end, the world operation index of a pattern type could be replaced by a list of binders. However this abstract presentation enables more flexible uses. For instance here is the de Bruijn version of the combinator `BinderP`:

$$\text{Binder}^{\text{Dp}} : \mathbb{P}$$

$$\text{Binder}^{\text{Dp}} _ _ Op = Op \equiv _ \uparrow 1$$

This de Bruijn combinator is useful in two ways. First one can imagine mixing `BinderP` and `BinderDp` as we have done in section 6.4. A second option would be to easily change the implementation from nominal to de Bruijn by just updating a single definition. This also shows that both nominal and de Bruijn are expressive enough to cover Caml data types.

A remark on linearity We can remark a subtle difference between Caml’s semantics for pattern types and the semantics provided by our definitions. In Caml if two binders in a same pattern are equal there is no shadowing. Such a pattern could be for instance: $\lambda (x, x) \rightarrow \dots$. In our setting there is shadowing and thus the order of binders in a pattern is relevant. This order is well-defined and specified on the combinator for products. In a pattern product, all the binders on the left comes before all the binders on the right. Thus those on the right can shadow those on the left. We actually recommend to deal with linearity before going to a well-scoped representation. If non-linear patterns have to be accepted, we recommend to turn the subsequent occurrences into free names and thus keep a linear pattern.

Conclusion In this section we explained how our embedded language to describe nominal signatures can scale to Caml signatures. This shows that all types that can be defined in Caml can also be defined in our system. However Caml also provides generated code to perform usual operations on data structures with names and binders. We do not address this part here but several generic programming techniques could be applied in our situation as well.

6.5.3 Connection with Binders Unbound

Stephanie Weirich, Brent Yorgey, and Tim Sheard develop a library called BINDERS UNBOUND [Weirich et al., 2011]. This library is written in HASKELL and provides type level combinators to express data structures with binders.

Then through a mechanism of “data generic programming”, usual functions are generically provided for any combination of types. The main feature of BINDERS UNBOUND is to get all of these operations for free right after the type definition. It enables rapid prototyping where we directly focus on the code that matters to our program. The safety features are similar to those in FRESHML, FRESH OCAML, and Caml. Name abstraction cannot be violated but names can escape their scope. However by providing many generic operations the number of bugs the programmer can introduce is reduced. Of course nothing prevents from combining both generic programming and a safer approach were names do not escape their scope. To achieve their goal of generic programming, an embedded language of combinators to define types is developed.

The language of combinators is modeled after expression and pattern types. Given our definitions of section 6.5.1 and section 6.5.2, we can adapt to the new combinators concisely. We focus on a single sort of atoms since we have shown how we can support multiple sorts. The first difference is the combinator `Bind P E` which combines a pattern type `P` and an expression type `E`. The binders in `P` scope over the expression type `E`. This is equivalent to `< P ×P InnerP E >` in Caml, but the combinator `<_>` is not part of BINDERS UNBOUND. The combinator `Embed` enables to embed an expression type into a pattern type. The combinator `Embed` is equivalent to `OuterP` from Caml. The real new combinators are `Rebind` and `Rec`. `Rebind P1 P2` is a pattern type which binds both `P1` and `P2` like `P1 ×P P2` does. However, binders of `P1` also scope over the expressions embedded in `P2`. `Rec P` is a pattern types where binders of `P` not only scope externally but scope over expressions embedded in `P`. As in Caml, functions cannot be used in data types.

We give definitions for `Bind`, `Embed`, `Rebind` and `Rec` using our definitions for Caml. The second world index (β) of pattern types is not used in BINDERS UNBOUND but would enable to combine them with Caml combinators.

```
Bind : P → E → E
Bind P E α = ∃[ Op ](P α (Op α) Op × E (Op α))
           -- ≅ < P ×P InnerP P E > α
```

```
Embed : E → P
Embed = OuterP
```

```

data Rebind (P1 P2 :  $\mathbb{P}$ ) :  $\mathbb{P}$  where
  -, - :  $\forall \{Op_1 Op_2 \alpha \beta\}$ 
    → P1  $\alpha$   $\beta$  Op1
    → P2 (Op1  $\alpha$ )  $\beta$  Op2
    → Rebind P1 P2  $\alpha$   $\beta$  (Op2  $\circ$  Op1)

```

```

Rec :  $\mathbb{P} \rightarrow \mathbb{P}$ 
Rec P  $\alpha$   $\beta$  Op = P (Op  $\alpha$ )  $\beta$  Op

```

Open/close based interface We have shown how the types provided by the library BINDERS UNBOUND can be expressed within our system. What about the operations? Their library generically provides the usual operations such as computing free variables, comparing two terms up to α -equivalence, substitutions and performing a renaming. However in order to let the user program the remaining specific operations, the library had to commit to a particular implementation or interface to program with binders. They chose to provide an interface with a default implementation using Locally Nameless [Charguéraud, 2011, Aydemir et al., 2008]. We discussed at length in section 6.1 the limitations of an interface based on open and close (here called unbind and bind respectively). In particular when the type of abstractions is kept abstract the convenience is reduced when pattern matching, performance is hurt due to unwanted traversals, and structural recursion is no longer possible.

Telescopes Here is an example of the usefulness of the combinator `Rebind`, namely telescopes. Telescopes are widely used in dependently typed programming languages. A telescope is a sequence of typed bindings. Let Γ be a meta-variable for such a telescope defined as $(A : \text{Set}) (x : A)$. Telescopes can be used to represent an environment in a typing judgement as in $\Gamma \vdash x : A$. They can be used to represent a sequence of arguments types as in $\Gamma \rightarrow A$. They can also be used to represent a sequence of typed arguments as in $\lambda \Gamma \rightarrow x$. While all these construction could be made by nesting `Bind`, this would tie the telescope to a particular usage. If we want the telescope to be a value on its own, its needs to be a pattern type. Hence the use of `Rebind` instead of `Bind`.

To illustrate `Rebind` and telescopes here are two mutually defined data types. The type `Exp` represents expressions in a small dependently typed language where types and expressions are fused in the same syntactic class. Expressions have variables, dependent function types, λ -abstractions, function applications, and a sort called `set`. The type `Tele` represents nested sequences of bindings associated with an expression, namely telescopes:

```

mutual
data Exp :  $\mathbb{E}$  where
  V    : Namee  $\mapsto^e$  Exp
   $\Pi$   : Bind Tele Exp  $\mapsto^e$  Exp
   $\lambda$  : Bind Tele Exp  $\mapsto^e$  Exp
   $\_ \cdot \_$  : Exp  $\times^e$  Liste Exp  $\mapsto^e$  Exp
  set  : 1e  $\mapsto^e$  Exp

data Tele :  $\mathbb{P}$  where
  []   : 1p  $\mapsto^p$  Tele
   $::$   : Rebind (Binderp  $\times^p$  Embed Exp) Tele  $\mapsto^p$  Tele

```

We show how to compute the free variables for expressions, lists of expressions, and telescopes. To do so we have to lift the function `rm` to telescopes as well:

```

mutual
fv : Exp  $\mapsto^e$  Liste Namee
fv (V x) = [ x ]
fv ( $\Pi$  (bind  $\Gamma$  t)) = fvTele  $\Gamma$  ++ rmTele  $\Gamma$  (fv t)
fv ( $\lambda$  (bind  $\Gamma$  t)) = fvTele  $\Gamma$  ++ rmTele  $\Gamma$  (fv t)
fv ( $\_ \cdot \_$  (t , us)) = fv t ++ fvL us
fv (set _) = []

fvL : Liste Exp  $\mapsto^e$  Liste Namee
fvL [] = []
fvL (t :: ts) = fv t ++ fvL ts

fvTele :  $\forall$  { $\alpha$   $\beta$  Op}  $\rightarrow$  Tele  $\alpha$   $\beta$  Op  $\rightarrow$  List (Name  $\alpha$ )
fvTele ( [] _) = []
fvTele ( $::$  ((binder b , embed  $\tau$ ) ,  $\Gamma$ )) = fv  $\tau$  ++ rm b (fvTele  $\Gamma$ )

rmTele :  $\forall$  { $\alpha$   $\beta$  Op}  $\rightarrow$  Tele  $\alpha$   $\beta$  Op
       $\rightarrow$  List (Name (Op  $\alpha$ ))  $\rightarrow$  List (Name  $\alpha$ )
rmTele ( [] (neutral _) ) = id
rmTele ( $::$  ((binder b , embed _) ,  $\Gamma$ )) = rm b  $\circ$  rmTele  $\Gamma$ 

```

Conclusion In this section on nominal types we have shown the expressiveness of our system to define various families of data types. Through a lightweight embedded language we can support nominal signatures with multiple sorts and pattern types as in FRESHML, Caml and BINDERS UNBOUND. Hence worlds are at least as expressive as these languages. Using our

embedded language to define data types is kept very light since the resulting data types are not much different from those written directly. However the expressiveness of worlds is also a weakness. Indeed generic programming as done in BINDERS UNBOUND is greatly simplified by having a language of combinators instead of our flexible use of worlds. To this end we advocate to have both flexible worlds with a mechanized notion of α -equivalence and combinators to get generic programming. We have not studied enough how to integrate data type generic programming since there are plenty of ways to build universes in AGDA and meta-programming tools are promising but still too young.

6.6 Related work

The difficulty of programming with, or reasoning about, names, binders, and α -equivalence has been known for a long time. It has recently received a lot of attention, due in part to the POPLMARK challenge [Aydemir et al., 2005]. Despite this attention, the problem is still largely unsolved: according to Guillemette and Monnier, for instance, “none of the existing representations of bindings is suitable” [Guillemette and Monnier, 2008].

In the following, we review several systems that are intended to facilitate the manipulation of names and binders. This review is not exhaustive: we focus on relatively recent related work.

In particular we have already covered nominal systems such as FRESHML [Shinwell et al., 2003], FRESH OCAML [Shinwell and Pitts, 2005], Caml [Pottier, 2006]; de Bruijn indices [de Bruijn, 1972]; systems based on well-scoped de Bruijn indices [Altenkirch, 1993, McBride and McKinna, 2004, Bellegarde and Hook, 1994, Bird and Paterson, 1999, Altenkirch and Reus, 1999]; dynamically stratified representations such as Locally Nameless [Charguéraud, 2011, Aydemir et al., 2008] and Locally Named [Pollack et al., 2011, Sato and Pollack, 2010]; and BINDERS UNBOUND [Weirich et al., 2011] which is internally based on Locally Nameless.

FreshML and Pure FreshML Pure FRESHML [Pottier, 2007] is built on top of FRESHML. The semantics of FRESHML dictates that pattern matching against a name abstraction silently replaces the bound atom with a fresh atom. FRESHML guarantees that “name abstractions cannot be broken”. However, FRESHML is unsafe: it is possible for a name to escape its scope. Put another way, FRESHML is impure: name generation is an observable side effect.

Pure FRESHML imposes additional static proof obligations, which ensure that freshly created atoms do not escape their scope, and correspond to Pitts’ *freshness condition for binders* [2006]. Because these proof obligations are expressed in a specialized logic, they can be discharged automatically.

Because it is safe, Pure FRESHML can be implemented either using atoms (such as the original FRESHML) or using de Bruijn indices. This is an implementation choice, which the programmer need not know about.

In contrast with Pure FRESHML, the approach proposed in the present work does require certain runtime checks: the operation `exportTm` fails if its argument contains a free name that cannot be exported. We claim that often those checks are required and thus would be discharged by Pure FRESHML only with equivalent checks. However in some cases, some invariants of the code would make these checks redundant. Pure FRESHML might be able to show that we can export without any check.

In Pure FRESHML, name abstraction is a primitive notion, and the fact that deconstructing an abstraction automatically freshens the bound atom is used to guarantee that all terms effectively live in a single world. In our work, in contrast, name abstraction is explained in terms of more basic notions, and it is possible to deconstruct a name abstraction without substituting a fresh name for the bound name. This leads to a finer-grained understanding of binding, and, in some cases, to greater runtime efficiency.

Nominal System T Pitts’ Nominal System T [2010] follows the tradition of FRESHML and guarantees that name abstractions cannot be violated. In order to ensure that names do not escape their scope, Pitts uses a dynamic technique: the ν construct, which can be applied to any term, turns every name occurrence that is about to escape its scope into a harmless “anonymous name”, known as `anon` (or `new`). One potential advantage of this approach is that the application of ν to a term can be lazily evaluated, whereas our `exportTm` operation requires eagerly traversing the entire term.

We can emulate this behavior in our system by adding a special value to the type for names to represent `anon`. We experimented with this option in our development [Pouillard, 2011a] by defining a type `Name? α` as `Maybe (Name α)` and `anon` as `nothing`. Most operations can be easily lifted. In particular `exportN?` gets the type `Name? ($\mathbf{b} \triangleleft \alpha$) \rightarrow Name? α` and so becomes “total”. Then, exporting terms gets the type `Tm ($\mathbf{b} \triangleleft \alpha$) \rightarrow Tm α` . However there is no direct equivalent for `¬Name \emptyset` , since the empty world is indeed inhabited by `anon`. Instead we have that only `anon` inhabits the empty world.

One may wonder if the use of the `Maybe` type could be avoided? Indeed we tried another option in our development [Pouillard, 2011a] where the type `Name? α` is equal to `Name ($\alpha \uparrow 1$)` and `anon` is equal to `zeroN`. We also lifted most of the operations from `Name` to `Name?`. We also gain “total” functions to export names and terms which would allow for a lazy exporting at the price of delaying “errors”.

A separate “data layer” Some systems are said to have a “data layer” separated from the “computation layer” for types. The data layer is generally a LF signature of a Logical Framework [Harper et al., 1993] (LF for short). This generally means that data types are only defined in the data layer. In the data layer there is an arrow type for functions but it should not be confused with computational functions that we expect from a programming language. A type constructor from the data layer can then be injected in the computation layer. The injection from the data to computation layer is a place where these systems differ. The data and computation layers being separated, this implies that we cannot mix first-class computational functions and data types.

In the case of a simply typed LF, we can map the data layer signatures to nominal signatures (section 6.5.1). To do so, we build an injection called $[-]$ from LF types to atom sorts. Usually this means an atom sort per data sort. Then the translation called $\llbracket _ \rrbracket$ of the signature is straightforward. The arrow case $\llbracket S \rightarrow T \rrbracket$ is translated to $\langle [S] \rangle^e \llbracket T \rrbracket$. Finally for each data type δ , a data constructor is added for variables. This constructor has type: $V\delta : \text{Name}^e [\delta] \mapsto^e \delta$. This translation can be used as an intuition of what it means to lift a type from the data layer to the computation layer and also as an initial step to support LF signatures in our system. In particular Licata and Harper follow a similar translation but using a single constructor for variables shared among the complete signature. Here is an example:

```
{- LF signature for System F:
Ty  : type.
arr : Ty -> Ty -> Ty.
all : (Ty -> Ty) -> Ty.

Tm  : type.
app : Tm -> Tm -> Tm.
lam : Ty -> (Tm -> Tm) -> Tm.
App  : Tm -> Ty -> Tm.
Lam  : (Ty -> Tm) -> Tm.
-}
```

Only two type expressions appear on the left of an arrow. The injection $[-]$ is thus defined as follow:

```
[ Ty ] = vty
[ Tm ] = vtm
```

Applying the previously mentioned translation we obtain the following data types:

```

data Ty :  $\mathbb{E}$  where
  VTy : Namee vty  $\mapsto^e$  Ty
  arr : Ty  $\times^e$  Ty  $\mapsto^e$  Ty
  all :  $\langle$  vty  $\rangle^e$  Ty  $\mapsto^e$  Ty

data Tm :  $\mathbb{E}$  where
  VTm : Namee vtm  $\mapsto^e$  Tm
  app : Tm  $\times^e$  Tm  $\mapsto^e$  Tm
  lam : Ty  $\times^e$   $\langle$  vtm  $\rangle^e$  Tm  $\mapsto^e$  Tm
  App : Tm  $\times^e$  Ty  $\mapsto^e$  Tm
  Lam :  $\langle$  vty  $\rangle^e$  Tm  $\mapsto^e$  Tm

```

Elphin, Delphin, Beluga, Beluga ^{μ} Elphin [Schürmann et al., 2005], Delphin [Poswolsky and Schürmann, 2008, 2009], and Beluga [Pientka, 2008] are closely related to one another in several ways. They aim at building programming languages where data types are built from LF types and terms.

Elphin only has a limited way to include types from the data layer in the computation layer. Both Delphin and Beluga have contexts but their handling is different. Delphin deals with a “current context” while Beluga has explicit contexts. Delphin allows incremental changes of the current context using ν whereas Beluga combines contexts and LF objects at the frontier of data and computation.

Beluga ^{μ} [Cave and Pientka, 2012] extends Beluga with richer computational types. In Beluga ^{μ} we can define data types in the computation layer as well. These data types can be recursive and indexed by LF objects of the data layer. While substitution comes for free at the data layer, substitution has to be user defined on data types from the computation layer. We think more connections between our system and Beluga ^{μ} are to be found. We also notice their use of our system to prove properties of their language in COQ proof assistant. They used a simplified version of our 2010 paper [Pouillard and Pottier, 2010] which corresponds to the simplifications we have made in section 6.3.

At the data level, Elphin, Delphin, Beluga provide substitution and higher-order matching as primitive operations. This ambitious approach can eliminate some boilerplate code, at the cost of a complex meta-theory. By contrast, the meta-theory of our proposal is very simple, as it only extends an existing logical relations argument with a few new primitive types and operations. We argue that some of the boilerplate can still be eliminated via generic programming as this is done by Licata and Harper but also in BINDERS UNBOUND.

Licata and Harper’s system Licata and Harper [2009] aim to provide substitution for free when possible; and they expose the use of well-scoped de Bruijn indices to the programmer, whereas we offer a choice between nominal and de Bruijn-based representation techniques.

This said, there are numerous similarities between the two systems. Both keep track of the context, or world, within which each name makes sense. Both offer flexible ways of parameterizing or quantifying types over worlds. Both offer ways of moving data from one world to another: Licata and Harper’s weakening and strengthening respectively correspond to our import (coerce) and export operations. Both systems support first-class computational functions. Not all functions can be imported or exported, but some can: for instance, in both systems, the example of normalization by evaluation (section 3.4), which requires importing a function into a larger world, is made type-correct by planning ahead and making this function polymorphic with respect to an arbitrary world extension.

Distinctions One traditionally distinguishes several broad approaches to the problem of names and binders, which employ seemingly different tools, namely: atoms and atom abstractions; well-scoped de Bruijn indices; higher-order abstract syntax. We believe that this distinction can be superficial. In fact, our work presents strong connections with all three schools of thought. Perhaps more important are the following questions:

Can α -equivalent terms be distinguished? Except for bare de Bruijn indices and plain nominal terms all the systems we have presented prevent α -equivalent terms to be distinguished. These systems offer an adequate encoding of “nominal” terms and guarantees that the identity of a bound name cannot be observed.

What hygiene properties are enforced by the system? The base hygiene property is prevent from α -equivalent terms to be distinguished. We now focus on more hygiene properties. In FRESHML, FRESH OCAML, Caml, Locally Nameless, Locally Named, and BINDERS UNBOUND generating names is an observable effect. Said differently functions can turn α -equivalent inputs into not α -equivalent outputs. Nominal System T repairs this problem by dynamically detecting an atom that attempts to escape and by turning it into a harmless “anonymous atom”. Systems based on well-scoped de Bruijn indices enforce the invariant that every name is within range, that is, every name refers to some binding site. However, this alone does not imply that indices are correctly adjusted where needed, or that comparisons between names are allowed only when they make sense. In systems based on higher-order abstract syntax, and in Pure FRESHML, name manipulation is hygienic by design: this is built in the syntax and semantics

of the programming language.

In the present document, hygiene is built in at a very low level. We provide a small number of abstract types, such as `Binder` and `Name`, together with a small number of operations. These operations are restricted on purpose: for instance, comparing two binders for equality is disallowed; comparing two names for equality is permitted only if they inhabit a common world. Through logical relations and parametricity, we are able to explore the consequences of these restrictions and to find out which hygiene properties can be expected of a well-typed program. In particular, we guarantee that two α -equivalent terms cannot be distinguished and that functions turn α -equivalent inputs into α -equivalent outputs. These results hold because *our* definition of α -equivalence *is* the logical relation, and that every well-typed program inhabits this relation. Moreover, our definition of α -equivalence corresponds to the “standard” notion of α -equivalence, at least for all nominal signatures like the type `Tm`.

Is there a type for names that can be freely used in types? The answer to this question is positive for most “first order” systems we have studied: de Bruijn indices, well-scoped de Bruijn indices, Locally Nameless, Locally Named, FRESHML, FRESH OCAML, `Caml`, Pure FRESHML, Nominal System *T*, BINDERS UNBOUND, and in the present work.

In systems based on LF and higher-order abstract syntax the type for names is not freely used. Instead, names inhabit all types issued from the data layer.

Does the system have different data and computation layers? Elphin, Delphin, Beluga, Beluga^μ clearly have two different layers. Since `Caml` does not support function while OCAML does we conclude that `Caml` has two different layers as well. In de Bruijn indices, well-scoped de Bruijn indices, Locally Nameless, Locally Named, FRESHML, FRESH OCAML, Pure FRESHML, Nominal System *T*, BINDERS UNBOUND and the present work there is no separation. However very few systems handle the mix of computational functions and data types with binders. Systems based on a generator or generic programming such as most Locally Nameless and Locally Named implementations, `Caml`, and BINDERS UNBOUND disallow types with computational functions. In Pure FRESHML, there is no support for first class functions at all. In de Bruijn indices, well-scoped de Bruijn indices, FRESHML and FRESH OCAML data and computation can be mixed since little operations are generically provided. However in FRESHML and FRESH OCAML the swap operation *is* generically provided and supports functions. In Licata and Harper’s work [2009] data and computation can be mixed. However they are both part of a “universe” which is then interpreted into computation types of the host language.

Is there a set of types where substitution can be defined? If so what are the restrictions? Some systems let the user program substitution functions. Some systems try to provide them for free either at all types or have restrictions. When a system has a specific “data layer” for types, the goal is to have substitution for free on these types once injected into the computation layer.

In Elphin, Delphin, and Beluga, substitution for data types comes for free but the data layer is restricted and cannot embed computational functions.

In Beluga^μ there is substitution has to be defined for data types in the computation layer and comes for free on the data layer.

In FRESHML, FRESH OCAML, Nominal System *T*, Pure FRESHML this is up to the user to program substitution functions. In C_αml (systems based on nominal signatures), Locally Nameless, Locally Named programmers have to define substitution functions but generators may be provided for usual cases. In BINDERS UNBOUND, substitution is generically provided for a family of types not including functions. In Licata and Harper’s work their is a generic definition for the substitution which is conditioned by criterion on the type.

In our system, substitution has to be programmed by the user and no particular set of types is prescribed for this task. Then as we have shown earlier a wide range of signatures can be embedded in our systems and generic programming could be applied on those.

How does the system keep track of the context or world in which a name makes sense? In Pure FRESHML and in Nominal System *T*, there is effectively just one world, within which every name makes sense; static or dynamic checks guarantee that no confusion can arise. In Elphin, Delphin, or in Licata and Harper’s system, the meaning of types is relative to a “current context”, and a number of modalities are provided to discard the current context, extend it with one new name, etc. In Beluga and Beluga^μ, contexts are explicit: a data layer type, once annotated with a context, becomes a computation layer type. In the present work, worlds are explicit, and are built into algebraic data type definitions by the programmer. There is no notion of a “current world”: multiple worlds can co-exist, and mechanisms are offered for transporting names (or whole data structures) from one world to another.

None of de Bruijn indices, Locally Nameless, Locally Named, FRESHML, FRESH OCAML, C_αml, or BINDERS UNBOUND keep track of the context in which a name makes sense.

Which high-level operations are involved in the semantics of the programming language? The semantics of the “nominal” systems (such as FRESHML, FRESH OCAML, C_αml, Pure FRESHML, Nominal System *T*)

involves renaming. In *Locally Nameless*, *Locally Named*, and *BINDERS UNBOUND* the open/close operations are costly and necessary to use to go under binders. The semantics of *Elphin*, *Delphin*, *Beluga*, and *Beluga^μ* involve higher-order matching. In the present work, as well as in Licata and Harper’s work, no costly operations are built into the semantics; high-level operations, such as our import and export operations, are programmed explicitly or obtained via generic programming.

Moving across representations It is arguably desirable to be able to offer several choices of representation within a single system, and to be able to migrate from one representation to another. Our system supports multiple representation styles, and combinations thereof (section 6.4). Furthermore, our implementation of normalization by evaluation (section 3.4) illustrates how to move back and forth between “syntactic” name abstractions and “semantic” name abstractions in the style of higher-order abstract syntax.

Robert Atkey and co-authors [Atkey, 2009, Atkey et al., 2009] investigate how to move back and forth between higher-order abstract syntax and de Bruijn indices. The translation out of higher-order abstract syntax produces well-scoped de Bruijn indices, but the proof of this fact is meta-theoretic. Atkey uses Kripke logical relations to argue that the current world at the time of application of a certain function must be larger than the world at the time of construction of this function. This seems somewhat related with our use of bounded polymorphism in the definition of semantic name abstractions (section 3.4). An exact connection remains to be investigated.

Chapter 7

Conclusion

7.1 Contributions

Our contributions are practical and fundamental. On the practical side, we have built a library (in AGDA) for programming with names and binders. The library comprises a safe programming interface and a sound implementation, which has been mechanically proved against a strong specification. The interface supports different styles for binders: nominal, de Bruijn indices, de Bruijn levels, and combinations of them.

Our interface is not only safe but expressive as well. Explicit worlds enable both multiple kinds of variables and precise control over the scoping rules. Name abstraction is *not* a primitive type nor an abstract type but concretely built out of more atomic types. While expressive and safe, the system is economical on the number of introduced types.

On a fundamental side, we re-use the logical relations defined for AGDA in [Bernardy et al., 2010], and build a logical relation argument which provides a strong specification of α -equivalence by building a unified model. This is the very same model that is used for all styles: nominal, de Bruijn indices, de Bruijn levels and the combinations. The framework of logical relations provides a systematic way of defining α -equivalence on complex data types. Moreover free-theorems arise for world-polymorphic functions defined by the user, yielding interesting properties about user code for free.

This work also establishes numerous connections between various systems of the three schools of thought (nominal, de Bruijn, HOAS). We built our library with a “nominal” insight and lightly integrated de Bruijn indices based on well-scoped de Bruijn indices approach. We then built a normalization by evaluation algorithm which shows how “syntactic” and “semantic” abstraction can be combined.

This unified approach does not stop here. De Bruijn levels, Locally Nameless, Locally Named, and combinations of them can be safely built on top of our abstract interface and programmed precisely. Through the

use of generic traversals and traversal kits, we can share most of the usual definitions we have on terms.

Moreover our types are very flexible, as shown by the embedding of nominal signatures, `Cam1` types, and `BINDERS UNBOUND` types. Not only flexible, our types automatically come with a definition for α -equivalence which is guaranteed to be preserved by the well-typed programs using our interface.

7.2 Future work

As future work we would like to make our library lighter to use. To do so we would like to infer most of the inclusion witnesses needed to move from a world to another.

While hand-written generic traversals already avoid code duplication, we would like to generically define this traversal for a fixed universe of types.

Studying how this system could support the presence of side-effects would be an interesting direction as well. This could be done either by considering a language design with our features and side effects or by studying how our types compose with particular effects as monads.

Finally we would like to explore how convenient this technique is, to not only program but reason about programs. Since the library is based on abstract types, this would require exporting properties about the functions of the interface. We can notice that this direction is promising since a simplified version of our approach has been successfully used by Cave and Pientka to prove properties of `Beluga`^u.

We would also like to better study the connection with systems based on LF and higher-order abstract syntax. A direction could be to combine the universe from Licata and Harper but using our library instead of well-scoped de Bruijn indices.

We also need to deal with a proper extraction of our library. We want to erase: worlds, membership proofs, world inclusion witnesses, etc. A real implementation should rely on machine integers; and useless traversal functions (like `coerceTm`) should be optimized away.

We focused on names and binders to guarantee that the binding structure of well-scoped terms is preserved by well-typed programs. However safe and expressive meta-programming requires that not only scopes but types are preserved. We worked on initial ideas in this direction but we left this to future work.

Bibliography

- Thorsten Altenkirch. A formalization of the strong normalization proof for System F in LEGO. In J.F. Groote M. Bezem, editor, *Typed Lambda Calculi and Applications*, LNCS 664, pages 13 – 28, 1993. URL <http://www.cs.nott.ac.uk/~txa/publ/tlca93.pdf>.
- Thorsten Altenkirch and Bernhard Reus. Monadic presentations of lambda terms using generalized inductive types. In *Computer Science Logic*, volume 1683 of *Lecture Notes in Computer Science*, pages 453–468. Springer, 1999. URL <http://www.cs.nott.ac.uk/~txa/publ/cs199.pdf>.
- Robert Atkey. Syntax for free: representing syntax with binding using parametricity. In *International Conference on Typed Lambda Calculi and Applications (TLCA)*, volume 5608 of *Lecture Notes in Computer Science*, pages 35–49. Springer, July 2009.
- Robert Atkey, Sam Lindley, and Jeremy Yallop. Unembedding domain-specific languages. In *Haskell symposium*, pages 37–48, September 2009. URL <http://personal.cis.strath.ac.uk/~raa/unembedding.pdf>.
- Brian Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. Engineering formal metatheory. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 3–15, January 2008. URL <http://www.cis.upenn.edu/~bcpierce/papers/binders.pdf>.
- Brian E. Aydemir, Aaron Bohannon, Matthew Fairbairn, J. Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich, and Steve Zdancewic. Mechanized metatheory for the masses: The POPLMARK challenge. In *International Conference on Theorem Proving in Higher-Order Logics (TPHOLs)*, volume 3603 of *Lecture Notes in Computer Science*, pages 50–65. Springer, August 2005. URL <http://research.microsoft.com/en-us/people/dimitris/poplmark.pdf>.
- Françoise Bellegarde and James Hook. Substitution: A formal methods case study using monads and transformations. *Sci. Comput. Program.*, 23(2-3): 287–311, 1994. doi: [http://dx.doi.org/10.1016/0167-6423\(94\)00022-0](http://dx.doi.org/10.1016/0167-6423(94)00022-0).

- Jean-Philippe Bernardy, Patrik Jansson, and Ross Paterson. Parametricity and dependent types. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional programming*, ICFP '10, pages 345–356, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-794-3. doi: <http://doi.acm.org/10.1145/1863543.1863592>. URL <http://doi.acm.org/10.1145/1863543.1863592>.
- Richard Bird and Ross Paterson. de Bruijn notation as a nested datatype. *Journal of Functional Programming*, 9(1):77–91, January 1999. URL <http://dx.doi.org/10.1017/S0956796899003366>.
- Andrew Cave and Brigitte Pientka. Programming with binders and indexed data-types. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 413–424, January 2012. URL <http://www.cs.mcgill.ca/~bpientka/papers/recursive-types.pdf>.
- Arthur Charguéraud. The locally nameless representation. *Journal of Automated Reasoning*, pages 1–46, 2011. ISSN 0168-7433. URL <http://www.chargueraud.org/research/2009/ln/main.pdf>. 10.1007/s10817-011-9225-2.
- Nicolaas G. de Bruijn. Lambda-calculus notation with nameless dummies: a tool for automatic formula manipulation with application to the Church-Rosser theorem. *Indag. Math.*, 34(5):381–392, 1972.
- Dawson R. Engler, Wilson C. Hsieh, and M. Frans Kaashoek. ‘c’: a language for high-level, efficient, and machine-independent dynamic code generation. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, POPL '96, pages 131–144, New York, NY, USA, 1996. ACM. ISBN 0-89791-769-3. doi: <http://doi.acm.org/10.1145/237721.237765>. URL <http://doi.acm.org/10.1145/237721.237765>.
- Louis-Julien Guillemette and Stefan Monnier. A type-preserving compiler in Haskell. In *ACM International Conference on Functional Programming (ICFP)*, pages 75–86, September 2008. URL <http://www-etud.iro.umontreal.ca/~guillelj/icfp08.pdf>.
- Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, pages 194–204, 1993.
- Graham Hutton. *Programming in Haskell*. Cambridge University Press, January 2007.
- Xavier Leroy, Damien Doligez, Jacques Garrigue, Didier Rmy, and Jérôme Vouillon. *The Objective Caml system*, October 2005. URL <http://caml.inria.fr/>.

- Daniel R. Licata and Robert Harper. A universe of binding and computation. In *ACM International Conference on Functional Programming (ICFP)*, pages 123–134, September 2009. URL <http://www.cs.cmu.edu/~drl/pubs/lh09unibind/lh09unibind.pdf>.
- Conor McBride and James McKinna. The view from the left. *J. Funct. Program.*, 14:69–111, January 2004. ISSN 0956-7968. doi: 10.1017/S0956796803004829. URL <http://www.cs.ru.nl/~james/RESEARCH/view-final2004.pdf>.
- Conor McBride and Ross Paterson. Applicative programming with effects. *Journal of Functional Programming*, 18(1):1–13, 2008.
- Ulf Norell. *Towards a practical programming language based on dependent type theory*. PhD thesis, Department of Computer Science and Engineering, Chalmers University of Technology, SE-412 96 Göteborg, Sweden, September 2007.
- Brigitte Pientka. A type-theoretic foundation for programming with higher-order abstract syntax and first-class substitutions. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 371–382, January 2008. URL <http://www.cs.mcgill.ca/~bpientka/papers/hoasfun-short.pdf>.
- Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- Benjamin C. Pierce, editor. *Advanced Topics in Types and Programming Languages*. MIT Press, 2005.
- Andrew M. Pitts. Alpha-structural recursion and induction. *Journal of the ACM*, 53:459–506, 2006. URL <http://www.cl.cam.ac.uk/~amp12/papers/alpsri/alpsri.pdf>.
- Andrew M. Pitts. Nominal System *T*. In *ACM Symposium on Principles of Programming Languages (POPL)*, pages 159–170, January 2010. URL <http://www.cl.cam.ac.uk/~amp12/papers/nomst/nomst-popl.pdf>.
- Randy Pollack, Masahiko Sato, and Wilmer Ricciotti. A canonical locally named representation of binding. pages 1–23, May 2011. URL <http://homepages.inf.ed.ac.uk/rpollack/export/PollackSatoRicciottiJAR.pdf>.
- Adam Poswolsky and Carsten Schürmann. Practical programming with higher-order encodings and dependent types. In *European Symposium on Programming (ESOP)*, volume 4960 of *Lecture Notes in Computer Science*, pages 93–107. Springer, March 2008. URL <http://cs-www.cs.yale.edu/homes/poswolsk/papers/delphinESOP08.pdf>.

- Adam Poswolsky and Carsten Schürmann. System description: Delphin – A functional programming language for deductive systems. *Electronic Notes in Theoretical Computer Science*, 228:113–120, 2009. URL <http://www.itu.dk/~carsten/papers/lfmtp-08.pdf>.
- François Pottier. An overview of Caml. In *ACM Workshop on ML*, volume 148 of *Electronic Notes in Theoretical Computer Science*, pages 27–52, March 2006. URL <http://gallium.inria.fr/~fpottier/publis/fpottier-alphacaml.pdf>.
- François Pottier. Static name control for FreshML. In *IEEE Symposium on Logic in Computer Science (LICS)*, pages 356–365, July 2007. URL <http://gallium.inria.fr/~fpottier/publis/fpottier-pure-freshml.pdf>.
- Nicolas Pouillard. NompA (Agda code), 2011a. <http://tiny.nicolaspouillard.fr/NomPa.agda>.
- Nicolas Pouillard. Nameless, painless. In *Proceedings of the 16th ACM SIGPLAN international conference on Functional programming, ICFP '11*, pages 320–332. ACM, September 2011b. URL <http://nicolaspouillard.fr/publis/nameless-painless.pdf>.
- Nicolas Pouillard and François Pottier. A fresh look at programming with names and binders. In *Proceedings of the 15th ACM SIGPLAN international conference on Functional programming, ICFP '10*, pages 217–228, New York, NY, USA, September 2010. ACM. ISBN 978-1-60558-794-3. doi: <http://doi.acm.org/10.1145/1863543.1863575>. URL <http://gallium.inria.fr/~fpottier/publis/pouillard-pottier-fresh-look.pdf>.
- John C. Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing 83*, pages 513–523. Elsevier Science, 1983. URL <ftp://ftp.cs.cmu.edu/user/jcr/typesabpara.pdf>.
- Masahiko Sato and Randy Pollack. External and internal syntax of the lambda-calculus. *J. Symb. Comput.*, 45(5):598–616, 2010. URL <http://homepages.inf.ed.ac.uk/rpollack/export/SatoPollack09.pdf>.
- Carsten Schürmann, Adam Poswolsky, and Jeffrey Sarnat. The ∇ -calculus: Functional programming with higher-order encodings. In *International Conference on Typed Lambda Calculi and Applications (TLCA)*, volume 3461 of *Lecture Notes in Computer Science*, pages 339–353. Springer, April 2005. URL <http://www.itu.dk/~carsten/papers/nabla.pdf>.
- Mark R. Shinwell and Andrew M. Pitts. Fresh Objective Caml user manual. Technical Report 621, University of Cambridge, February 2005. URL <http://www.cl.cam.ac.uk/TechReports/UCAM-CL-TR-621.pdf>.

- Mark R. Shinwell, Andrew M. Pitts, and Murdoch J. Gabbay. FreshML: Programming with binders made simple. In *ACM International Conference on Functional Programming (ICFP)*, pages 263–274, August 2003. URL <http://www.cl.cam.ac.uk/~amp12/papers/frepbm/frepbm.pdf>.
- Walid Taha. *Multi-stage Programming: Its Theory and Applications*. PhD thesis, Oregon Graduate Institute, November 1999. URL <http://www.cs.rice.edu/~taha/publications/thesis/thesis.pdf>.
- Christian Urban, Andrew Pitts, and Murdoch Gabbay. Nominal unification. *Theoretical Computer Science*, 323:473–497, 2004. URL <http://citeseerx.ist.psu.edu/viewdoc/download;jsessionid=D14EB8679102796E01FBD20257BB5F34?doi=10.1.1.3.7275&rep=rep1&type=pdf>.
- Philip Wadler. Views: a way for pattern matching to cohabit with data abstraction. In *ACM Symposium on Principles of Programming Languages (POPL)*, POPL '87, pages 307–313, New York, NY, USA, 1987. ACM. ISBN 0-89791-215-2. doi: <http://doi.acm.org/10.1145/41625.41653>. URL <http://doi.acm.org/10.1145/41625.41653>.
- Philip Wadler. Theorems for free! In *Conference on Functional Programming Languages and Computer Architecture (FPCA)*, pages 347–359, September 1989. URL <http://homepages.inf.ed.ac.uk/wadler/papers/free/free.ps.gz>.
- Stephanie Weirich, Brent Yorgey, and Tim Sheard. Binders unbound. In *Proceedings of the 16th ACM SIGPLAN international conference on Functional programming*, ICFP '11, pages 333–345. ACM, September 2011. URL <http://www.cis.upenn.edu/~byorgey/papers/binders-unbound.pdf>.

List of Figures

2.1	The NOMPA interface	49
3.1	Lifting <code>trName</code>	66
4.1	Logical relations for core types	83
4.2	Logical relations for data types	84
4.3	Relations for NOMPA types	90
5.1	Core operations on names	99
5.2	Operations on names	99
5.3	New rules for world inclusion	100
5.4	Shifting versus adding example	110
5.5	Various term types: from concrete to abstract	111
6.1	Various forms of name abstraction	127

Index

- α -equivalent, 26
- Abstract Syntax Trees, 20
- algebraic data types, 30
- assuming, 27
- AST, 20

- binder, 26
- binding-links, 140
- bound, 26

- Code
 - β -red^D, 108
 - β -red, 70, 147
 - \perp , 37
 - \circ -Kit, 68
 - \emptyset , 55, 96
 - \equiv , 38
 - $[-]$, 169
 - \mathbb{N} , 30
 - \mathbb{P} , 159
 - \neg Name \emptyset , 53, 55, 96
 - \neg A, 38
 - ν^e , 152, 154
 - \subseteq - \emptyset +1, 100
 - \subseteq - \emptyset , 54, 96
 - \subseteq -Reasoning, 55
 - \subseteq -#, 54
 - \subseteq - \triangleleft , 54
 - \subseteq -dist-+- \triangleleft , 137
 - \subseteq -dist-+1- \triangleleft , 137
 - \subseteq -dist- \triangleleft -+1, 137
 - \subseteq -dist- \triangleleft -+, 137
 - \subseteq -refl, 54, 96
 - \subseteq -trans, 54, 55, 96
 - T, 37
 - 1^e, 152, 154
 - AbsPreTm, 126, 128
 - AbsTm, 121
 - Add^L.kit, 137
 - Atom, 41, 120
 - Binder^{DP}, 163
 - Binder^P, 163
 - Binder, 50, 96
 - Bind, 164
 - Bool, 30
 - B, 85
 - Coerce, 67
 - Comm, 145
 - DataEnv, 61
 - DeBruijn, 142
 - Exp, 153, 154, 160, 165
 - Fin, 95, 100
 - FreshLink, 143
 - Fresh, 143
 - Inner^P, 162
 - Let, 42, 140
 - Link, 140
 - List, 32
 - Name^e, 157
 - Name, 50, 96
 - Neutral1^e, 155, 157
 - Neutral2^e, 155, 157
 - Neutral^e, 155, 157
 - NomFresh, 143
 - Nom, 142
 - Outer^P, 160, 162
 - Pat, 73, 161
 - PreTm, 126
 - Preserve \equiv , 88
 - Program, 35
 - Rebind, 164
 - Rec, 164

- RelatedFunctions, 81
- RenameEnv^L, 135
- SemAbs, 74, 75
- Sem, 74
- Set₀, 37
- Set₁, 37
- Sort, 155
- SubstEnv+ \subseteq , 128
- SubstEnv^L, 135
- SubstEnv, 65, 66
- Supply, 65
- SynAbs^D, 126
- SynAbs^N, 74, 128
- Tele, 165
- Tm^A, 41, 121
- Tm^D, 97
- Tm^L⇒Tm^N, 138
- Tm^L, 133
- Tm^M, 94
- Tm, 50, 140
- TraversePreTm, 127
- Tree, 34
- V, 41, 94
- World, 96
- , 41, 133
- ÷-, 36
- ≡₀-, 38
- \subseteq -, 53, 96
- ×^e-, 152, 154
- ↑₁, 97, 98
- ↑-, 98
- *- , 31
- ++-, 36
- +₁, 55, 97
- +-, 31
- +^W-, 98
- , ↪-, 75
- /2, 31
- #∅, 54
- #-, 54, 79
- $\overset{\circ}{=}$ ₀-, 39
- (<-because-), 53
- [[→]-, 81
- [[\subseteq]-, 89, 92
- [[↑₁]-, 109
- [[+₁]-, 109
- [[<]-, 89, 91
- [[V]-, 86
- [[==^N]-, 91
- ↪^e-, 152, 154
- ↪^P-, 159
- ÷-, 36
- ⊗-, 40, 67
- <-, 50, 78, 96
- V-, 85
- |↪|-, 67
- >>-, 36
- ^N, 101
- .: -, 32
- [.:=-], 157
- ==-, 155
- ==^A-, 120
- ==^N-, 63, 79, 96, 108
- ==Tm-, 88
- $\overset{\circ}{=}$, 39
- (<-)^N, 115
- (<-)Tm, 115
- (<-)^W, 115
- [[II]], 81
- [[∅]], 89
- [[¬Name∅]], 91
- [[B]], 86
- [[Binder]], 88
- [[Name]], 88
- [[Set₀]], 82
- [[Tm^D]]↔rename, 117
- [[Tm^D]], 114
- [[_]], 81
- [[add^N]], 110
- [[cmp^N]], 110
- [[false]], 86
- [[name^P]], 91
- [[true]], 86
- ℰ, 154
- , 55
- |Cmp|, 62
- <ν>^e-, 152
- <->^e-, 157

<->, 159, 164
 addTm^D, 107
 addTm^L, 138
 add^N, 98
 add^s, 101
 appTm^D, 104
 appTm^M, 95
 ba, 44, 47
 binder^N, 52
 closeTm?, 70
 closeTm^D?, 107
 closeTm^L, 136
 closeTm, 121, 127, 129
 cmp-ba, 44, 47
 cmpName[↑], 108
 cmpTm^D, 109
 cmpTm^L, 133
 cmpTm, 63, 88
 cmp^N, 100
 coeTm, 64
 coerceEnv, 76
 coerceKit, 64
 coerceTm^L, 138
 coerceTm, 70
 coerce^N, 53, 55, 79, 96
 commEnv, 145
 compTm^D, 104
 convTm^N∅ ⇒ Tm^L∅, 139
 convTm^N∅ ⇒ Tm^L, 139
 conv, 71
 env-length-Tm^J, 48
 eval, 75
 exportWith, 60
 export^N?, 57, 79
 export^N, 56
 extEnv, 63, 71
 extendNameCmp, 62
 extend, 60
 falseTm, 52
 false, 30, 85
 fast-fv, 48
 fa, 125
 filter, 33
 fv^D, 104
 fv, 43, 44, 47, 61, 133, 140, 148,
 153, 154
 idTm^D, 104
 idTm, 52
 id, 32
 importFun, 145
 init, 143
 is42?, 85
 kitVar, 146
 lam, 154, 160
 length, 33
 let', 122, 161
 lookup, 61
 mapKit, 68
 map, 33
 match, 161
 nameOf, 140
 name^B, 52, 79
 next, 143
 node, 34
 nomview, 150
 not, 30
 occurs, 60
 occ, 60
 openSubstTm, 122, 127, 128
 openTm, 121, 150
 plug, 73
 pred^N?, 102
 print, 36
 protect[↑]A, 106
 protect[↑], 103
 protect[↵]1, 149
 protectedAdd[↑], 113
 protectedAdd, 112
 pure, 40, 67
 reify, 76
 renameAKit^L, 136
 renameAKit, 67
 renameKit+_⊆, 128
 renameKit, 66
 renameTm?, 70
 renameTm^D?, 106
 renameTm^DA, 106
 renameTm^D, 106

- renameTm^L⇒Tm^N, 139
- renameTm^L?, 136
- renameTm^LA, 136
- renameTm^L, 136
- renameTm, 146
- ren, 149
- rerootTm^L, 138
- rm₀, 104
- rm^A, 43
- rm, 59, 140
- seed, 65
- set, 165
- shiftName', 103
- shiftName, 102
- shiftTm^D, 107
- show, 36
- size, 51, 123, 124
- starKit, 68
- strengthen?, 140
- substKit, 67
- substPreTm, 128
- substTm^D, 108
- substTm^L, 136
- substTm, 70, 147
- substVarTm^D, 108
- subtractTm^D?, 107
- subtractTm^D, 107
- subtractTm^L, 138
- subtract^N?, 102
- subtract^N, 100
- subtract^s, 101
- suc#, 54, 79
- suc^B, 51, 96
- suc^N↑, 101
- suc^N, 101
- suc^s, 65
- suc, 30
- sumForest, 34
- sumTree, 34
- t^f, 43
- t^x, 43
- t^y, 43
- trAtom, 71
- trBinder, 63
- trName, 63, 65, 145
- trTm', 69
- trTm^{D'}, 105
- trTm, 68
- triple, 31
- trueTm, 53, 55, 56
- true, 30, 85
- unprotectedAdd, 113
- vr, 154
- weakenTm, 147
- weaken, 140
- zero^B, 51, 96
- zero^N, 98
- zero^s, 65
- zero, 30
- [], 32
- κ', 122
- λ, 41, 94, 95
- complete, 17
- context, 26
- data constructors, 30
- definitional equality, 38
- definitionally equal, 38
- equal pointwise, 39
- expression types, 152
- free, 26
- fresh, 26
- GADT, 30
- has type, 18
- higher-order, 33
- inductive families, 29
- infix, 16
- inner, 159, 160
- Keyword
 - ..., 34
 - Set, 37
 - with, 33
- lexing, 19

links, [140](#)
locally closed, [120](#)
Locally Named, [120](#)
Locally Nameless, [120](#), [126](#)

mutually recursive, [34](#)

name, [26](#)
nominal, [27](#)
nominal signatures, [152](#)
Nominal terms, [152](#)

occurrence, [26](#)
outer, [159](#), [160](#)

parsing, [20](#)
pattern, [151](#)
pattern type, [159](#)
pattern types, [152](#)
pattern matching, [30](#)
pattern types, [152](#), [159](#)
pointwise equality, [40](#)
polymorphism, [32](#)
propositional equality, [38](#)

recursive, [30](#)
regular tree types, [29](#)

sequenced, [14](#)
sound, [17](#)
sum and product types, [29](#)

tail, [33](#)
token class, [20](#)
tokens, [20](#)
traversal kit, [63](#)
types, [17](#)
typing judgements, [18](#)

variable, [26](#)
wildcard, [33](#)