



**POLARITIES & FOCUSSING:
A JOURNEY FROM
REALISABILITY TO AUTOMATED REASONING**

STÉPHANE GRAHAM-LENGRAND

**DISSERTATION SUBMITTED TOWARDS THE DEGREE OF
HABILITATION À DIRIGER DES RECHERCHES
UNIVERSITÉ PARIS-SUD**

**Thesis prepared at École Polytechnique, with support from CNRS and INRIA,
and publicly defended on 17th December 2014 before**

LAURENT REGNIER	Chairman
WOLFGANG AHRENDT	Referee
HUGO HERBELIN	Referee
FRANK PFENNING	Referee
SYLVAIN CONCHON	Examiner
DAVID DELAHAYE	Examiner
DIDIER GALMICHE	Examiner
CHRISTINE PAULIN-MOHRING	Examiner

Abstract

This dissertation explores the roles of polarities and focussing in various aspects of Computational Logic.

These concepts play a key role in the the interpretation of proofs as programs, a.k.a. the Curry-Howard correspondence, in the context of classical logic. Arising from linear logic, they allow the construction of meaningful semantics for cut-elimination in classical logic, some of which relate to the Call-by-Name and Call-by-Value disciplines of functional programming. The first part of this dissertation provides an introduction to these interpretations, highlighting the roles of polarities and focussing. For instance: proofs of positive formulae provide structured data, while proofs of negative formulae consume such data; focussing allows the description of the interaction between the two kinds of proofs as pure pattern-matching. This idea is pushed further in the second part of this dissertation, and connected to realisability semantics, where the structured data is interpreted algebraically, and the consumption of such data is modelled with the use of an orthogonality relation. Most of this part has been proved in the **Coq** proof assistant.

Polarities and focussing were also introduced with applications to logic programming in mind, where computation is proof-search. In the third part of this dissertation, we push this idea further by exploring the roles that these concepts can play in other applications of proof-search, such as theorem proving and more particularly automated reasoning. We use these concepts to describe the main algorithm of SAT-solvers and SMT-solvers: **DPLL**. We then describe the implementation of a proof-search engine called **PSYCHE**. Its architecture, based on the concept of focussing, offers a platform where smart techniques from automated reasoning (or a user interface) can safely and trustworthily be implemented via the use of an API.

Acknowledgements

It is difficult to determine when to write an habilitation thesis, what it will include and which format it will have. In the case of a Ph.D. thesis, such issues are often entirely resolved by the end of the Ph.D. funding and the input of the Ph.D. adviser. Therefore, encouragements to write an habilitation thesis are all the more important, and for this reason I am grateful to Olivier Bournez, head of our research laboratory (LIX), who first planted the idea in my head, as well as to Benjamin Werner, head of Polytechnique's C.S. department, for his friendly support.

In fact I surprisingly found, in the habilitation process, no other obstacles than those I encountered on my own, as everybody that I interacted with helped me overcome them:

In particular, I am indebted to my team leader Dale Miller and to my Ph.D. adviser Roy Dyckhoff, who recommended me when I enrolled. I am grateful to Stéphanie Druetta and Marie-Christine Mignier¹ at the Paris-Sud administration, as well as to Dominique Gouyou-Beauchamps, whose work at the C.S. School of Doctoral Studies guarantees the high standards of the degree towards which this dissertation is submitted. They all worked impressively efficiently, especially given the tight schedules that characterised my application process.

Christine Paulin-Mohring was kind enough to sponsor my application and be one of the first people to look into my dissertation. I thank her, as well as Sylvain Conchon, David Delahaye, Didier Galmiche and Laurent Regnier, for accepting to sit on the defence panel. For the same reason, but also for having reported on my dissertation, I wish to thank Wolfgang Ahrendt, Hugo Herbelin and Frank Pfenning: they honoured me with their time and interest.

The work described in this dissertation benefitted from many years of interactions within my departement and elsewhere: I am very grateful to all past and present assistants at LIX² and of course I am also very grateful to all of my colleagues (whether or not they are inclined towards Logic) for making Polytechnique's C.S. department and my research community such a great working environment. I am particularly thankful to Assia Mahboubi for the many years of scientific interaction and friendship that have passed since we studied at ENS Lyon.

I understand that an *Habilitation à Diriger des Recherches* also has to do with student supervision. Students in general have been a very important part of my work ever since I came back to France. The first part of this dissertation results from my teaching at MPRI. The PSYCHE engine was inspired by a few lines of code from a undergraduate students' project at ESIEA, while its latest development results from Damien Rouhling's internship at LIX. I also learnt a lot from interacting, in very different styles, with my two Ph.D. students Mahfuza Farooque and Alexis Bernadet; congratulations again for your work and for having successfully defended your theses before mine. In brief, I am very grateful to all of the students I have worked with.

Besides Academia, I thank my parents and my friends who supported me in the habilitation process, particularly the Rémi(e)s who repeatedly dared to ask me for news updates, no matter how uninteresting to them the topics of *Polarities and Focussing* must have been. Finally, no-one supported me more than my wife Claire, whose patience I challenged by opening my laptop most nights and on every holiday: for your unwavering love and cups of coffee I am forever grateful.

¹who successfully convinced me I once was a student at Paris-Sud in the 90s, which I had no recollection of

²This includes Martine Thirion for her help in making the defence happen.

*To all of those who are not computer scientists or logicians,
here is to inverted witch-hats and squash courts.*

Table of Contents

Introduction	1
Notations and prerequisites	9
I The Curry-Howard view of classical logic - a short introduction	11
1 Classical proofs as programs	13
1.1 Curry-Howard correspondence: concepts and instances	14
1.1.1 Simply-typed combinators	14
1.1.2 Simply-typed λ -calculus	16
1.1.3 The categorical aspect	18
1.1.4 Applying the methodology to other systems	20
1.2 Continuations and control	22
1.3 Contributions in the 90s	25
1.4 System L	30
1.5 Non-confluence of cut-elimination in classical logic	35
1.6 Continuations, Call-by-Name and Call-by-Value	37
1.7 Classical logic and CBN/CBV	42
1.7.1 Identifying CBN and CBV in System L	43
1.7.2 Two stable fragments	46
1.7.3 Denotational semantics of CBN and CBV	47
Conclusion	49
2 Orthogonality, normalisation and witness extraction	51
2.1 Revisiting Proofs of Strong Normalisation for System <i>F</i>	52
2.1.1 Orthogonality models and the Adequacy Lemma	53
2.1.2 Applicative orthogonality models and Strong Normalisation	54
2.2 Adapting the approach to classical calculi	55
2.2.1 The case of a confluent calculus	56
2.2.2 The case of a non-confluent calculus	57
2.3 Orthogonality models for extracting witnesses from classical proofs	60

Conclusion	64
3 Polarisation and focussing	65
3.1 Recovering confluence by polarisation	66
3.1.1 Symmetry, asymmetry, and η -expansions	66
3.1.2 Towards polarised System L	68
3.1.3 Focussing	71
3.1.4 Weak η -conversion	72
3.1.5 Related works	73
3.2 Computational interpretation of a focussed calculus	74
3.2.1 Informal relation to System L	76
3.2.2 Identifying phases as atomic steps	77
3.2.3 Functional interpretation as pattern-matching	82
Conclusion	84
II Abstract focussing	85
4 An abstract focussed sequent calculus - without quantifiers	89
4.1 Presentation of the system	90
4.1.1 Atoms, molecules, typing decompositions and typing contexts	90
4.1.2 Logical connectives	92
4.1.3 Definition of the system	92
4.2 Capturing existing systems	93
4.3 Examples in propositional logic	95
4.3.1 Polarised classical logic - one-sided	95
4.3.2 Polarised classical logic - two-sided	99
4.3.3 Polarised intuitionistic logic	100
4.4 Examples of labels implementation: De Bruijn's indices and levels	104
4.4.1 Labels for classical logic	104
4.4.2 Labels for intuitionistic logic	105
5 An abstract focussed sequent calculus - with quantifiers	107
5.1 Presentation of the system	108
5.1.1 Quantifying structure	108
5.1.2 Atoms and Molecules	108
5.1.3 Typing decompositions and typing contexts	109
5.1.4 Logical connectives	112
5.1.5 Definition of the system	112
5.2 Extending LAF_{K1} with quantifiers	112

6	Realisability models of abstract focussing	117
6.1	Model structures and the interpretation of proof-terms	118
6.2	Realisability algebras, interpretation of types & Adequacy	119
6.3	A more concrete class of LAF instances	121
6.3.1	LAF instances with eigenlabels	122
6.3.2	LAF_{K1} is a LAF instance with eigenlabels	124
6.3.3	LAF instances with eigenlabels are LAF instances	125
6.4	A more concrete class of realisability algebras	127
6.5	Example: boolean models to prove Consistency	129
7	Transforming proofs in the abstract focussed sequent calculus	133
7.1	Head reduction	134
7.2	Head normalisation	136
7.3	Re-using proofs	137
7.4	Cut-elimination	140
7.5	Conclusion and further work: Strong normalisation	144
III	Theorem proving	147
8	DPLL($\mathcal{T}$) as proof-search in a focussed sequent calculus	151
8.1	A version of LKF to work modulo a theory: $\text{LK}^p(\mathcal{T})$	152
8.1.1	Background	152
8.1.2	Definitions	153
8.2	Bisimulation with the DPLL($\mathcal{T}$) procedure	156
8.2.1	The elementary DPLL($\mathcal{T}$) procedure	156
8.2.2	Simulation of the elementary DPLL($\mathcal{T}$) procedure in $\text{LK}^p(\mathcal{T})$	158
8.2.3	Completing the bisimulation	161
8.2.4	More advanced features	163
8.3	Future work: Relation to abstract focussing	163
8.3.1	On-the-fly polarisation	164
8.3.2	Extending LAF to $\text{LAF}(\mathcal{T})$	164
9	The PSYCHE system	167
9.1	Motivation	168
9.2	Overview and general architecture	170
9.3	PSYCHE's Kernel	171
9.4	Plugins	172
9.4.1	Specifications and implemented instances	172
9.4.2	Memoisation and lemma learning	173
9.5	Decision procedures	175
	Conclusion: Testing and perspectives	176

10 Conclusion and further work	179
10.1 Summary of the topics covered by this dissertation	179
10.2 Further work	180
Bibliography	182
Index	195
A Basic definitions for categories	197

Introduction

This dissertation concerns two fundamental ways in which mathematical proofs relate to computation: *proof-normalisation* and *proof-search*.

- The key idea, in the view of “computation as proof-normalisation”, is that mathematical proofs can be composed in a modular way and that composed proofs can (sometimes) be “simplified” into *normal forms* by a normalisation procedure. The most well-known tool to compose proofs (though not the only one) is a specific reasoning step known as *cut* in the proof formalism of *Sequent Calculus* and known as *cut* or *detour* in the proof formalism of *Natural Deduction* [Gen35]. The normalisation process that turns proofs with cuts into cut-free proofs, known as *cut-elimination*, strongly relates to the computational paradigm of Functional Programming, as shown by the Curry-Howard correspondence [CF58, How80].
- The view of “computation as proof-search”, on the other hand, considers a mathematical formula as the input of computation, and a proof of that formula as its output. This strongly relates to the computational paradigm of Logic programming, as described for instance by the seminal paper on *uniform proofs* [MNPS91].

Interestingly enough, investigating normal forms for proofs is useful for both views: for the former, to understand to which proofs all other proofs should reduce; for the latter, to only search for proofs in normal form and thus restrict the search space in efficient ways. For example, both kinds of computation are often taken to produce cut-free proofs (though not always).

In fact, two key concepts in the study of normal forms have proved useful for both views: *polarities* and *focussing*. In proof-search, they were used to design variants and generalisations of logic programming languages [AP89, And92, LM09]. In proof-normalisation, they were used to understand the semantics of, and design meaningful variants of, cut-elimination procedures [Lau02, LQdF05, MM09] (building on previous work [DJS95, DJS97]).

Roughly speaking, polarities and focussing generalise the idea that a formula of the form

$$\forall x_1 \exists y_1 \forall x_2 \exists y_2 \dots \forall x_i \exists x_i \dots$$

suggests a two-player game: the opponent gets to choose x_1 , and depending on x_1 the proponent gets to choose y_1 , after which the opponent gets to choose x_2 , etc until some criterion (determined by the final ‘...’) decides who has won, given all the choices that have been made. Of course, what the exact rules of the game are, what a winning strategy is, etc depends on the logic considered and its proofs (for instance in classical logic, one can backtrack on a previous choice, using the input of the adversary). Polarities and focussing generalise this idea to all connectives (not only quantifiers), with some *positive connectives* “corresponding to” proponent’s moves and *negative connectives* “corresponding to” opponent’s moves.

The range of fields that build on, or benefit from, the two computational aspects of mathematical proofs, is broad. This dissertation engages in two of them which may seem distant from each other: program semantics and theorem proving. More specifically, it proposes the use of polarities and focussing as the core concepts to approach a field of topics ranging from *realisability semantics* to *automated reasoning*. We can briefly illustrate how polarities impact those two areas.

Realisability semantics (see e.g. [Kle45, VO02]) is a way to interpret a mathematical formula (in a broad sense, including a program type) as a *specification* that an object of a certain kind (such as a computer program or a mathematical proof) may or may not satisfy. This interpretation as specifications (i.e. what it means for an object to satisfy them) is defined by induction on the syntax of formulae, and refers to the object either by its internal structure or by the way it behaves when placed in a well-chosen environment. Realisability semantics have been studied for various logics and systems, and a particular approach emerged from classical logic and Girard’s *linear logic* [Gir87], namely *orthogonality*. This approach is sometimes described as *classical realisability* [DK00, Kri01] (even though it may be used for other logics than classical logic).

This dissertation aims at the very essence of orthogonality-based realisability, by building an abstract semantics only based on polarities and focussing:

- if a formula starts with a positive connective, then the criterion determining whether an object satisfies the formula’s specification refers primarily to the object’s internal structure;
- if a formula starts with a negative connective, then the criterion refers to the object’s behaviour when placed in a well-chosen environment.

Automated reasoning (see e.g. [RV01]) concerns the numerous algorithmic techniques by which the validity or the satisfiability of mathematical formulae can be determined. Since a formula is valid if and only if it has a proof, an obvious approach to automated reasoning is proof-search. The basic core of logic programming, for instance, can be understood as proof-search on *Horn clauses*, and in that respect it can be seen as a very specific area of automated reasoning. Now the reason why proof-search on Horn clauses also provides a meaningful computational paradigm is because this class of formulae makes a simple goal-directed proof-search strategy logically complete, with well-identified backtrack points and a reasonable covering of the proof-search space. This still holds when the class is extended to *hereditary Harrop formulae* [MNPS91], and can hold on a wider class of formulae if logical connectives (and atoms) are tagged with polarities:

- Negative connectives can be decomposed with *invertible* inference rules: a goal-directed proof-search strategy performs the bottom-up application of those rules as basic proof-search steps, without loss of generality;³ in other words, no backtracking is necessary on the application of such steps, even though other steps were possible.
- Positive connectives are the (De Morgan’s) duals of negative connectives, and their decomposition rules are not necessarily invertible, so a goal-directed proof-search procedure creates backtrack points when applying them bottom-up.

To what extent these ideas can be useful for a wider area of automated reasoning (than logic programming) remains a recent field, with numerous open questions but already with a

³If the goal was provable, it remains provable after applying the step.

couple of implementations available, such as Imogen [MP08] (using the *inverse method*) and Tac [BMS10] (using bottom-up proof-search). This dissertation explores a particular take on this, with its own implementation: PSYCHE [Psy].

This dissertation is therefore a journey through the above topics, trying to connect them with e.g. common formalisms. It is organised in three parts.

Part I of this dissertation is a short introduction to the adaptation, to classical logic, of the Curry-Howard correspondence, already mentioned above in the view of “computation as proof-normalisation”. Also known as the “proofs-as-programs paradigm”, the correspondence emerged with a strong flavour of constructive mathematics, so its adaptation to classical logic only emerged in the past 25 years [Gri90]. This part explores (some of) the contributions that have been made in that period, where we shall see the important roles of polarities and focussing. While it starts from Parigot’s $\lambda\mu$ -calculus [Par92] and ends with a Zeilberger-style system [Zei08a, Zei08b], this part mostly uses Curien and Herbelin’s System L [CH00] as a common framework to express and connect the concepts pertaining to the computational interpretations of classical proofs.

Chapter 1 describes the basic set-up, viewing classical proofs as programs. In particular, we give an overview of how classical reasoning corresponds to the use of control operators [Rey72, SW00, Fel87] that let programs capture the contexts within which they are being evaluated. We show standard ways of building meaningful operational and denotational semantics for cut-elimination, which correspond to the Call-by-Name and Call-by-Value evaluation strategies in programming [Plo75], and to control and co-control categories in category theory [Sel01].

Chapter 2 explores the concepts and techniques based on orthogonality: orthogonality models form the classical version of realisability semantics [DK00, Kri01], as well as providing methodology to prove strong normalisation results [Par97, LM08], i.e. the termination of well-typed programs. We also illustrate another use of orthogonality models for extracting, out of a classical proof of an existential formula of arithmetic (more precisely, a Σ_1^0 -formula), a term witnessing the existence; this technique is due to Miquel [Miq09, Miq11]. Out of orthogonality techniques we shall see the notion of polarity naturally emerge.

Chapter 3 formalises this concept, inspired by a discussion on η -conversion and observational equivalence. A new semantics for the evaluation of classical programs is inferred from the use of polarities (as in [MM09]), and three different notions of normal forms are identified, out of which the concept of focussing naturally emerges. The strongest version of focussing, namely system LKF [LM09], organises each proof into an alternation of phases (similar to the alternation between proponent’s moves and opponent’s moves in the intuitive view of the formula $\forall x_1 \exists y_1 \forall x_2 \exists y_2 \dots \forall x_i \exists y_i \dots$). The chapter then describes how each phase can be collapsed into one inference step, giving rise to a presentation of LKF in the style of “big-step focussing”. It then describes the computational interpretation of this in terms of pattern-matching, along the lines of [Zei08a, Zei08b].

Part II of this dissertation takes this last idea further and presents new material. Stripping focussed systems off the concept of connective and off the inductive structure of mathematical formulae, we only keep the core mechanisms of focussing to define a highly abstract system for big-step focussing, called LAF, whose computational interpretation is pure pattern-matching.

One of the main goals is to formalise the strong ties between Zeilberger-style systems and orthogonality models.

Chapter 4 presents the syntax and the typing system of a quantifier-free version of the LAF abstract system, which is modular in its syntax for atoms and formulae, in its logical connectives, in the logic used, and in the implementation of variables. The chapter shows how the abstract system can be instantiated to capture existing focussed systems such as LKF and its intuitionistic variant LJF.

Chapter 5 presents the extension with quantifiers of that abstract focussed system LAF. Different approaches may lead to either a treatment of quantifiers along the lines of the ω -rule [Hil31, Sch50] (where a formula “ $\forall n \in \mathbb{N}, A(n)$ ” may be proved by providing a proof for each natural number), or a treatment that forces to prove a universal formula in a uniform way: for this we need to extend LAF with a mechanism that generalises eigenvariables.

Chapter 6 presents the realisability models of the abstract system LAF, finally formalising the connection between big-step focussing and orthogonality models: indeed we lift the orthogonality models of Chapter 2 to our abstract framework, and we prove the Adequacy Lemma, that relates typing to realisability, i.e. syntax to semantics. We present instances of orthogonality models such that the Adequacy Lemma immediately provides the logical consistency of the LAF system.

Chapter 7 explores proof transformations in the abstract system LAF; we start with an abstract machine to perform head-reduction, thus revealing the actual pattern-matching mechanism of the proof-term calculus. Using the realisability models of Chapter 6, we show that head-normalisation terminates on typed proof-terms. We then describe, via a notion similar to that of free variables, how to identify the parts of a sequent that have actually been used in its proof (which is useful if the proof is to be re-used for a sequent that is similar). We also use this to extend the abstract machine into a big-step operational semantics that evaluates a proof-term as a normal form that is cut-free. Adapting the orthogonality model to this big-step operational semantics, we prove that every typed proof-term does evaluate as a cut-free form, and conclude the cut-elimination result for the LAF system.

Part III of this dissertation concerns the roles that polarities and focussing can have in automated reasoning and more generally theorem proving (i.e. proof construction may also be interactive). Originally aiming at classical logic (and therefore departing from the Imogen [MP08] and Tac [BMS10] provers), we investigated one of the most popular automated reasoning techniques for classical propositional logic (a.k.a. SAT-solving): DPLL [DP60, DLL62], as well as its extension known as DPLL($\mathcal{T}$) [NOT06] for solving SAT-modulo-theories problems (SMT).

Chapter 8 aims at describing and simulating DPLL($\mathcal{T}$) runs as bottom-up proof-search in a focussed system for classical logic. We therefore present an extension of system LKF that allows atoms to be assigned polarities *on-the-fly* during proof-search, and that integrates the possibility to call a procedure that decides whether a conjunction of (ground) atoms is consistent with a given input theory $\mathcal{T}$. The resulting system, $\text{LK}^p(\mathcal{T})$, is used to establish a bisimulation result between proof-search and DPLL($\mathcal{T}$) runs. Based on the fact that LKF can be seen as an instance of LAF, the chapter then discusses how $\text{LK}^p(\mathcal{T})$ could be seen as an instance of a generalisation of LAF that could work *modulo* the theory $\mathcal{T}$.

Chapter 9 describes a small prototype called PSYCHE [Psy] implementing bottom-up proof-search in an extension of $\text{LK}^p(\mathcal{T})$ with quantifiers and meta-variables. Highly modular with

respect to the decision procedure and the proof-search strategy it can run with, PSYCHE comes with a strategy *plugin* that implements the simulation of $\text{DPLL}(\mathcal{T})$ described in Chapter 8, and can also perform pure first-order reasoning.

Chapter 10 concludes this dissertation, in particular by giving an informal description of PSYCHE’s mechanisms for quantifiers, and hinting at what could be achieved with them, in particular in the combination of first-order reasoning with decision procedures. It finally presents the LAF system as the theoretical foundations for the next version of PSYCHE.

Note that the whole of Part II is admittedly technical, which is due to two reasons:

- The first reason is the systematic search for the greatest generality (and therefore strength) in the definitions and theorems. It was a goal in itself to determine exactly which ingredients are necessary and which are disposable for the system to make sense, for the models to be built, and for the theorems to be proved. Hypotheses are systematically weakened and structures are systematically parameterised to achieve this. The result of course is a highly parameterised framework with complex yet precise specifications. In order to digest this technicality with confidence, most of the proofs have been formalised [GL14] in the proof assistant Coq [Coq], which was particularly useful to refine the definitions and theorems according to the above methodology.
- The second reason is that the development of this abstract framework was not only done for the sake of it, but also to provide the foundations of (the next version of) our PSYCHE implementation. Abstraction in the theoretical framework translates to genericity in the code, making the implementation more versatile, decomposing its architecture into smaller modules that could more easily be shown to be correct. Therefore, when introducing as a mathematical structure a tuple such as

$$(\mathbb{S}, \text{Lab}_e, \mathbb{T}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \mathbb{R}, \text{Co}, \text{Pat}, \Vdash)$$

satisfying a long list of axioms, we really have in mind an OCaml module providing the corresponding types and functions and satisfying the corresponding specifications. Hence the verbosity of our axiomatic structures in Part II.

Personal note

This section aims at relating this dissertation to the papers I have published in the recent years.

Firstly, this dissertation lies within the very broad field of *Computational Logic*, on which Didier Galmiche and I edited a special issue of the Journal of Logic and Computation, in honour of Roy Dyckhoff [GGL14].

My interest for the topics developed in this dissertation can be traced back to the first paper I wrote as the sole author [Len03], relating Curien and Herbelin’s work on the computational interpretations of classical logic [CH00] to Urban’s [Urb00].⁴ However, such topics stayed in the slow-cooker at the back of my mind, as my next contributions mostly concerned intuitionistic systems, which could more simply be related to the λ -calculus, the Curry-Howard correspondence, and Type Theory:

In [KL08] we explored a Call-by-Name cut-elimination procedure for the intuitionistic sequent calculus, in [DL07] we explored the focussed sequent calculus LJQ, and in [Len08] I

⁴Despite its critical typos, it surprisingly appears to be my most cited paper.

proved some conjecture about the termination of a Call-by-Value λ -calculus. Although these contributions are not directly included in this dissertation, the work that I did around that time greatly contributed to my understanding of focussing and cut-elimination strategies.

Still in intuitionistic logic, two more recent contributions broached the topics that this dissertation approaches under the focussing angle, namely realisability and automated reasoning: In [BGL12] we develop a simple presentation of Hyland’s effective topos [Hyl82] which is based on realisability concepts; in [LDM11] we developed a focussed sequent calculus that can describe proof-search in the type theory behind the proof-assistant Coq [Coq].

More directly included in this dissertation are the publications [LM08] and [BL11b, BL11c, BGL13], all of which formalise proofs of strong normalisation with orthogonality techniques, aiming at genericity. In [LM08] we compare Barbanera and Berardi’s technique based on symmetric reducibility candidates [BB96] with the basic orthogonality technique. In [BL11b, BL11c, BGL13], we formalise an abstract notion of orthogonality model and describe, as instances of this notion, several variants of proofs for the strong normalisation of System F [Gir72].

Chapter 2 of this dissertation covers these contributions with a systematic orthogonality model construction. It also uses the same orthogonality framework to describe an interesting application of classical realisability to witness extraction (due to Miquel [Miq09, Miq11]).

Over the recent years, a greater proportion of my research was devoted to the use of focussing for proof-search, in the context of our ANR-funded project on *Proof-Search Control in Interaction with domain-specific methods* [PSI]. Since the concept of focussing emerged with motivations for logic programming, it was natural to explore whether the concept could also impact automated reasoning. We first explored propositional problem solving, which can be seen either as proof-search or as satisfiability (SAT) solving: More precisely, we described in [FLM12b, FLM12a, FGLM13] how one of the main procedures, namely DPLL [DP60, DLL62], can be seen as the gradual construction of proof-trees in a focussed sequent calculus. We actually did this in combination with decision procedures, so as to describe SMT-solving in terms of proof-search; this required the extension of sequent calculus with such procedures [FL11, FGL13]. All this was put together in Farooque’s Ph.D. thesis [Far13] which I supervised, and where another class of automated reasoning techniques, namely *tableaux* methods, are also simulated in the same focussed sequent calculus.

In the present dissertation, the above contributions are not developed in as many details as in [Far13]. However, they form the theoretical basis of the PSYCHE prototype [Psy], of which I am the main developer; and the software is the topic of Chapter 9, which covers the system description [GL13].

The material presented in this dissertation does not only come from publications. Some of it relates to an active teaching activity: in particular, Part I approximately covers the material that I teach at M.Sc. level in Paris, with the most advanced parts inspired by Munch-Maccagnoni’s work relating focussing and classical realisability [MM09], and Zeilberger’s work on big-step focussing and pattern-matching interpretations [Zei08a, Zei08b].

Part II presents entirely new material, rather than published work (or survey thereof), that builds on those two inspirational topics: Zeilberger’s framework seemed particularly appropriate to relate focussing and classical realisability at a particular abstract level. The proposal is to make this the theoretical foundation of PSYCHE’s next version, and in that it connects to Part III this dissertation.

On the other hand, several publications do not (yet) relate (or only very remotely) to this dissertation, since they are too disconnected from its topic: In [GL08, GL09], we developed a λ -calculus inspired by Nominal Logic [Pit03], with a special construct in order to represent binding in data-structures; the goal is to allow incomplete terms within the scope of binders, without blocking α -conversion or computation. In [BL11a] we studied intersection type systems [CD78] for the λ -calculus, in a non-idempotent version similar to de Carvalho's [dC05, dC09], but such that the length of the longest β -reduction sequence starting from a strongly normalising term, can be directly read from its typing tree.

Finally, the careful reader will note that this dissertation not only has little material in common with my Ph.D. thesis [Len06], but it is not even in its direct continuation: I do not present here refinements or developments of its contributions, but rather a thesis that complements my doctoral work in the topics of my interest.

Notations and prerequisites

In this dissertation, we assume the reader to be already familiar with some areas and concepts of logic and computer science. Unless specifically given, the notations and definitions used in this dissertation are rather standard, and formally follow [Len06]. The areas and concepts are:

- Set theory; see e.g. [Kri71].
In particular we will use the concepts of, and notations for, subsets, power sets, union, intersection and difference of sets, relations, functions, injectivity, surjectivity, etc. Our notation for the power set of A is $\mathbb{P}(A)$. Our notation for the set of total functions from A to B is denoted $A \rightarrow B$; the set of partial functions from A to B is denoted $A \rightharpoonup B$. We also assume the reader to be familiar with natural numbers, lists and trees.
- The standard difference between object-level and meta-level.
In particular, variables of the meta-level are called meta-variables and (unless otherwise stated) “rules” and “systems” are meta-level devices (i.e. a rule has no existence at the object level, but its instances do -and the collection of them, for example).
- Trees and derivations.
We use inference rules and systems to define sets of (valid) derivations and derivability of judgements, as well as partial derivations; when we state that a rule is derivable/admissible/invertible (in a system) we actually mean that its instances are derivable/admissible/invertible (in the collection of derivations defined by the system).
- Rewriting (first-order and higher-order); see e.g. [Ter03].
In particular, the notations $\rightarrow^n$, $\rightarrow^+$, $\rightarrow^*$, $\leftrightarrow^*$, denote the composition n times of a (binary) relation $\rightarrow$, the transitive closure, the transitive and reflexive closure, and the transitive, reflexive and symmetric closure, of the relation $\rightarrow$, respectively.

We assume that the reader is familiar with the properties of confluence and Church-Rosser, weak normalisation, strong normalisation, and the usual techniques to prove them, in particular the simulation techniques.

Following [Len06], the notation

$$(\gamma) \quad M \longrightarrow N$$

introduces a rewrite rule whose contextual closure (or more precisely, the contextual closure of its instances) is denoted $M \longrightarrow_\gamma N$. We also use this notation when γ is a system of rules.

Our languages will often be made of terms whose syntax is defined by a BNF-grammar. Some of its syntactic categories may contain *variables*. We assume the reader is familiar with variable binding, α -conversion and *equivariance*; specifying binders and their scopes automatically defines what the *free variables* of a term, denoted $FV(t)$, are; *capture-*

avoiding substitution of u for x in t is denoted

$$\{^u/_x\} t$$

where x is a variable of some syntactic category with variables and u is a term of that syntactic category.

- Basic proof theory; see e.g. [TS00]. In particular, standard proof formalisms such as Natural Deduction and Sequent Calculus, for intuitionistic and classical logic (propositional and first-order).

Part I

The Curry-Howard view of classical logic - a short introduction

Chapter 1

Classical proofs as programs

Contents

1.1	Curry-Howard correspondence: concepts and instances	14
1.1.1	Simply-typed combinators	14
1.1.2	Simply-typed λ -calculus	16
1.1.3	The categorical aspect	18
1.1.4	Applying the methodology to other systems	20
1.2	Continuations and control	22
1.3	Contributions in the 90s	25
1.4	System L	30
1.5	Non-confluence of cut-elimination in classical logic	35
1.6	Continuations, Call-by-Name and Call-by-Value	37
1.7	Classical logic and CBN/CBV	42
1.7.1	Identifying CBN and CBV in System L	43
1.7.2	Two stable fragments	46
1.7.3	Denotational semantics of CBN and CBV	47
Conclusion		49

The Curry-Howard correspondence [CF58, How80] has been one of the most fruitful connections between proofs and computation: As one of the embodiments of constructivism, where mathematical proofs bear computational content, the correspondence naturally emerged in the context of minimal and intuitionistic logic, and gave rise to the field of Type Theory [ML82, ML84].

Despite the non-constructive character of proofs in classical logic, arising from the Law of Excluded Middle, or the Double Negation Elimination etc, it is natural to investigate what part of the Curry-Howard correspondence can still be built for that logic.

In this chapter we review the foundations of the correspondence in the framework of classical logic, along the main lines of investigation that were explored over the past 25 years since Griffin's seminal work [Gri90].

The first step in this programme is to turn a proof format for classical logic into a typing system for a language. For such a language to be of computational nature, an operational semantics and/or a denotational semantics has to be designed.

Section 1.1 reviews the basic concepts of the Curry-Howard correspondence, both in their original framework and at a more abstract level. Section 1.2 present some concepts in programming, namely continuations and control, which will prove useful to understand classical proofs as programs. Section 1.3 presents early formalisations of the above concepts as proof-term calculi for classical logic while Section 1.4 presents in more details one of the most convenient ones, which relates to Gentzen’s classical sequent calculus. Section 1.6 uses continuations to describe the evaluation strategies known as Call-by-Name and Call-by-Value while Section 1.7 explains how this can be used to build semantics for classical proofs.

1.1 Curry-Howard correspondence: concepts and instances

The correspondence relates logic to programming languages, and is sometimes taken to involve a third aspect, namely category theory (as it forms a popular framework to build the semantics of programming languages). Table 1.1 gives a high-level view of the correspondence,¹ which operates at several levels: mathematical formulae, or propositions, correspond to the types of a given programming language; proofs of such propositions correspond to programs that can be given the corresponding type; the way proofs can be composed corresponds to the way programs can be composed/applied; finally (and this is where we adopt the view of computation as proof-normalisation), cut-elimination corresponds to program execution.

Logic	Programming language	Categories
Propositions	Types	Objects
Proofs	Typed programs	Morphisms
Cut/Composition	Program composition	Morphism composition
Cut-elimination	Program execution	Equality of morphisms (commuting diagrams)

Table 1.1: High-level view of the Curry-Howard correspondence

The rest of this section gives a brief overview of the correspondence in the framework of minimal and intuitionistic logic. An in-depth presentation of the correspondence can be found in the book [SU06].

1.1.1 Simply-typed combinators

The original instance of the correspondence was given in the study of combinators [CF58], which yields a simple language made of three basic programs **I**, **K**, **S** and with program application as its only construct:

DEFINITION 1 (The (**I**, **K**, **S**)-combinatoric system)

The syntax is given by the following grammar:

$$M, N, \dots ::= \mathbf{I} \mid \mathbf{K} \mid \mathbf{S} \mid M N$$

The last construct, *program application*, is associative to the left, i.e. $(M N) P$ can be abbreviated as $M N P$.

¹We use the expression (Curry-Howard) “correspondence” rather than the popular (Curry-Howard) “isomorphism”, as it is difficult to specify what the isomorphism exactly is before specifying exactly what formal systems we intend to relate.

and its operational semantics is given by the following first-order rewrite system

$$\begin{aligned} \mathbf{I} M &\longrightarrow M \\ \mathbf{K} M N &\longrightarrow M \\ \mathbf{S} M N P &\longrightarrow M P (N P) \end{aligned}$$

※

Clearly, the operational semantics defines $\mathbf{I}$ as the identity, while $\mathbf{K}$ provides erasure and $\mathbf{S}$ provides duplication. The reduction relation is confluent and defines a model of computation that turns out to be Turing-complete. At the cost of losing that property, the language can be given an intuitive typing system, using *simple types*:

DEFINITION 2 (Simple types) *Simple types* are defined by the following grammar:

$$A, B, \dots ::= a \mid A \rightarrow B$$

where a ranges over a fixed set of elements called *atomic types*. The symbol $\rightarrow$ is associative to the right, i.e. $A \rightarrow (B \rightarrow C)$ can be abbreviated as $A \rightarrow B \rightarrow C$. ※

The typing system is defined as follows:

DEFINITION 3 (Simple types for the ($\mathbf{I}$, $\mathbf{K}$, $\mathbf{S}$)-combinatoric system)

Typing is a binary relation between terms and simple types, denoted with expressions such as $\vdash M : A$. That relation is defined for the combinators as follows:

$$\begin{aligned} \vdash \mathbf{I} : A \rightarrow A \\ \vdash \mathbf{K} : A \rightarrow B \rightarrow A \\ \vdash \mathbf{S} : (A \rightarrow (B \rightarrow C)) \rightarrow (A \rightarrow B) \rightarrow (A \rightarrow C) \end{aligned}$$

and program application is typed by the following rule:

$$\frac{\vdash M : A \rightarrow B \quad \vdash N : A}{\vdash M N : B}$$

Derivability of the typing statement $\vdash M : A$, from the above axioms and using the program application rule, is denoted $\vdash_{\text{stC}} M : A$. ※

The reduction defined by the rewrite system preserves types, a property called Subject Reduction:

THEOREM 1 (Subject reduction for simply-typed combinatoric system)

If $\vdash M : A$ and $M \longrightarrow N$ then $\vdash N : A$. ※

Proof: By induction on the derivation of $M \longrightarrow N$, with the base cases corresponding to the 3 rewrite rules themselves. □

The essence of the Curry-Howard correspondence, is the simple remark that, viewing the functional type construct $\rightarrow$ as the logical symbol for implication, simple types are isomorphic² to the syntax of formulae for propositional minimal logic [Joh36] and that typing derivations are isomorphic to proofs in a particular Frege-Hilbert system [Fre79, Hil28] for minimal logic.

²The isomorphism with simple types assumes that atomic types are isomorphic to atomic formulae.

DEFINITION 4 (Propositional minimal logic)

Formulae of minimal logic are defined by the following grammar:

$$A, B, \dots ::= a \mid A \Rightarrow B$$

where a ranges over a fixed set of elements called *atomic formulae*.

Proofs are the derivations built with the *Modus Ponens* rule

$$\frac{\vdash A \Rightarrow B \quad \vdash A}{\vdash B}$$

from the axioms:

$$\begin{aligned} A &\Rightarrow A \\ A &\Rightarrow B \Rightarrow A \\ (A &\Rightarrow (B \Rightarrow C)) \Rightarrow (A \Rightarrow B) \Rightarrow (A \Rightarrow C) \end{aligned}$$

The symbol $\Rightarrow$ is associative to the right.

Derivability of $\vdash A$, from the above axioms and using the program application rule, is denoted $\vdash_{\text{FH}} A$. ※

Via the correspondence, the Subject Reduction property allows the view of the reduction relation as a proof-transforming procedure.

1.1.2 Simply-typed λ -calculus

Curry's view about minimal logic was extended by Howard to intuitionistic first-order arithmetic [How80]. A different format was also proposed, both for proofs and for programs, and became perhaps the most popular setting for the Curry-Howard correspondence: Natural Deduction [Gen35] was the formalism used for proofs, and the λ -calculus [Chu41] was the formalism used for programs. This instance of the Curry-Howard correspondence that we present is a version of natural deduction using *sequents* and a version of the simply-typed λ -calculus using typing contexts.

DEFINITION 5 (λ -calculus) The syntax of the λ -calculus is given by the following grammar:

$$M, N, \dots ::= x \mid \lambda x.M \mid M N$$

where x ranges over a denumerable set of *variables*, and the construct $\lambda x.M$ binds x in M .³

Standard conventions are used for parentheses [Bar84]: the scopes of binders extend as much as parentheses allow (i.e. $\lambda x.M N$ abbreviates $\lambda x.(M N)$); program application is associative to the left (i.e. $M N P$ abbreviates $(M N) P$); moreover, binder can be grouped, so that $\lambda xy.M$ abbreviates $\lambda x.\lambda y.M$.

The following rewrite rules

$$\begin{aligned} (\beta) \quad (\lambda x.M) N &\longrightarrow \left\{ N / x \right\} M \\ (\eta) \quad \lambda x.M x &\longrightarrow M \quad \text{if } x \notin \text{FV}(M) \end{aligned}$$

define the reduction relations $\longrightarrow_\beta$, $\longrightarrow_\eta$ and $\longrightarrow_{\beta\eta}$. ※

As for the combinatoric system from Section 1.1.1, the reduction relations are confluent:

THEOREM 2 (Confluence) $\longrightarrow_\beta$, $\longrightarrow_\eta$ and $\longrightarrow_{\beta\eta}$ are confluent. ※

³As mentioned in the section about notations, specifying binders and their scopes automatically defines free variables, α -conversion, capture-avoiding substitution, etc.

Proof: See for instance [Bar84]. $\square$

DEFINITION 6 (Simply-typed λ -calculus) *Typing contexts* are finite maps from variables to simple types, with $()$ denoting the empty context (sometimes the notation $()$ is completely omitted), Γ, Γ' denoting the union of contexts Γ and Γ' (assuming it is defined), and $x:A$ denoting the singleton context mapping variable x to the simple type A .

The typing rules of the simply-typed λ -calculus are given in Fig. 1.

Derivability of the typing statement $\Gamma \vdash M : A$ in that system is denoted $\Gamma \vdash_{\text{st}\lambda} M : A$. $\ast$

$$\begin{array}{c} \hline \Gamma, x:A \vdash x:A \\ \hline \frac{\Gamma, x:A \vdash M:B}{\Gamma \vdash \lambda x.M : A \rightarrow B} \quad \frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N:A}{\Gamma \vdash M N : B} \end{array}$$

Figure 1: Simply-typed λ -calculus

As for the combinatoric system from Section 1.1.1, the reduction relations satisfy Subject Reduction:

THEOREM 3 (Subject reduction for simply-typed λ -calculus)

1. If $\Gamma, x:A \vdash M : B$ and $\Gamma \vdash N : A$ then $\Gamma \vdash \left\{ \frac{N}{x} \right\} M : B$.
2. If $\Gamma \vdash M : A$ and $M \rightarrow_{\beta\eta} N$ then $\Gamma \vdash N : A$.

$\ast$

Proof: See for instance [Bar84]. $\square$

Our second instance of the Curry-Howard correspondence relates the simply-typed λ -calculus with the Natural Deduction system $\text{NJ}_{\Rightarrow}$ for minimal logic.

DEFINITION 7 (Natural Deduction for minimal logic - $\text{NJ}_{\Rightarrow}$)

System $\text{NJ}_{\Rightarrow}$ is the inference system given in Fig. 2, where

- A, B range over formulae of minimal logic;
- Γ stands for a “collection” of formulae. By collection we mean either set or multiset,⁴ with Γ, Γ' denoting the union of Γ and Γ' , A denoting either the formula A itself or the singleton $\{A\}$ (or $\llbracket A \rrbracket$), while the empty set (or multiset) is sometimes omitted;
- $\Gamma \vdash A$ is a structure called *sequent*.

Derivations in that system are called *proofs* in $\text{NJ}_{\Rightarrow}$.

Derivability in $\text{NJ}_{\Rightarrow}$ of a sequent $\Gamma \vdash A$ is denoted $\Gamma \vdash_{\text{NJ}_{\Rightarrow}} A$. $\ast$

$$\begin{array}{c} \hline \Gamma, A \vdash A \\ \hline \frac{\Gamma, A \vdash B}{\Gamma \vdash A \Rightarrow B} \quad \frac{\Gamma \vdash A \Rightarrow B \quad \Gamma \vdash A}{\Gamma \vdash B} \end{array}$$

Figure 2: Natural Deduction for minimal logic - $\text{NJ}_{\Rightarrow}$

⁴That choice will change the number of proofs of a given formula.

Comparing Fig. 1 and Fig. 2 reveals our second instance of the Curry-Howard correspondence, although the exact meaning of the word ‘correspondence’ is in this case more subtle than for our first instance:

Clearly, the bijective aspect that pertains to the word “isomorphism” is jeopardised as $\overline{x:A, y:A \vdash x:A}$ and $\overline{x:A, y:A \vdash y:A}$ are clearly two distinct typing derivations which would both ‘correspond to’ the proof $\overline{A, A \vdash A}$ (whether we use sets or multisets). Moreover, binding introduces an ambiguity in the way we count typing derivations: is there one or infinitely many derivations of $\vdash \lambda x.x:A \rightarrow A$?⁵

For this reason, this dissertation takes the view that the interesting aspects of the Curry-Howard correspondence do not include the bijective aspect of an encoding from one system into another, but rather its compositionality (for trees), and the soundness and completeness properties:

In the present case, the forgetful encoding that maps every typing derivation to a proof is compositional with respect to the tree-structure of derivations; its surjectivity provides completeness of type inhabitation -whether there exists a λ -term of a given type- with respect to the provability of the corresponding formula; soundness is simply the fact that the tree obtained by forgetting variables and terms from a typing derivation is a correct proof.

These are the properties that we will aim at when investigating the variants of the Curry-Howard correspondence.

As for the combinatoric system from Section 1.1.1, the Subject Reduction property allows the view of the reduction relations $\rightarrow_\beta$, $\rightarrow_\eta$ and $\rightarrow_{\beta\eta}$ as proof-transforming procedures.

In summary, the most well-known settings for the Curry-Howard correspondence are:

Frege-Hilbert system	$\leftrightarrow$	Combinators (S,K,I)	[CF58]
Natural Deduction	$\leftrightarrow$	Typed λ -terms	[How80]

1.1.3 The categorical aspect

We now briefly mention what is sometimes considered a third aspect of the Curry-Howard correspondence, in category theory.

Categories can be used to shed a semantical light on the Curry-Howard correspondence. In our case, a particular kind of category provides models of the simply-typed λ -calculus: cartesian closed categories (CCC). In brief, CCC feature a terminal object, products, and exponential objects. We start with a few notational conventions:

NOTATION 8 (Category)

The class of morphisms from object A to object B is denoted $\text{hom}(A, B)$, and the expression $f: A \rightarrow B$ denotes that f is a morphism from A to B . Identity morphisms are denoted id_A , and the composition of $f: A \rightarrow B$ and $g: B \rightarrow C$ is denoted $f \cdot g: A \rightarrow C$.

⁵Formally, the typing system allows infinitely many premisses for that typing judgement, depending on the variable that we pick to place in the typing context with type A .

In a cartesian closed category (CCC),

- the *terminal object* is denoted 1 , with morphisms $1_A: A \rightarrow 1$
- the *product* of A and B is denoted $A \times B$, with projections denoted $\pi_1: A \times B \rightarrow A$ and $\pi_2: A \times B \rightarrow B$ (and more generally, the i^{th} projection from n objects is denoted $\pi_{i/n}$) and morphism pairing denoted $\langle f_1, f_2 \rangle: C \rightarrow A \times B$ for every $f_1: C \rightarrow A$ and $f_2: C \rightarrow B$.
- the *exponential* of A and B is denoted B^A , with morphisms $\text{eval}: B^A \times A \rightarrow B$ and a curried morphism $\Lambda g: X \rightarrow B^A$ for every $g: X \times A \rightarrow B$.

※

For a formal definition of the above concepts, see Appendix A.

DEFINITION 9 (Semantics of the simply-typed λ -calculus in a CCC)

Consider a cartesian closed category.

Consider a mapping that interprets every atomic type a as an object $\llbracket a \rrbracket$ of the CCC, and extend it to all simple types by defining $\llbracket A \rightarrow B \rrbracket$ as $\llbracket B \rrbracket^{\llbracket A \rrbracket}$.

Consider a total order on the λ -calculus variables; when writing a typing context as $x_1:A_1, \dots, x_n:A_n$ we now follow the convention that $x_1, \dots, x_n$ is an increasing sequence; we then define the semantics of any typing context by

$$\llbracket x_1:A_1, \dots, x_n:A_n \rrbracket := 1 \times \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket$$

The semantics of a typing derivation π for the typing judgement $\Gamma \vdash M:A$ is defined according to Fig. 3, by induction on π , as a morphism $\llbracket \pi \rrbracket: \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket$. ※

$$\begin{aligned} \left[\frac{}{x_1:A_1, \dots, x_n:A_n \vdash x_i:A_i} \right] &:= \pi_{i+1/n+1}: 1 \times \llbracket A_1 \rrbracket \times \dots \times \llbracket A_n \rrbracket \rightarrow \llbracket A_i \rrbracket \\ \left[\frac{\begin{array}{c} \vdots \pi \\ \Gamma, x:A \vdash M:B \end{array}}{\Gamma \vdash \lambda x.M:A \rightarrow B} \right] &:= \Lambda g: \llbracket \Gamma \rrbracket \rightarrow \llbracket B \rrbracket^{\llbracket A \rrbracket} \\ \text{where } g = \left[\frac{\begin{array}{c} \vdots \pi \\ \Gamma, x:A \vdash M:B \end{array}}{} \right] &: \llbracket \Gamma \rrbracket \times \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket \\ \left[\frac{\begin{array}{cc} \vdots \pi_1 & \vdots \pi_2 \\ \Gamma \vdash M:A \rightarrow B & \Gamma \vdash N:A \end{array}}{\Gamma \vdash M N:B} \right] &:= \langle g_1, g_2 \rangle \cdot \text{eval}: \llbracket \Gamma \rrbracket \rightarrow \llbracket B \rrbracket \\ \text{where } g_1 = \left[\frac{\begin{array}{c} \vdots \pi_1 \\ \Gamma \vdash M:A \rightarrow B \end{array}}{} \right] &: \llbracket \Gamma \rrbracket \rightarrow \llbracket B \rrbracket^{\llbracket A \rrbracket} \\ \text{and } g_2 = \left[\frac{\begin{array}{c} \vdots \pi_2 \\ \Gamma \vdash N:A \end{array}}{} \right] &: \llbracket \Gamma \rrbracket \rightarrow \llbracket A \rrbracket \end{aligned}$$

Figure 3: Semantics of the simply-typed λ -calculus in a CCC

Note that in the case of a λ -abstraction, we assume that x is (strictly) greater than any

variable in Γ . Any derivation in which this is not the case can easily be turned into one satisfying this condition: variables can always be renamed⁶ so that they are introduced in the typing context in increasing order.⁷

Now as mentioned earlier, we can relate the reductions in the simply-typed λ -calculus to the equality of morphisms in CCC:

THEOREM 4 (Soundness and completeness)

Assume $\frac{\vdots \pi}{\Gamma \vdash_{\text{st}\lambda} M : A}$ and $\frac{\vdots \pi'}{\Gamma \vdash_{\text{st}\lambda} N : A}$.
 $M \xrightarrow{\beta\eta}^* N$ if and only if in every CCC we have $\llbracket \pi \rrbracket = \llbracket \pi' \rrbracket$. ※

The equality theorems that can be derived from the axioms of CCC (and that therefore hold in every CCC) are reflected syntactically in the simply-typed λ -calculus, and the simply-typed λ -calculus is therefore said to form an *internal language* for CCC.

Note that it is easy to define the semantics of the **(I, K, S)**-combinatoric system (with simple types) in a CCC, for instance by encoding combinators as simply-typed λ -terms:

THEOREM 5 (Semantics of the (I, K, S)-combinatoric system)

The encoding of Fig. 4 satisfies the following properties:

- If $\vdash_{\text{stC}} M : A$ then $\vdash_{\text{st}\lambda} \overline{M} : A$, with a function $\pi \mapsto \overline{\pi}$ transforming a derivation of the former into a derivation of the latter.
- If $M \rightarrow M'$ then $\overline{M} \rightarrow_{\beta} \overline{M}'$.

The above properties allow the definition of the semantics $\llbracket \pi \rrbracket$ of a typing derivation π of $\vdash_{\text{stC}} M : A$ as the morphism $\llbracket \pi \rrbracket : 1 \rightarrow \llbracket A \rrbracket$, such that the following holds:

If $\pi \xrightarrow{*} \pi'$ then $\llbracket \pi \rrbracket = \llbracket \pi' \rrbracket$. ※

$$\begin{array}{ll} \overline{I} & := \lambda x.x \\ \overline{K} & := \lambda xy.x \\ \overline{S} & := \lambda xyz.x \ z \ (y \ z) \\ \overline{M \ N} & := \overline{M} \ \overline{N} \end{array}$$

Figure 4: **(I, K, S)**-combinators as λ -terms

1.1.4 Applying the methodology to other systems

The approach of the Curry-Howard correspondence can be, and has been, generalised with the following methodology:

- The first step is to decorate proofs with proof-terms: $\Gamma \vdash A$ becomes $\Gamma' \vdash M : A$, with Γ' being a typing context whose co-domain (i.e. the types which have been assigned to variables) is Γ ;

⁶Using equivariance of typing derivations.

⁷Alternative presentations of the simply-typed λ -calculus may be more convenient to define its semantics in a CCC: the use of De Bruijn indices (see e.g. [Bar84]) instead of named variables provides a natural way of ordering the objects in the interpretation of a typing environment (without resorting to ordering the set of variables); if variables carry their own type (or when each type comes with its own set of variables), the interpretation can be defined on the terms themselves rather than their typing derivations.

- the second is to express proof transformations in terms of proof-term reduction, denoted $M \longrightarrow_{\mathcal{S}} N$, often given by a rewrite system $\mathcal{S}$.

The desired properties of reduction are

- *Progress*, i.e. any term containing “undesirable structures” can be reduced.
- *Subject reduction* property, i.e. preservation of typing:
If $\Gamma \vdash M : A$ and $M \longrightarrow_{\mathcal{S}} N$ then $\Gamma \vdash N : A$
- possibly *Confluence*, programs are deterministic.
- possibly *Normalisation*, i.e. the fact that the execution of programs terminates.

The notion of “undesirable structures” is of course one of the concepts to identify in an interesting way; for instance in the simply-typed λ -calculus, a structure of the form $(\lambda x.M) N$ corresponds to the introduction of implication followed by the elimination of the introduced implication, a situation which we may consider undesirable from a proof-theoretic point of view.

To illustrate this methodology, we show how the correspondence from Section 1.1.2 can be extended to intuitionistic logic with both the implication connective and the logical constant $\perp$.

First note that by identifying $\perp$ simply as one of the atomic formulae, intuitionistic negation can be defined as follows: $\neg A := A \Rightarrow \perp$. With this definition, the following rules are instances of those of the simply-typed λ -calculus:

$$\frac{\Gamma, x : A \vdash M : \perp}{\Gamma \vdash \lambda x.M : \neg A} \quad \frac{\Gamma \vdash M : \neg A \quad \Gamma \vdash N : A}{\Gamma \vdash M N : \perp}$$

and these reflect the usual Natural Deduction rules for negation, but what is missing, to have intuitionistic logic, is the rule named *Ex falso quodlibet* (EFQ):

$$\frac{\Gamma \vdash \perp}{\Gamma \vdash A}$$

We now see what that rule may become in the Curry-Howard correspondence.

EXAMPLE 1 (Extension of the Curry-Howard correspondence for $\mathbf{NJ}_{\Rightarrow, \perp}$)

We extend the syntax of the λ -calculus with the following construct:

$$M, N, \dots ::= \dots \mid \text{abort}(M)$$

and we add to the simply-typed λ -calculus a typing rule corresponding to EFQ:

$$\frac{\Gamma \vdash M : \perp}{\Gamma \vdash \text{abort}(M) : A}$$

We may then add to the λ -calculus a rule such as

$$(b) \quad \text{abort}(M) N \longrightarrow \text{abort}(M)$$

to computationally interpret the new construct as a greedy consumer of arguments, and $\longrightarrow_b$, $\longrightarrow_{\beta b}$, $\longrightarrow_{\eta b}$, and $\longrightarrow_{\beta \eta b}$ are all confluent.

With these definitions, we still have Subject Reduction:

If $\Gamma \vdash M : A$ and $M \longrightarrow_{\beta \eta b} N$ then $\Gamma \vdash N : A$.

We can also interpret EFQ in category theory by requiring from a CCC the extra axiom that there is an *initial object* $\perp$, i.e. an object such that, from every object A there is a unique morphism $0_A : \perp \longrightarrow A$ (which is the dual of the terminal object 1 of the CCC). $\ast$

From there, it is natural to try to extend the above correspondence to classical logic, which can be obtained from intuitionistic logic by adding any one of the three axiom schemes:

$$\begin{array}{ll} \textit{Elimination of double negation (EDN):} & (\neg\neg A) \Rightarrow A \\ \textit{Peirce's law (PL):} & ((A \Rightarrow B) \Rightarrow A) \Rightarrow A \\ \textit{Law of excluded middle (LEM):} & A \vee \neg A \end{array}$$

In presence of EFQ (in short, the axiom scheme $\perp \Rightarrow A$), the above schemes are all equivalent in terms of formula provability. Interestingly enough and as noted in [AH03], without EFQ, we only have the following implications between the schemes:

$$\begin{array}{c} \text{EDN} \Rightarrow \text{PL} \Rightarrow \text{LEM} \\ \text{EDN} \Rightarrow \text{EFQ} \end{array}$$

Alternatives to adding axiom schemes is to add inference rules such as

$$\frac{\Gamma, \neg A \vdash \perp}{\Gamma \vdash A} \text{ for EDN}$$

or even to change the structure of the proof formalism, for instance by using the classical sequent calculus [Gen35] with *right-contraction*.

When thinking about classical logic, we have a tendency to identify a formula A with $\neg\neg A$, as suggested not only by the elimination of double negation but also by models of classical provability in boolean algebras.

Now, attempts to apply the Curry-Howard methodology to, say, the above axiom schemes or inference rule, are limited by the following fact:

A CCC with initial object $\perp$ and such that every object A is naturally isomorphic to $\perp^{\perp^A}$, collapses to a boolean algebra: there is at most 1 morphism between any 2 objects (see the proof in [LS86] or [Str11]).

That means that such a category would not distinguish two proofs of the same theorem, which is rather useless for a theory of proofs, or for the proofs-as-programs paradigm.

At that point, the natural question to ask is whether classical logic has computational content? To that question, and based on the above remarks, the book *Proofs and Types* [GTL89] answers in 1989:

“[The Curry-Howard] interpretation is not possible with classical logic: there is no sensible way of considering proofs as algorithms. In fact, classical logic has no denotational semantics, except the trivial one which identifies all the proofs of the same type.”

In the rest of this chapter we explore the alternative answers that have been given to the question since then.

1.2 Continuations and control

For this we start with some concepts which at first sight may seem unrelated: continuations and control.

A freeze-frame shot taken at one point of a program's execution flow could be represented, in a high-level view, as follows:

$$\begin{array}{l} \downarrow \text{ code } P \text{ that has been executed, producing data } v \\ v \text{ its output} \\ \downarrow \text{ code } E \text{ that remains to be executed, consuming data } v \end{array}$$

The code E that remains to be executed, and more generally the programming environment or programming context within which some code is executed, is called *continuation*.

The concept is also useful for compiling recursive calls: consider the following pseudo code

```
myfunction(a1,...,an){
  some code;
  x = myfunction(a1',...,an');
  some code possibly using x;
}
```

When executing the recursive call, the code that remains to be executed (i.e. `some code possibly using x`), together with the values of the local variables, needs to be stored in order to resume computation after the recursive call has returned with a value for `x`. But this is not needed in the case of *tail recursion*, in which `some code possibly using x` just returns (the value for) `x`.

The above code can be transformed into a tail-recursive code by modelling the remaining code `some code possibly using x` as a “continuation” function `c'` taking the value of `x` as input, and by *passing that continuation c'* as an extra argument to the recursive call, now in charge of executing it once the value of `x` is determined:

```
myfunction(a1,...,an,c){
  some code;
  return myfunction(a1',...,an',c');
}
```

The function `myfunction` now generally takes an extra argument `c` (the continuation function describing what to do with the result) and uses it in e.g. `some code` or in the construction of `c'`. Originally, the continuation is often the identity function (“do nothing with your result, just return it”), but it usually becomes more and more sophisticated with each recursive call.

Now we see what these concepts become in the case where the programming language is the λ -calculus. An instance of the above program execution flow picture

```

      ↓ code P that has been executed, producing data v
      v  its output
      ↓ code E that remains to be executed, consuming data v
```

can be seen by considering

- P to be a λ -term that is reduced,
- v to be the value to which P reduces,
- E to be the context, in the syntactic sense: a term with a hole $E[]$ (with the original λ -term being $E[P]$, i.e. the context $E[]$ whose hole has been filled with the λ -term P).

This is only a general idea: whether that view accurately describes program execution depends on the evaluation strategy for λ -terms; in particular, whether λ -terms are reduced inside-out, what notion of “value” is considered (is it a normal form?), what grammar for contexts $E[]$ ranges over, etc.

But in pure λ -calculus, it is clear that P has no knowledge of $E[]$ while being evaluated.

Control is about letting a program know and manipulate its evaluation context. Originally, the concept was used to model `goto` instructions, and other features that are not pure functional programming.

In the case of λ -calculus, the evaluation context $E[]$ within which a term is evaluated gives rise to a continuation function $\lambda x.E[x]$ (for a fresh variable x) that could be passed as

an argument.

Reynolds [Rey72], Strachey-Wadsworth [SW00] (re-edition of 74) explored continuations and control along those lines, letting a program capture its evaluation context with a feature known as *call-with-current-continuation* (call-cc): `cc`. This was added to the programming language **Scheme**.

Felleisen's PhD work [Fel87] on the Syntactic Theory of Control introduced another control operator: `C`.

The general idea of these control operators is given by the following reduction rules:

$$\begin{aligned} E[\text{cc } M] &\longrightarrow E[M (\lambda x. E[x])] \\ E[\text{C } M] &\longrightarrow M (\lambda x. E[x]) \end{aligned}$$

In presence of `abort( )` and its (slightly modified) rule

$$E[\text{abort}(M)] \longrightarrow M$$

`cc` and `C` are interdefinable:

$$\begin{aligned} \text{C } M &:= \text{cc } (\lambda k. \text{abort}(M k)) \quad k \notin \text{FV}(M) \\ \text{cc } M &:= \text{C } (\lambda k. k (M k)) \quad k \notin \text{FV}(M) \end{aligned}$$

Indeed,

$$\begin{aligned} E[\text{cc } M] &= E[\text{C } (\lambda k. k (M k))] \\ &\longrightarrow (\lambda k. k (M k)) (\lambda x. E[x]) \\ &\longrightarrow (\lambda x. E[x]) (M \lambda x. E[x]) \\ &\longrightarrow E[M (\lambda x. E[x])] \\ E[\text{C } M] &= E[\text{cc } (\lambda k. \text{abort}(M k))] \\ &\longrightarrow E[(\lambda k. \text{abort}(M k)) (\lambda x. E[x])] \\ &\longrightarrow E[\text{abort}(M \lambda x. E[x])] \\ &\longrightarrow M (\lambda x. E[x]) \end{aligned}$$

Of course, the above rules are not “standard” rewrite rules (clearly not first-order rewrite rules) and remain informal because we have not specified what $E[\]$ exactly stands for or ranges over.

More fundamentally, a central question about control is: what kind of continuation can be captured by a control operator and how? Is the capture delimited? undelimited? etc.

But what is interesting is that the above intuitions are sufficient to start seeing a connection with classical logic, as initiated by Griffin in [Gri90]:

$$\begin{aligned} \text{cc can be typed by PL: } & ((A \rightarrow B) \rightarrow A) \rightarrow A \\ \text{C can be typed by EDN: } & (\neg \neg A) \rightarrow A \end{aligned}$$

With these types, the rewrite rules satisfy the following Subject Reduction properties: The following (again, informal) typing tree for $E[\text{cc } M]$

$$\frac{\frac{\Gamma \vdash \lambda x. E[x] : A \rightarrow B}{\Gamma \vdash \text{cc } ((A \rightarrow B) \rightarrow A) \rightarrow A} \quad \frac{\Gamma \vdash M : (A \rightarrow B) \rightarrow A}{\Gamma \vdash \text{cc } M : A}}{\Gamma \vdash E[\text{cc } M] : B}$$

can be transformed into a typing tree for $E[M (\lambda x. E[x])]$

$$\frac{\frac{\Gamma \vdash \lambda x.E[x]: A \rightarrow B}{\Gamma \vdash \lambda x.E[x]: A \rightarrow B} \quad \frac{\Gamma \vdash M:(A \rightarrow B) \rightarrow A \quad \overline{\Gamma \vdash \lambda x.E[x]: A \rightarrow B}}{\Gamma \vdash M(\lambda x.E[x]): A}}{\Gamma \vdash E[M(\lambda x.E[x])]: B}$$

while the following (again, informal) typing tree for $E[\mathcal{C} M]$

$$\frac{\frac{\Gamma \vdash \lambda x.E[x]: A \rightarrow \perp}{\Gamma \vdash \lambda x.E[x]: A \rightarrow \perp} \quad \frac{\frac{\Gamma \vdash \mathcal{C}: ((A \rightarrow \perp) \rightarrow \perp) \rightarrow A \quad \Gamma \vdash M:(A \rightarrow \perp) \rightarrow \perp}{\Gamma \vdash \mathcal{C} M: A}}{\Gamma \vdash E[\mathcal{C} M]: \perp}$$

can be transformed into a typing tree for $E[M(\lambda x.E[x])]$

$$\frac{\frac{\Gamma \vdash M:(A \rightarrow \perp) \rightarrow \perp}{\Gamma \vdash M:(A \rightarrow \perp) \rightarrow \perp} \quad \frac{\Gamma \vdash \lambda x.E[x]: A \rightarrow \perp}{\Gamma \vdash \lambda x.E[x]: A \rightarrow \perp}}{\Gamma \vdash M(\lambda x.E[x]): \perp}$$

We can already see that, for the Subject Reduction property to hold in the case of $\mathcal{C}$, the context $E[\]$ cannot be any context: it has to produce something of type $\perp$. Similarly, for the generalised abort rule to satisfy Subject Reduction, the context $E[\]$ also needs to produce something of type $\perp$.

In fact, $\mathcal{C}$ generalises **abort**($_$), since we can define

$$\mathbf{abort}(M) := \mathcal{C}(\lambda x.M)$$

where x is a fresh (and therefore dummy) variable.

This reflects what we have already seen in pure logic: $\text{EDN} \Leftrightarrow (\text{PL} \wedge \text{EFQ})$

1.3 Contributions in the 90s

One possible formalisation of the above informal concepts was proposed by Parigot [Par92] in the form of the $\lambda\mu$ -calculus.

DEFINITION 10 ($\lambda\mu$ -calculus) The syntax of terms extends that of λ -calculus as follows:

$$\begin{array}{ll} \text{Terms} & M, N, P \dots ::= x \mid \lambda x.M \mid M N \mid \mu\alpha.c \\ \text{Commands} & c ::= [\alpha]M \end{array}$$

where α ranges over a new set of variables called *continuation variables*, and $\mu\alpha.c$ binds α in c . The scope of this binder, as well as that of the unique *command* construct $[\alpha]M$, extend as much as parentheses allow, so that $\mu\alpha.M N$ stands for $\mu\alpha.(M N)$ and $[\alpha]M N$ stands for $[\alpha](M N)$.

The typing rules extend those of λ -calculus as follows:

$$\begin{array}{c}
\overline{\Gamma, x:A \vdash x:A; \Delta} \\
\\
\frac{\Gamma, x:A \vdash M:B; \Delta}{\Gamma \vdash \lambda x.M:A \rightarrow B; \Delta} \quad \frac{\Gamma \vdash M:A \rightarrow B; \Delta \quad \Gamma \vdash N:A; \Delta}{\Gamma \vdash M N:B; \Delta} \\
\\
\frac{c:(\Gamma \vdash ; \alpha:A, \Delta)}{\Gamma \vdash \mu\alpha.c:A; \Delta} \quad \frac{\Gamma \vdash M:A; \alpha:A, \Delta}{[\alpha]M:(\Gamma \vdash ; \alpha:A, \Delta)}
\end{array}$$

where Γ is a typing context for term-variables and Δ is a typing context for continuation-variables. Derivability of sequents in this system is respectively denoted $\Gamma \vdash_{\lambda\mu} M:A ; \Delta$ and $c:(\Gamma \vdash_{\lambda\mu} ; \Delta)$.

The reduction rules extend the β -reduction of λ -calculus as follows:

$$\begin{array}{lcl}
(\lambda x.M) N & \longrightarrow & \{N/x\} M \\
(\mu\alpha.c) N & \longrightarrow & \mu\beta. \{[\beta]M N / [\alpha]M\} c \\
[\beta]\mu\alpha.c & \longrightarrow & \{\beta/\alpha\} c
\end{array}$$

where $\{[\beta]M N / [\alpha]M\} c$ is an unconventional substitution operation, consisting in replacing, in c , every subcommand (i.e. subterm that is a command) of the form $[\alpha]M$ by $[\beta]M N$, with the usual capture-avoiding conditions pertaining to substitution.

The rules define a reduction relation $\longrightarrow_{\lambda\mu}$ on both terms and commands. *

A basic intuition of the syntax is that each continuation variable α represents a “place” where various sub-terms of a given type (that of α) can be “stored” with a construct such as $[\alpha]M$. The construct $\mu\alpha.c$ retrieves what is stored under the continuation variable α and presents it as if it was a simple term. The second rewrite rule distributes for instance an argument to every sub-term stored under the variable α .

This calculus provides a computational interpretation of classical logic. Indeed, the typing system, when forgetting variables and terms, turns into the proof system of Fig. 5, where Γ and Δ now stand for sets or multisets of formulae. We see that the system generalises $\text{NJ}_{\Rightarrow}$, in particular with a more general form of sequent: $A_1, \dots, A_n \vdash A; B_1, \dots, B_m$, and a new form of sequent $A_1, \dots, A_n \vdash ; B_1, \dots, B_m$.

$$\begin{array}{c}
\overline{\Gamma, A \vdash A; \Delta} \\
\\
\frac{\Gamma, A \vdash B; \Delta}{\Gamma \vdash A \Rightarrow B; \Delta} \quad \frac{\Gamma \vdash A \Rightarrow B; \Delta \quad \Gamma \vdash A; \Delta}{\Gamma \vdash B; \Delta} \\
\\
\frac{\Gamma \vdash ; A, \Delta}{\Gamma \vdash A; \Delta} \quad \frac{\Gamma \vdash A; A, \Delta}{\Gamma \vdash ; A, \Delta}
\end{array}$$

Figure 5: The proof system corresponding to the simply-typed $\lambda\mu$ -calculus

Just like a sequent $A_1, \dots, A_n \vdash A$ of $\text{NJ}_{\Rightarrow}$ can be interpreted as the formula $(A_1 \wedge \dots \wedge A_n) \Rightarrow A$, the two sequent forms above can respectively be interpreted as the formulae $(A_1 \wedge \dots \wedge A_n) \Rightarrow (A \vee B_1 \vee \dots \vee B_m)$ and $(A_1 \wedge \dots \wedge A_n) \Rightarrow (B_1 \vee \dots \vee B_m)$. With this interpretation, the system of Fig. 5 can be easily checked to be sound with respect to classical logic, and for completeness we can see that

- the rules of $\text{NJ}_{\Rightarrow}$ are particular instances of the first three rules;
- the system features a right-contraction rule, which allows Peirce's Law to be proved, as we see below.

As with the simply-typed λ -calculus, the rewrite rules satisfy Subject Reduction, which allows the $\lambda\mu$ -calculus to describe a proof-transforming procedure for the system of Fig. 5:

THEOREM 6 (Subject reduction for the simply-typed $\lambda\mu$ -calculus)

1. If $\Gamma \vdash_{\lambda\mu} M : A; \Delta$ and $M \rightarrow_{\lambda\mu} M'$ then $\Gamma \vdash_{\lambda\mu} M' : A; \Delta$.
2. If $c : (\Gamma \vdash_{\lambda\mu} ; \Delta)$ and $c \rightarrow_{\lambda\mu} c'$ then $c' : (\Gamma \vdash_{\lambda\mu} ; \Delta)$.

✱

REMARK 7 This calculus integrates Peirce's law: By defining

$$\text{cc} := \lambda x. \mu \alpha. [\alpha](x \lambda y. \mu \beta. [\alpha]y)$$

we can build the following typing tree:

$$\frac{\frac{\frac{x : (A \rightarrow B) \rightarrow A, y : A \vdash y : A; \alpha : A, \beta : B}{[\alpha]y : (x : (A \rightarrow B) \rightarrow A, y : A \vdash ; \alpha : A, \beta : B)}}{x : (A \rightarrow B) \rightarrow A, y : A \vdash \mu \beta. [\alpha]y : B; \alpha : A}}{x : (A \rightarrow B) \rightarrow A \vdash \lambda y. \mu \beta. [\alpha]y : A \rightarrow B; \alpha : A} \quad \frac{x : (A \rightarrow B) \rightarrow A \vdash x : (A \rightarrow B) \rightarrow A; \alpha : A}{x : (A \rightarrow B) \rightarrow A \vdash x \lambda y. \mu \beta. [\alpha]y : A; \alpha : A} \quad \frac{[\alpha](x \lambda y. \mu \beta. [\alpha]y) : (x : (A \rightarrow B) \rightarrow A \vdash ; \alpha : A)}{x : (A \rightarrow B) \rightarrow A \vdash \mu \alpha. [\alpha](x \lambda y. \mu \beta. [\alpha]y) : A;}$$

$$\vdash \text{cc} : ((A \rightarrow B) \rightarrow A) \rightarrow A;$$

Now, consider that contexts are of the form $E[\] = [\gamma]([\] N_1 \dots N_n)$. We can perform the following reduction:

$$\begin{aligned} E[\text{cc } M] &= [\gamma](\lambda x. \mu \alpha. [\alpha](x \lambda y. \mu \beta. [\alpha]y)) M N_1 \dots N_n \\ &\rightarrow [\gamma](\mu \alpha. [\alpha](M \lambda y. \mu \beta. [\alpha]y)) N_1 \dots N_n \\ &\rightarrow [\gamma](\mu \alpha. [\alpha](M \lambda y. \mu \beta. [\alpha]y N_1) N_1) N_2 \dots N_n \\ &\rightarrow \dots \\ &\rightarrow [\gamma]\mu \alpha. [\alpha](M \lambda y. \mu \beta. [\alpha]y N_1 \dots N_n) N_1 \dots N_n \\ &\rightarrow [\gamma](M \lambda y. \mu \beta. [\gamma]y N_1 \dots N_n) N_1 \dots N_n \\ &= E[M (\lambda y. \mu \beta. E[y])] \end{aligned}$$

✱

Notice that what is passed to M as an argument is not exactly $\lambda y. E[y]$, since $E[\]$ forms a command and $\lambda y. E[y]$ is not correct syntax, but $\mu \beta. E[y]$ turns the command $E[y]$ into a term (of any type).

REMARK 8 If given a top-level continuation variable $\text{top} : \perp$ (Ariola-Herbelin [AH03, AH08]), then the $\lambda\mu$ -calculus integrates *Ex falso quodlibet* and the elimination of double negation:

We can build the following typing tree:

$$\frac{\frac{\frac{x:\perp \vdash x:\perp; \alpha:A}{[\text{top}]x:(x:\perp \vdash ; \alpha:A)}}{x:\perp \vdash \mu\alpha.[\text{top}]x:A;}}{\vdash \lambda x.\mu\alpha.[\text{top}]x:\perp \rightarrow A;}$$

and perform the following reduction:

$$\begin{aligned} E[(\lambda x.\mu\alpha.[\text{top}]x) M] &= [\gamma](\lambda x.\mu\alpha.[\text{top}]x) M N_1 \dots N_n \\ &\longrightarrow [\gamma](\mu\alpha.[\text{top}]M) N_1 \dots N_n \\ &\longrightarrow [\gamma](\mu\alpha.[\text{top}]M) N_2 \dots N_n \\ &\longrightarrow \dots \\ &\longrightarrow [\gamma](\mu\alpha.[\text{top}]M) \\ &\longrightarrow [\text{top}]M \end{aligned}$$

And by defining

$$\mathcal{C} := \lambda x.\mu\alpha.[\text{top}](x \lambda y.\mu\beta.[\alpha]y)$$

we can build the following typing tree:

$$\frac{\frac{\frac{\frac{x:\neg\neg A, y:A \vdash y:A; \alpha:A, \beta:\perp}{[\alpha]y:(x:\neg\neg A, y:A \vdash ; \alpha:A, \beta:\perp)}}{x:\neg\neg A, y:A \vdash \mu\beta.[\alpha]y:\perp; \alpha:A}}{x:\neg\neg A \vdash x:\neg\neg A; \alpha:A} \quad \frac{x:\neg\neg A \vdash \lambda y.\mu\beta.[\alpha]y:\neg A; \alpha:A}{x:\neg\neg A \vdash x \lambda y.\mu\beta.[\alpha]y:\perp; \alpha:A}}{\frac{[\text{top}](x \lambda y.\mu\beta.[\alpha]y):(x:\neg\neg A \vdash ; \alpha:A)}{x:\neg\neg A \vdash \mu\alpha.[\text{top}](x \lambda y.\mu\beta.[\alpha]y):A;}}{\vdash \lambda x.\mu\alpha.[\text{top}](x \lambda y.\mu\beta.[\alpha]y):(\neg\neg A) \rightarrow A;}$$

and we can perform the following reduction:

$$\begin{aligned} E[\mathcal{C} M] &= [\gamma](\lambda x.\mu\alpha.[\text{top}](x \lambda y.\mu\beta.[\alpha]y)) M N_1 \dots N_n \\ &\longrightarrow [\gamma](\mu\alpha.[\text{top}](M \lambda y.\mu\beta.[\alpha]y)) N_1 \dots N_n \\ &\longrightarrow [\gamma](\mu\alpha.[\text{top}](M \lambda y.\mu\beta.[\alpha]y N_1)) N_2 \dots N_n \\ &\longrightarrow \dots \\ &\longrightarrow [\gamma]\mu\alpha.[\text{top}]M \lambda y.\mu\beta.[\alpha]y N_1 \dots N_n \\ &\longrightarrow [\text{top}]M \lambda y.\mu\beta.[\gamma]y N_1 \dots N_n \\ &= [\text{top}]M (\lambda y.\mu\beta.E[y]) \end{aligned}$$

※

Notice that what is eventually produced by the rewrites is not M and $M (\lambda y.\mu\beta.E[y])$, respectively, but $[\text{top}]M$ and $[\text{top}]M (\lambda y.\mu\beta.E[y])$, since reducing a command has to produce a command. But since M is of type $\perp$ (respectively produces a term of type $\perp$), it can be stored in the top-level continuation top .

Now, when thinking about classical logic, we often have in mind concepts of symmetry or duality:

Inversing the order in a boolean algebra provides another boolean algebra where e.g. the

top and bottom elements have been swapped, the meet and join operations have been swapped.

Very related to this are De Morgan’s rules, which show a duality, via negation, between $\wedge$ and $\vee$:

$$\begin{aligned}\neg(A \wedge B) &= \neg A \vee \neg B \\ \neg(A \vee B) &= \neg A \wedge \neg B\end{aligned}$$

In terms of proof formalisms, the classical sequent calculus LK [Gen35] shows a symmetry between the left-hand side and the right-hand of sequents, of the form $A_1, \dots, A_n \vdash B_1, \dots, B_m$: whatever can be done on the left-hand side can be done on the right-hand side, and vice versa. For instance, the left-introduction rule for $\wedge$ is symmetric to the right-introduction rule for $\vee$ and vice versa; left-contraction symmetric to right-contraction (and this is very different from intuitionistic logic).

But so far, such symmetries and dualities are not explicitly reflected in our proof-term approach to classical proofs.

However, before even Griffin made the connection between control operators and classical logic, Filinski [Fil89] formalised a duality between

- functions as values
- functions as continuations

in the form of a “symmetric λ -calculus”, with explicit conversions from one view of functions to the other. Yet there was no explicit connection with classical logic.

In [BB96], Barbanera and Berardi formalised their own symmetric λ -calculus, with a typing system providing a Curry-Howard interpretation of classical proofs. The classical proof system depicted by their calculus is a one-sided version of the classical sequent calculus [Gen35], with a proof of normalisation for typed terms (we will see such a proof in Chapter 2).

Since then, two calculi emerged to provide Curry-Howard interpretations of the two-sided sequent calculus LK (or variants thereof), with the reduction rules describing the famous proof-transformation procedure known as *cut-elimination*:

- Urban’s calculus [Urb00],
- Curien and Herbelin’s $\bar{\lambda}\mu\tilde{\mu}$ [CH00] for $\Rightarrow$,
later extended by Wadler [Wad03] for $\wedge$ and $\vee$ (explicitly connecting the symmetries of the calculus to De Morgan’s duality).

These two independent (sets of) contributions had different aims: Curien and Herbelin’s was to expose, as the syntactic symmetry of the classical sequent calculus, a *duality* in computation based on Filinski’s ideas about continuations and on the call-by-value and call-by-name evaluation strategies; they gave semantics to their calculus, but with no proof of normalisation. Urban’s aim was to have a typing system as close as possible to LK and have a reduction system as close as possible to basic cut-elimination procedures; his Ph.D. adapted Barbanera and Berardi’s proof of strong normalisation to his calculus, but gave no (denotational) semantics.

Several papers formalise the links between the various proof calculi for classical logic: in particular, [Len03] relates $\bar{\lambda}\mu\tilde{\mu}$ and Urban’s calculus, [Roc05] relates $\bar{\lambda}\mu\tilde{\mu}$ and $\lambda\mu$, and [Her05] presents an extensive exploration of the relations between the various calculi.

In the rest of this chapter, we focus on Curien and Herbelin’s calculus to explore some more semantical concepts, but many of them can be transposed to other calculi for classical logic (in particular, Parigot’s $\lambda\mu$ -calculus).

1.4 System L

System L is the new name of Curien and Herbelin's $\bar{\lambda}\mu\tilde{\mu}$, extended with other connectives. We start with its syntax:

DEFINITION 11 (Syntax) The syntax of L is made of three syntactic categories:

terms	$t ::= x \mid \mu\beta.c \mid \lambda x.t \mid (t_1, t_2) \mid \text{inj}_i(t)$
continuations	$e ::= \alpha \mid \mu x.c \mid t :: e \mid (e_1, e_2) \mid \text{inj}_i(e)$
commands	$c ::= \langle t \mid e \rangle$

where i ranges over $\{1, 2\}$, x and α respectively range over term variables and continuation variables,⁸ $\mu\beta.c$ binds β in c , $\mu x.c$ binds x in c , and $\lambda x.t$ binds x in t . The scope of binders extends as much as parentheses allow. *

A (somewhat shallow) intuition of the syntax can be given as follows:

Term variables $x, y, \dots$	denote inputs
Continuation variables $\alpha, \beta, \dots$	denote outputs
A <i>term</i> has	one main output
	some inputs (free term variables)
	some “alternative” outputs (free continuation variables)
A <i>continuation</i> has	one main input
	some “additional” inputs (free term variables)
	some possible outputs (free continuation variables)
A <i>command</i> is	a term facing a continuation (the interaction is computation)

DEFINITION 12 (Typing)

We consider the following grammar for *types* (extending that of simple types):

$$A, B, \dots ::= a \mid A \rightarrow B \mid A \wedge B \mid A \vee B$$

where a ranges over a fixed set of elements called *atomic types*. The symbol $\rightarrow$ is associative to the right, i.e. $A \rightarrow (B \rightarrow C)$ can be abbreviated as $A \rightarrow B \rightarrow C$.

The typing system for System L is given for three kinds of sequents corresponding to the three syntactic categories of the syntax:

$$\Gamma \vdash t : A ; \Delta \qquad \Gamma ; e : A \vdash \Delta \qquad c : (\Gamma \vdash \Delta)$$

where Γ is a typing context for term-variables and Δ is a typing context for continuation-variables.

The system is presented in Fig. 6. Derivability of sequents in this system is respectively denoted $\Gamma \vdash_L t : A ; \Delta$, $\Gamma ; e : A \vdash_L \Delta$, and $c : (\Gamma \vdash_L \Delta)$. *

As we can see, forgetting about variables and proof-terms does not give the sequent calculus LK exactly as we know it from [Gen35] or as the popular variants described in [TS00] (for this one can look at Urban's Ph.D. [Urb00]), if only because there are three types of sequents. However, it is a variant with a bit more structure, which defines the same notion of provability as LK, and which will prove useful for the computational interpretation of classical logic.

An intuition about this interpretation can be given as follows: similarly to the Curry-Howard correspondence in intuitionistic logic, each connective in the syntax of formulae corresponds to a type construct in programming; term constructs offer basic ways in which such

⁸As in Parigot's $\lambda\mu$ -calculus [Par92].

$$\begin{array}{c}
\frac{}{\Gamma, x:A \vdash x:A; \Delta} \qquad \frac{}{\Gamma; \alpha:A \vdash \alpha:A, \Delta} \\
\\
\frac{\Gamma, x:A \vdash t:B; \Delta}{\Gamma \vdash \lambda x.t:A \rightarrow B; \Delta} \qquad \frac{\Gamma \vdash t:A; \Delta \quad \Gamma; e:B \vdash \Delta}{\Gamma; t::e:A \rightarrow B \vdash \Delta} \\
\\
\frac{\Gamma \vdash t_1:A_1; \Delta \quad \Gamma \vdash t_2:A_2; \Delta}{\Gamma \vdash (t_1, t_2):A_1 \wedge A_2; \Delta} \qquad \frac{\Gamma; e:A_i \vdash \Delta}{\Gamma; \text{inj}_i(e):A_1 \wedge A_2 \vdash \Delta} \\
\\
\frac{\Gamma \vdash t:A_i; \Delta}{\Gamma \vdash \text{inj}_i(t):A_1 \vee A_2; \Delta} \qquad \frac{\Gamma; e_1:A_1 \vdash \Delta \quad \Gamma; e_2:A_2 \vdash \Delta}{\Gamma; (e_1, e_2):A_1 \vee A_2 \vdash \Delta} \\
\\
\frac{c:(\Gamma \vdash \alpha:A, \Delta)}{\Gamma \vdash \mu \alpha.c:A; \Delta} \qquad \frac{c:(\Gamma, x:A \vdash \Delta)}{\Gamma; \mu x.c:A \vdash \Delta} \\
\\
\frac{\Gamma \vdash t:A; \Delta \quad \Gamma; e:A \vdash \Delta}{\langle t \mid e \rangle : (\Gamma \vdash \Delta)}
\end{array}$$

Figure 6: Typing system for L

types can be inhabited, while continuation constructs offer basic ways in which inhabitants of such types are consumed:

- A conjunction $A_1 \wedge A_2$ corresponds to a product type, so basic inhabitants are pairs (t_1, t_2) of terms (with the first component inhabiting A_1 and the second inhabiting A_2); basic continuations that consume such a pair start by extracting either the first or the second component (in other words, they start with one of the two projections), which corresponds to the continuation constructs $\text{inj}_1(e)$ and $\text{inj}_2(e)$.
- A disjunction $A_1 \vee A_2$ corresponds to a sum type, so basic inhabitants are the injections $\text{inj}_1(t)$ and $\text{inj}_2(t)$ (with t inhabiting A_1 or inhabiting A_2 , respectively); basic continuations that consume such an injection must handle both cases, so the case analysis leads to providing a pair $\langle e_1, e_2 \rangle$ of two continuations: the former can consume inhabitants of A_1 and the latter can consume inhabitants of A_2 .
- An implication $A_1 \Rightarrow A_2$ corresponds to a function type, with the basic inhabitants being constructed with λ -abstractions just like in the λ -calculus; we do not have the construct that directly applies a function to an argument, but a basic way in which a continuation consumes a function is to offer an argument t as the input of the function, together with a continuation e that can consume the output of the function; hence the continuation construct $t::e$ (which is simply the usual stacking construct that can be found in abstract machines to implement computation in the λ -calculus).

This intuition will be strengthened by the reduction rules for System L, but we first start with an example.

The following story is borrowed from Phil Wadler [Wad03] (who might have borrowed it from Peter Selinger), and illustrates the computational contents of classical proofs:

EXAMPLE 2 (The devil, the fool, and the \$1,000,000)

The Devil meets a man and says:

“- I have an offer for you! I promise you that

either I offer you \$1,000,000 or, if you give me \$1,000,000, then I will grant you any wish.

Actually, I choose to offer you the latter.”

The man then goes back home and, motivated by the Devil’s promise, strives to gather \$1,000,000. Ten years later, he finally succeeds; he goes back to the Devil and, handing him the money, says:

“- Here’s \$1,000,000! I want immortality.”

The Devil takes the money and says:

“- Well done and thank you!

Actually, I’ve changed my mind. I’ve now decided to fulfil my promise by offering you \$1,000,000. Here is your money back!” ※

The reason why that short story illustrates the computational contents of classical logic is that the Devil behaves as a proof of the Law of Excluded Middle: Imagine that

- the money (\$1,000,000) can be seen as an atomic proposition a
- the part of the promise “If you give me \$1,000,000, I’ll grant you any wish” can be seen as the formula $a \Rightarrow \perp$, i.e. $\neg a$;

the Devil is then the proof of $a \vee \neg a$ shown in Fig. 7. Indeed, following the bottom-up

$$\begin{array}{c}
 \hline
 y:a \vdash y:a ; \alpha:a\vee\neg a, \beta:\perp \\
 \hline
 y:a \vdash \text{inj}_1(y):a\vee\neg a ; \alpha:a\vee\neg a, \beta:\perp \quad y:a ; \alpha:a\vee\neg a \vdash \alpha:a\vee\neg a, \beta:\perp \\
 \hline
 \frac{\langle \text{inj}_1(y) \mid \alpha \rangle : (y:a \vdash \alpha:a\vee\neg a, \beta:\perp)}{y:a \vdash \mu\beta. \langle \text{inj}_1(y) \mid \alpha \rangle : \perp ; \alpha:a\vee\neg a} \\
 \hline
 \frac{\vdash \lambda y. \mu\beta. \langle \text{inj}_1(y) \mid \alpha \rangle : \neg a ; \alpha:a\vee\neg a}{\vdash \text{inj}_2(\lambda y. \mu\beta. \langle \text{inj}_1(y) \mid \alpha \rangle) : a\vee\neg a ; \alpha:a\vee\neg a} \quad \frac{}{; \alpha:a\vee\neg a \vdash \alpha:a\vee\neg a} \\
 \hline
 \frac{\langle \text{inj}_2(\lambda y. \mu\beta. \langle \text{inj}_1(y) \mid \alpha \rangle) \mid \alpha \rangle : (\vdash \alpha:a\vee\neg a)}{\vdash \mu\alpha. \langle \text{inj}_2(\lambda y. \mu\beta. \langle \text{inj}_1(y) \mid \alpha \rangle) \mid \alpha \rangle : a\vee\neg a ;} \\
 \hline
 \end{array}$$

Figure 7: A proof of LEM

construction of the left-hand branch, we see that

- the proof (the Devil) starts by choosing to prove $\neg a$, as reflected by the $\text{inj}_2(_)$ construct;
- that requires an input of type a (the \$1,000,000 earned by the fool), namely y , as reflected by the $\lambda y._$ construct;
- given the impossibility to prove $\perp$ directly (the immortality wish, or for that matter, any wish), the proof re-attacks the original formula to prove, namely $a\vee\neg a$ (the Devil returns to his original promise), but this time with the input $y:A$ (the \$1,000,000 that the fool gave him);
- this time, the proof chooses to prove a , which is trivially done by returning y (the Devil chooses to give \$1,000,000, by returning the money that the man earned).

We see here that the proof works because of the possibility to construct an inhabitant of

$a \vee \neg a$, twice along the same branch (we inhabit it the first time with the second injection, then with the first one), which is technically allowed by the *right-contraction* implicitly featured in the bottom two steps of the proof. While in intuitionistic logic it is possible to contract on the left but not contract on the right, classical logic allows both symmetrically.

This allows to also build a proof-term of type PL and, allowing again (as in Parigot's $\lambda\mu$) a top-level continuation variable **top** of type $\perp$, we can build proof-terms for *Ex falso quodlibet* and the elimination of double negation.

In summary, we have seen that it is easy enough to introduce proof-terms to represent classical proofs, such that the symmetry of classical logic reflects the symmetry between programs and continuations.

The use of classical reasoning corresponds to the use of control features allowing programs to capture their continuation, as we now see by looking at reductions:

DEFINITION 13 (Reductions) The reductions are given by the following rewrite system:

$$\begin{array}{lll}
 (\rightarrow) & \langle \lambda x. t_1 \mid t_2 :: e \rangle & \longrightarrow \langle t_2 \mid \mu x. \langle t_1 \mid e \rangle \rangle \\
 (\wedge) & \langle (t_1, t_2) \mid \text{inj}_i(e) \rangle & \longrightarrow \langle t_i \mid e \rangle \\
 (\vee) & \langle \text{inj}_i(t) \mid (e_1, e_2) \rangle & \longrightarrow \langle t \mid e_i \rangle \\
 (\overleftarrow{\mu}) & \langle \mu \beta. c \mid e \rangle & \longrightarrow \{e/\beta\} c \\
 (\overrightarrow{\mu}) & \langle t \mid \mu x. c \rangle & \longrightarrow \{t/x\} c
 \end{array}$$

※

Now, while it was very clear that Parigot's $\lambda\mu$ forms an extension of λ -calculus, we should emphasise the fact that the λ -calculus can be encoded in System L:

DEFINITION 14 (Encoding of λ -calculus)

We encode λ -terms as terms of System L by first encoding values, then all terms:

$$\begin{array}{ll}
 \overline{x}^V & := x \\
 \overline{\lambda x. M}^V & := \lambda x. \overline{M}
 \end{array}
 \qquad
 \begin{array}{l}
 \overline{V M_1 \dots M_n} := \mu \alpha. \langle \overline{V}^V \mid \overline{M_1} :: \dots \overline{M_n} :: \alpha \rangle \\
 \text{where } V \text{ is not an application and } n \geq 0
 \end{array}$$

※

LEMMA 9 (Simulation of λ -calculus)

1. $\mu \alpha. \langle \overline{M} \mid \overline{M_1} :: \dots \overline{M_n} :: \alpha \rangle \longrightarrow \overline{M M_1 \dots M_n}$.
2. $\{ \overline{N} / x \} \overline{M} \longrightarrow^* \{ \overline{N} / x \} \overline{M}$.
3. If $M \longrightarrow_\beta N$ then $\overline{M} \longrightarrow^* \overline{N}$.

※

Proof: The first point is a simple $\overleftarrow{\mu}$ -reduction, the second point is by induction on M , the third point is by induction on the rewrite derivation. $\square$

LEMMA 10 (Preservation of simple types)

1. If $\Gamma \vdash_{\text{st}\lambda} V : A$ then $\Gamma \vdash_{\text{L}} \overline{M}^V : A$;
2. If $\Gamma \vdash_{\text{st}\lambda} M : A$ then $\Gamma \vdash_{\text{L}} \overline{M} : A$;

※

Since System L contains cuts, a proof of LEM, and it can encode simply-typed λ -terms, it is clearly complete for classical logic (in the same sense as for the $\lambda\mu$ -calculus: EDN and EFQ require the presence of a top-level continuation variable **top**: $\perp$). Soundness can be trivially

checked by checking that all the inference rules are sound when forgetting about variables and proof-terms.

Substitution behaves well with respect to typing:

THEOREM 11 (Substitution Lemma)

1. If $c : (\Gamma, x : A \vdash_{\mathbf{L}} \Delta)$ and $\Gamma \vdash_{\mathbf{L}} t : A ; \Delta$ then $\{t/x\}c : (\Gamma \vdash_{\mathbf{L}} \Delta)$.
2. If $c : (\Gamma \vdash_{\mathbf{L}} \alpha : A, \Delta)$ and $\Gamma ; e : A \vdash_{\mathbf{L}} \Delta$ then $\{e/\alpha\}c : (\Gamma \vdash_{\mathbf{L}} \Delta)$.

※

Proof: By induction on c , simultaneously proving the two analogous properties for both terms and continuations. $\square$

And the reduction relation satisfies Subject Reduction:

THEOREM 12 (Subject reduction for System L)

1. If $c : (\Gamma \vdash_{\mathbf{L}} \Delta)$ and $c \longrightarrow c'$ then $c' : (\Gamma \vdash_{\mathbf{L}} \Delta)$.
2. If $\Gamma \vdash_{\mathbf{L}} t : A ; \Delta$ and $t \longrightarrow t'$ then $\Gamma \vdash_{\mathbf{L}} t' : A ; \Delta$.
3. If $\Gamma ; e : A \vdash_{\mathbf{L}} \Delta$ and $e \longrightarrow e'$ then $\Gamma ; e' : A \vdash_{\mathbf{L}} \Delta$.

※

Proof: Straightforward induction on the rewrite derivations. $\square$

Again, Subject Reduction allows the rewrite system to describe a proof transformation procedure in the classical sequent calculus, and in this case it is *cut-elimination* [Gen35].

Let us see the other properties we mentioned when introducing the Curry-Howard methodology:

Progress depends of course on what we consider an “undesirable structure”. In the case of sequent calculus, the natural concept of undesirable structure is the cut, which in the typing system of $\mathbf{L}$ is (at least at first sight) represented as the bottom-most rule of Fig. 6. And at this point we notice that some cuts cannot be reduced, as no rewrite rule applies to their proof-terms, namely those of the form $\langle x \mid e \rangle$ and $\langle t \mid \alpha \rangle$ when e is not of the form $\mu x.c$ and t is not of the form $\mu \alpha.c$. We may think progress fails (in terms of cut-elimination), but we should also notice that cuts of that form are very peculiar: they do nothing but respectively implement a left-contraction or a right-contraction, two rules that the extra structure of the system requires for completeness (compared to e.g. G3ii [TS00]):

$$\frac{\Gamma, x : A \vdash x : A ; \Delta \quad \Gamma, x : A ; e : A \vdash \Delta}{\langle x \mid e \rangle : (\Gamma, x : A \vdash \Delta)} \quad \frac{\Gamma \vdash t : A ; \alpha : A, \Delta \quad \Gamma ; \alpha : A \vdash \alpha : A, \Delta}{\langle t \mid \alpha \rangle : (\Gamma \vdash \alpha : A, \Delta)}$$

We actually used two of these special “cuts” in the proof of LEM showed in Fig. 7, and we would not expect to eliminate them (unless we had specific constructs for contractions and for the axiom represented as $\langle x \mid \alpha \rangle$).

Concerning normalisation, it can be proved that typed commands (resp. terms, continuations) are strongly normalising. This was inferred from Urban’s calculus in [Len03], but can be more simply obtained as the direct application of Barbanera and Berardi’s technique, as shown in [Pol04] for a variant of System $\mathbf{L}$ with explicit substitutions. This will be the topic of Chapter 2.

Finally, we look at the confluence property.

1.5 Non-confluence of cut-elimination in classical logic

As the reduction relation of System L specifies a cut-elimination procedure, we should note that cut-elimination in classical logic, at a purely logical level, can easily be defined as a non-confluent transformation procedure. A typical example of this is Lafont’s example, which we first express in the original sequent calculus LK, with explicit rules for weakenings and contractions and “context-splitting” rules (see e.g. [TS00]):

EXAMPLE 3 (Lafont’s example for non-confluence)

Consider the following cut that we would like to eliminate

$$\frac{\frac{\frac{\vdots \pi}{\Gamma \vdash \Delta}}{\Gamma \vdash \Delta, A} \quad \frac{\frac{\vdots \pi'}{\Gamma' \vdash \Delta'}}{\Gamma', A \vdash \Delta'}}{\Gamma, \Gamma' \vdash \Delta, \Delta'}$$

There are two ways to eliminate the cut:

$$\frac{\frac{\vdots \pi}{\Gamma \vdash \Delta}}{\Gamma, \Gamma' \vdash \Delta, \Delta'} \quad \text{or} \quad \frac{\frac{\vdots \pi'}{\Gamma' \vdash \Delta'}}{\Gamma, \Gamma' \vdash \Delta, \Delta'}$$

※

This obviously leads to non-confluence as soon as π and π' are two distinct proofs (say, cut-free). Note that we could, somewhat artificially, avoid the choice between π and π' by considering the following *mix* rule [FR94]:

$$\frac{\Gamma \vdash \Delta \quad \Gamma' \vdash \Delta'}{\Gamma, \Gamma' \vdash \Delta, \Delta'}$$

which would allow the symmetric combination of π and π' into a single proof (and what would be the semantics of this combination?). The question is whether we accept this derivation as a normal proof. Let us look at the same example in a sequent calculus (such as G3ii [TS00]) where rules are “context-sharing”:

EXAMPLE 4 (Lafont’s example in a context-sharing sequent calculus)

The following cut:

$$\frac{\frac{\frac{\vdots \pi}{\Gamma \vdash \Delta}}{\Gamma \vdash \Delta, A} \quad \frac{\frac{\vdots \pi'}{\Gamma \vdash \Delta}}{\Gamma, A \vdash \Delta}}{\Gamma \vdash \Delta}$$

can be reduced to:

$$\frac{\vdots \pi}{\Gamma \vdash \Delta} \quad \text{or} \quad \frac{\vdots \pi'}{\Gamma \vdash \Delta}$$

※

What is even more striking in this example is that π and π' are two proofs of the same sequent, which we probably do not want to consider denotationally equal and whose combination via the following rule

$$\frac{\Gamma \vdash \Delta \quad \Gamma \vdash \Delta}{\Gamma \vdash \Delta}$$

looks even more artificial than with the context-splitting mix. Again, do we want this derivation as a normal proof?

Unsatisfying though the mix may seem, it does technically solve Lafont's non-confluence problem based on two weakenings. Unfortunately, it cannot solve the even more problematic example obtained with contractions instead of weakenings:

EXAMPLE 5 (Example with contractions) The following cut

$$\frac{\frac{\frac{\vdots \pi}{\Gamma \vdash \Delta, A, A} \quad \frac{\vdots \pi'}{\Gamma, A, A \vdash \Delta}}{\Gamma \vdash \Delta, A} \quad \frac{\vdots \pi'}{\Gamma, A \vdash \Delta}}{\Gamma \vdash \Delta}$$

can be reduced to

$$\begin{array}{ccc} \begin{array}{c} \vdots \pi' \\ \bullet \\ \vdots \pi(\simeq) \\ \Gamma \vdash \Delta \end{array} & \text{or} & \begin{array}{c} \vdots \pi \\ \bullet \\ \vdots \pi'(\simeq) \\ \Gamma \vdash \Delta \end{array} \end{array}$$

where $\pi(\simeq)$ (resp. $\pi'(\simeq)$) denotes the proof π (resp. π') modified by the propagation of π' (resp. π) into its structure.

We can give a concrete instance of the above:

$$\frac{(A \rightarrow B) \rightarrow A \vdash A \quad A, A \rightarrow C, A \rightarrow D \vdash C \wedge D}{(A \rightarrow B) \rightarrow A, A \rightarrow C, A \rightarrow D \vdash C \wedge D}$$

Peirce's Law requires a right-contraction on the cut-formula A while the right-hand side proof requires a left-contraction on the cut-formula A . ※

Coming back to the proof-term side, both examples would appear in System L as instances of the following scheme:

$$\frac{c: (\Gamma \vdash \alpha: A, \Delta) \quad c': (\Gamma, x: A \vdash \Delta)}{\Gamma \vdash \mu\alpha.c: A; \Delta \quad \Gamma; \mu x.c': A \vdash \Delta} \frac{}{\langle \mu\alpha.c \mid \mu x.c' \rangle: (\Gamma \vdash \Delta)}$$

α (resp. x) could be used 0 (weakening), 1, or several (contraction) times in c (resp. c').

That cut could be reduced to

$$\frac{c: (\Gamma \vdash \alpha: A, \Delta) \quad \frac{c': (\Gamma, x: A \vdash \Delta)}{\Gamma; \mu x.c': A \vdash \Delta}}{\dots \left\{ \mu x.c' / \alpha \right\} c: (\Gamma \vdash \Delta)} \quad \text{or} \quad \frac{c: (\Gamma \vdash \alpha: A, \Delta) \quad c': (\Gamma, x: A \vdash \Delta)}{\dots \left\{ \mu\alpha.c / x \right\} c': (\Gamma \vdash \Delta)}$$

where dotted lines do not represent primitive inference rules, but inference rules that have been shown admissible in the typing system (Lemma 11).

In the case of weakening, and reflecting Example 4, $\alpha \notin \text{FV}(c)$ and $x \notin \text{FV}(c')$ and we can reduce $\langle \mu\alpha.c \mid \mu x.c' \rangle$ to two arbitrary commands c and c' with the same type.

This makes it hard to give a denotational semantics of classical proofs or of typed proof-terms: if we require $\langle \mu\alpha.c \mid \mu x.c' \rangle$, $\left\{ \mu x.c' / \alpha \right\} c$, and $\left\{ \mu\alpha.c / x \right\} c'$, to have the same denotation,

Example 4 leads to giving the same denotation to every proof of the same sequent.

This of course relates to the fact that we have already mentioned: a CCC with initial object where every object A naturally isomorphic to $\neg\neg A$ collapses to a boolean algebra. As identified in [Str11], there are three natural ways to resolve this:

1. Break the symmetry between $\wedge$ and $\vee$
2. Break the cartesian product (as studied for instance in [FP06, LS05, Lam07, Str11])
3. Break curryfication (as studied for instance in [DP04, CS09])

In this dissertation, we break the symmetry between $\wedge$ and $\vee$, since out of the three solutions it is the one for which the Curry-Howard correspondence with programs is best understood.

One way of breaking the non-confluence problem

$$\frac{\begin{array}{c} \vdots \pi \\ \Gamma \vdash \Delta, A \end{array} \quad \begin{array}{c} \vdots \pi' \\ \Gamma, A \vdash \Delta \end{array}}{\Gamma \vdash \Delta}$$

is simply to give systematic priority to

- the right (push π into π')
- or to the left (push π' into π)

Almost by definition, both solutions make the calculus confluent.

They also break the $\wedge\vee$ symmetry: Giving systematic priority to the right, say, makes a term t of type $A\wedge B$ have the same behaviour as $(\text{inj}_1(t), \text{inj}_2(t))$, whereas a continuation e of type $A\vee B$ will not necessarily have the same behaviour as $(\text{inj}_1(e), \text{inj}_2(e))$.

More details on this will be given in Chapter 3, but we shall also see by semantical means that the $\wedge\vee$ symmetry is broken. The two reduction strategies suggest to construct two denotational semantics $\llbracket c \rrbracket_N$ and $\llbracket c \rrbracket_V$ with the hope that:

$$\begin{aligned} \llbracket c_0 \rrbracket_N = \llbracket c_1 \rrbracket_N & \quad \text{iff} \quad "c_0 \longleftrightarrow^* c_1 \text{ with systematic priority to the right}" \\ \llbracket c_0 \rrbracket_V = \llbracket c_1 \rrbracket_V & \quad \text{iff} \quad "c_0 \longleftrightarrow^* c_1 \text{ with systematic priority to the left}" \end{aligned}$$

The use of the letters N and V reflects the fact that the strategies relate to the notions of Call-by-name and Call-by-value, as investigated for instance by Plotkin [Plo75], Moggi [Mog89], and others.

In conclusion of this section, we have seen that it is easy enough to give a rewrite system on proof-terms to represent cut-elimination (and the system follows the intuitions of continuations and control), but it gives a non-confluent calculus because cut-elimination is non-confluent in classical logic (via the Curry-Howard correspondence, because programs and continuations fight for the control of computation).

The rest of this chapter is devoted to the construction of the above CBN and CBV semantics.

1.6 Continuations, Call-by-Name and Call-by-Value

Call-by-name and call-by-value are two strategies for evaluating programs. Imagine the definition of a function (in pseudo-code):

```
MyFavoriteFunction(x){
  ... x ...
}
```

and later a call to that function with argument A :

```
MyFavoriteFunction(A)
```

The main question is whether A should be evaluated before entering the (code of the) function (CBV) or when it is actually used (CBN)? This is a question of interpretation or compilation of programs and, especially in presence of side-effects, knowing which of the two the compiler implements, is vital for the determinism of evaluation.

In general, we call *values* what evaluation should produce (e.g. booleans `true`, `false`). In functional programming, functions are particular values and can be passed as arguments. In general, functions are therefore not reduced.

The λ -calculus is both a core functional language and a theory of functions.

As a core functional language, it is equipped with an operational semantics, close to implementation, which can be expressed by an evaluation strategy that selects a unique β -redex to reduce:

- Never reduce a λ -abstraction, as it is a “value” (this is called *weak reduction*)
- Always reduce M first in an application $M N$. Then:
 - If M is an abstraction: reduce the β -redex first (CBN)
reduce N first (CBV)
 - Otherwise, reduce N (never happens with closed terms)

We denote those strategies $\longrightarrow_{\text{CBN}}$ and $\longrightarrow_{\text{CBV}}$.⁹

As a theory of functions, the λ -calculus is equipped with a denotational semantics close to the mathematical notion of functions: in particular, equalities are congruences (e.g. if $M = N$ then $\lambda x.M = \lambda x.N$) and reductions are congruences (this is called *strong reduction*). In [Plo75], Plotkin investigated the concepts of call-by-name and call-by-value by identifying particular λ -terms as values:

DEFINITION 15 (Value, β_v , Call-by-Name and Call-by-Value)

λ -terms of the form $\lambda x.M$ and x are called *values*, and denoted V , V' , etc, while λ -terms of the form MN are not values.¹⁰

- “Call-by-name” evaluation is given by general β -reduction

$$(\beta) \quad (\lambda x.M) N \longrightarrow \left\{ \frac{N}{x} \right\} M$$

- “Call-by-value” evaluation is given by the restriction of β -reduction where the argument is a value

$$(\beta_v) \quad (\lambda x.M) V \longrightarrow \left\{ \frac{V}{x} \right\} M$$

※

Now, natural questions to raise are

CBN: whether there is a relation between $\longrightarrow_{\text{CBN}}$ and $\longrightarrow_{\beta}$;

CBV: whether there is a relation between $\longrightarrow_{\text{CBV}}$ and $\longrightarrow_{\beta_v}$?

⁹For instance, Haskell implements CBN while OCaml implements CBV.

¹⁰The intuition is that, by evaluating MN , you may get a λ -term of a completely different shape.

Clearly, $\rightarrow_{\text{CBN}} \subseteq \rightarrow_{\beta}$ and $\rightarrow_{\text{CBV}} \subseteq \rightarrow_{\beta_v}$, but what about the converse? A bridge between weak and strong reductions was given by Plotkin [Plo75]:

THEOREM 13 (CBN and CBV: weak and strong reductions)

CBN: $\rightarrow_{\beta}^*$ is the closure of $\rightarrow_{\text{CBN}}^*$ under the rules

$$\frac{M_1 \rightarrow_{\text{CBN}}^* \lambda x.M_2 \quad M_2 \rightarrow M_2'}{M_1 \rightarrow \lambda x.M_2'} \quad \frac{M_1 \rightarrow_{\text{CBN}}^* M_2 \quad M_2 M_3 \quad M_2 \rightarrow M_2' \quad M_3 \rightarrow M_3'}{M_1 \rightarrow M_2' M_3'}$$

CBV: $\rightarrow_{\beta_v}^*$ is the closure of $\rightarrow_{\text{CBV}}^*$ under the rules

$$\frac{M_1 \rightarrow_{\text{CBV}}^* \lambda x.M_2 \quad M_2 \rightarrow M_2'}{M_1 \rightarrow \lambda x.M_2'} \quad \frac{M_1 \rightarrow_{\text{CBV}}^* M_2 \quad M_2 M_3 \quad M_2 \rightarrow M_2' \quad M_3 \rightarrow M_3'}{M_1 \rightarrow M_2' M_3'}$$

※

The point of this result is that we shall now call CBN and CBV, not some operational semantics of some functional programming language, but some rewriting theories in λ -calculus.

As we have already mentioned compilation in reference to CBN/CBV, it is interesting to see that λ -calculus can be compiled into (a fragment of) itself: this is based on the idea of the program transformation presented in Section 1.2 using continuations. As continuations are passed as an extra argument to every call, such transformations are known as *Continuation Passing Style* (CPS)-translations.

DEFINITION 16 (CPS-translations)

Two important CPS-translations were defined for CBN and CBV:

CBN-translation (Plotkin [Plo75]) $\begin{aligned} \underline{x} &:= \lambda k.x \ k \\ \underline{\lambda x.M} &:= \lambda k.k \ (\lambda x.\underline{M}) \\ \underline{M \ N} &:= \lambda k.\underline{M} \ (\lambda y.y \ \underline{N} \ k) \end{aligned}$	CBV-translation (Reynolds [Rey72]) $\begin{aligned} \bar{x} &:= \lambda k.k \ x \\ \overline{\lambda x.M} &:= \lambda k.k \ (\lambda x.\lambda k'.\bar{M} \ k') \\ \overline{M \ N} &:= \lambda k.\bar{M} \ (\lambda y.\bar{N} \ (\lambda z.y \ z \ k)) \end{aligned}$
--	---

where the variables k and k' are always chosen to be fresh.

※

One main feature of these translations is their target fragment of the λ -calculus: in this fragment, arguments are always values! This fragment is stable under $\rightarrow_{\beta}$ and $\rightarrow_{\beta_v}$, which actually coincide. The evaluation of a CPS-translated term is *strategy-indifferent*. How this evaluation relates to the evaluation of the original term is given by the following simulation properties:

THEOREM 14 (CPS-translations preserve reductions)

Soundness:

CBN If $M \rightarrow_{\beta} N$ then $\underline{M} \rightarrow_{\beta}^* \underline{N}$

CBV If $M \rightarrow_{\beta_v} N$ then $\bar{M} \rightarrow_{\beta_v}^* \bar{N}$

Completeness:

CBN If $\underline{M} \leftrightarrow_{\beta}^* \underline{N}$ then $M \leftrightarrow_{\beta}^* N$

CBV It is **not** the case for CBV, that if $\bar{M} \leftrightarrow_{\beta_v}^* \bar{N}$ then $M \leftrightarrow_{\beta_v}^* N$.

※

Proof: It is interesting to look at soundness to see how or why the CPS-translations make sense; for complete proofs, see [Plo75].

$$\begin{aligned}
(\lambda x.M) N &= \lambda k.(\lambda k'.(k' (\lambda x.M))) (\lambda y.y \ N \ k) \\
&\longrightarrow \lambda k.(\lambda y.y \ N \ k) (\lambda x.M) \\
&\longrightarrow \lambda k.(\lambda x.M) \ N \ k \\
&\longrightarrow \lambda k.(\{N/x\} \ M) \ k \\
&= \lambda k.(\{N/x\} \ M) \ k \\
&\longrightarrow \{N/x\} \ M
\end{aligned}$$

The second equality of course relies on the property that the CPS-translation behaves well with substitution: $(\{N/x\} \ M = \{N/x\} \ M)$. The last rewrite is an instance of β -reduction because $\{N/x\} \ M$ necessarily starts with a λ -abstraction (i.e. we are not using η -reduction).

$$\begin{aligned}
(\lambda x.M) \bar{V} &= \lambda k.(\lambda k'.(k' (\lambda x k''. \bar{M} \ k''))) (\lambda y. \bar{V} (\lambda z.y \ z \ k)) \\
&\longrightarrow \lambda k.(\lambda y. \bar{V} (\lambda z.y \ z \ k)) (\lambda x k''. \bar{M} \ k'') \\
&\longrightarrow \lambda k. \bar{V} (\lambda z.(\lambda x k''. \bar{M} \ k'') \ z \ k) \\
&\longrightarrow \lambda k. \bar{V} (\lambda z.(\lambda k''. \{z/x\} \ \bar{M} \ k'') \ k) \\
&= \lambda k. \bar{V} (\lambda x.(\lambda k''. \bar{M} \ k'') \ k) \\
&\longrightarrow \lambda k. \bar{V} (\lambda x. \bar{M} \ k) \\
&\longrightarrow \lambda k.(\lambda x. \bar{M} \ k) \ H \quad \text{where } \bar{V} = \lambda k.k \ H \\
&\longrightarrow \lambda k.(\{H/x\} \ \bar{M}) \ k \\
&= \lambda k.(\{V/x\} \ M) \ k \\
&\longrightarrow \{V/x\} \ M
\end{aligned}$$

The second equality is simply the renaming of z into x ; the third one relies again on the property that the CPS-translation behaves well with substitution by values $(\{H/x\} \ M = \{V/x\} \ M \text{ if } V = \lambda k.kH)$. The last rewrite is again an instance of β -reduction, not η -reduction.

The point here is to realise that if V had not been a value, then $\bar{V}$ would not be of the form $\lambda k.kH$, and the simulation of this specific β -reduction would be stuck. $\square$

The above simulations give some intuition about the encodings: the translation of any term M starts with a λ -abstract on a fresh variable k that is used exactly once. The variable k stands for the current continuation (hence the expression *continuation-passing style*). In the encoding of an abstraction $\lambda x.M$ (which is a value), the current continuation is applied to the encoding of the body M under a λ -abstraction on x . In case of an application $M \ N$, the current continuation is not directly applied, but wrapped in a bigger continuation that is passed as an argument to the encoding of M ; what this wrapping exactly is depends on whether we do CBN or CBV and will determine whether we reflect the evaluation of N as a value V before we reflect the reduction of $M \ N$.

Now, the fact that $\bar{M} \longleftrightarrow_{\beta}^* \bar{N}$ does not imply $M \longleftrightarrow_{\beta_v}^* N$ is slightly disappointing: one way to look at it is to consider that $M \longleftrightarrow_{\beta_v}^* N$ is too weak, or incomplete, for Call-by-Value. Indeed, the inspiration from monads, and Moggi's monadic λ -calculus [Mog89], has allowed the extension of the Call-by-Value λ -calculus into a sound and complete calculus with respect to the CBV CPS-translation (see for instance [Len06]).

We now turn to the behaviour of the CPS-translations with respect to typing: Assume we have $\Gamma \vdash M : A$. Do we have: $\Gamma' \vdash \bar{M} : A'$ (for some Γ', A') and $\Gamma'' \vdash \bar{M} : A''$ (for some Γ'', A'')?

The CPS-translations reveal two classes of terms in the target: *values* & *continuations* (like k). The types of values and continuations in the translated terms depend on CBN or CBV:

DEFINITION 17 (CPS-translations of simple types)

We choose or we add a particular atomic type R , called the *response type*, then we define

CBN	CBV
$\underline{a} := a$	$\bar{a} := a$
$\underline{A \rightarrow B} := ((\underline{A} \rightarrow R) \rightarrow R) \rightarrow (\underline{B} \rightarrow R) \rightarrow R$	$\overline{A \rightarrow B} := \overline{A} \rightarrow (\overline{B} \rightarrow R) \rightarrow R$

※

Intuitively, a type A in the original calculus will give rise to a type $\underline{A}$ (resp. $\overline{A}$) of “ A -values”; continuations are functions consuming those and returning something in the response type R (which is abstract in the sense that we will never need to know what it is), so continuations will therefore be of type $\underline{A} \rightarrow R$ (resp. $\overline{A} \rightarrow R$).

The encoding $\underline{M}$ (resp. $\overline{M}$) of a term M of type A will take the current continuation, of type $\underline{A} \rightarrow R$ (resp. $\overline{A} \rightarrow R$), and using that continuation, it will eventually output a response in the response type (as we have seen, the encoding starts with $\lambda k. \dots$). It will therefore be of type $(\underline{A} \rightarrow R) \rightarrow R$ (resp. $(\overline{A} \rightarrow R) \rightarrow R$).

This is formalised as the following theorem:

THEOREM 15 (CPS-translations preserve types)

If $\Gamma \vdash M : A$ then $(\underline{\Gamma} \rightarrow R) \rightarrow R \vdash \underline{M} : (\underline{A} \rightarrow R) \rightarrow R$ and $\overline{\Gamma} \vdash \overline{M} : (\overline{A} \rightarrow R) \rightarrow R$,
 where $\overline{x_1 : A_1, \dots, x_n : A_n}$ stands for $x_1 : \overline{A_1}, \dots, x_n : \overline{A_n}$
 and $((\underline{x_1 : A_1, \dots, x_n : A_n}) \rightarrow R) \rightarrow R$ stands for $x_1 : (\underline{A_1} \rightarrow R) \rightarrow R, \dots, x_n : (\underline{A_n} \rightarrow R) \rightarrow R$. ※

Proof: Straightforward induction on M . □

Variants of CPS-translations exist, of which we mention two that are related to CBN and CBV:

DEFINITION 18 (Variants)

- Fischer’s translation for CBV [Fis72]

$\overline{x} := \lambda k. k \ x$	$\bar{a} := a$
$\overline{\lambda x. M} := \lambda k. (k \ (\lambda k'. \lambda x. \overline{M} \ k'))$	$\overline{A \rightarrow B} := (\overline{B} \rightarrow R) \rightarrow \overline{A} \rightarrow R$
$\overline{M \ N} := \lambda k. \overline{M} \ (\lambda y. \overline{N} \ (\lambda z. y \ k \ z))$	

- Hofmann & Streicher’s translation for CBN [HS97], using product types

$\underline{x} := \lambda k. x \ k$	$\underline{a} := a \rightarrow R$
$\underline{\lambda x. M} := \lambda(x, k). \underline{M} \ k$	$\underline{A \rightarrow B} := (\underline{A} \rightarrow R) \times \underline{B}$
$\underline{M \ N} := \lambda k. \underline{M} \ (\underline{N}, k)$	

※

Fischer’s CBV-translation is very similar to Reynolds’s: they only differ in the order in which arguments are passed in the encoding of λ -abstractions and applications (e.g. for the abstraction, Reynolds’s translation binds x first, then binds the continuation variable k' , whereas Fischer’s binds k' first, then x). This is reflected in the encoding of the function type $A \rightarrow B$: the two arguments are swapped.

Hofmann & Streicher’s CBN-translation differs more importantly from Plotkin’s, as fewer “continuation wrappings” are introduced, reflected in the number of $\dots \rightarrow R$ in the encoding

of types: that encoding works “negatively”, as $\underline{A}$ is directly the type of A -continuations which are not necessarily functions consuming an A -value and returning in the response type.

Again, these translations allow the same simulations as Plotkin’s and Reynolds’s, and of course preserve typing, with a slightly different formulation in the case of CBN:

THEOREM 16 (Hofmann & Streicher’s CBN-translation preserves types)

If $\Gamma \vdash M : A$ then $\underline{\Gamma} \rightarrow R \vdash \underline{M} : \underline{A} \rightarrow R$

where $(\underline{x_1 : A_1}, \dots, \underline{x_n : A_n}) \rightarrow R$ stands for $x_1 : \underline{A_1} \rightarrow R, \dots, x_n : \underline{A_n} \rightarrow R$. ※

Let us now look at CPS-translations with respect to denotational semantics: Remember that simply-typed λ -terms have a semantics in a Cartesian Closed Category. CPS-translations compile the simply-typed λ -calculus into itself (preserving types in the sense of Theorem 15), so we can now assign to a simply-typed λ -term M , the semantics (in a CCC) of $\underline{M}$ or $\overline{M}$ (so that semantics now depends on CBN/CBV). By the simulation theorem (Theorem 14), reductions are sound w.r.t. their corresponding semantics.

More interestingly, notice that we do not need the whole structure of a CCC to build those two semantics, as $\underline{M}$ or $\overline{M}$ live in a fragment of the simply-typed λ -calculus (the CPS-fragment), where in particular the types of $\underline{M}$ or $\overline{M}$ are functional types. More than this, every functional type that we ever need for that fragment is of the form $A \rightarrow R$.¹¹ Therefore, in order to build the categorical semantics of the CPS-fragment, we do not need as strong axioms as those of a CCC: on top of asking for cartesian products we only require the existence of exponential objects of the form R^A . This is called a *response category*.

Now given a response category, the sub-category made of the objects of the form R^A is called a *continuation category*, a.k.a. *control category* (Selinger [Sel01]). Such a category turns out to have a rich structure that proves very useful for classical logic: not only it is a CCC (with exponential objects $(R^A)^{(R^B)}$ defined as $R^{A \times (R^B)}$) but objects of the form $R^{A \times B}$, denoted $R^A \wp R^B$, will play an important role.

1.7 Classical logic and CBN/CBV

We now relate classical logic to the above notions. We first review known translations from classical logic into intuitionistic logic: The intuition is that we can always turn P into P' by adding (enough) double negations, to get the property that

$$\text{If } \vdash_c P \text{ then } \vdash_i P'.$$

where $\vdash_c$ denotes classical provability and $\vdash_i$ denotes intuitionistic provability. Obviously, $\vdash_c P \leftrightarrow P'$, since the two formulae only differ by some double negations.

A potential question is then: If it suffices to add double negations in a classically provable formula to make it intuitionistically provable, are the two logics *really* different? Well, they differ at least in the sense that double negations break some of the nice properties of intuitionistic logic:

$$\begin{aligned} &\text{If } \vdash_i A_1 \vee A_2 \text{ then either } \vdash_i A_1 \text{ or } \vdash_i A_2. \\ &\text{If } \vdash_i \exists x A \text{ then there is } t \text{ such that } \vdash_i \{t/x\} A \end{aligned}$$

¹¹To be precise, we did use types such as $A_1 \rightarrow \dots \rightarrow A_n \rightarrow R$, but if we have products we can consider this to be the type $(A_1 \times \dots \times A_n) \rightarrow R$.

Getting t from the proof of $\vdash_i \exists x A$ is called *witness extraction*. This can also be done in some theories, like (Heyting) arithmetics:

If $HA \vdash_i \exists x A$ then there is t such that $HA \vdash_i \{t/x\} A$.

But in the most general case we cannot have the same properties when $\vdash_i \neg\neg(A_1 \vee A_2)$ or $\vdash_i \neg\neg\exists x A$. So what to do with a classical proof of $\vdash \exists x A$ is unclear. However, it is known that if A satisfies some specific property, a witness may be obtained from a classical proof of $\vdash \exists x A$; this is called classical witness extraction, and we will see this in Chapter 2.

The principle of inserting double negations gives rise to double negation translations (or $\neg\neg$ -translations), of which we present two, remembering that $\neg A$ is $A \Rightarrow \perp$:

DEFINITION 19 (Double negation translations)

$$\left| \begin{array}{ll} a^\bullet & := a \\ (A \Rightarrow B)^\bullet & := ((A^\bullet \Rightarrow \perp) \Rightarrow \perp) \Rightarrow (B^\bullet \Rightarrow \perp) \Rightarrow \perp \end{array} \right. \quad \left| \begin{array}{ll} a^\star & := a \\ (A \Rightarrow B)^\star & := A^\star \Rightarrow (B^\star \Rightarrow \perp) \Rightarrow \perp \end{array} \right. \quad \ast$$

We realise here that these translations, via the Curry-Howard correspondence, are exactly the translations of types from Definition 17 that make Plotkin's and Reynolds's translations “preserve types”: The response type previously denoted R corresponds to the formula $\perp$, and a continuation is a proof of negation.

1.7.1 Identifying CBN and CBV in System L

The fact that double negation translations allow the construction of an intuitionistic proof of $A^\bullet$ (resp. $A^\star$) from a classical proof of A , suggests that we can adapt the CPS-translations of Definitions 16 and 18 to encode classical proof-terms, say of System L, into the simply-typed λ -calculus. If this encoding not only preserves types but also reductions (as in e.g. Theorem 14), then we could assign to a classical proof-term the categorical semantics of its CPS-encoding (which, as a simply-typed λ -term, is well-understood).

It remains to identify which reductions of System L will be reflected in the CPS-encoding.

Inspired by Theorem 14, we remark that the β -reductions that can be reflected by the CBV-encoding are of the form β_v , i.e. those reductions where every substitution that is computed substitute a variable by a value. In System L, we can impose similar restrictions: a CBV-reduction should only allow a substitution $\{t/x\} c$ to be computed if t is a “value”; and by symmetry, we could expect CBN-reduction to only allow a substitution $\{e/\alpha\} c$ to be computed if e is a “continuation value”. But we still need to identify what the notions of values and continuation values are for System L. Considering the non-confluence situation $\langle \mu\alpha.c \mid \mu x.c' \rangle$ described in Section 1.4 (which causes so much difficulty for building semantics for System L), ruling out $\mu\alpha.c$ as value and ruling out $\mu x.c'$ as continuation value solves the problem: CBN-reduction would allow the reduction $\langle \mu\alpha.c \mid \mu x.c' \rangle \rightarrow \{ \mu\alpha.c / \alpha \} c'$ and disallow $\langle \mu\alpha.c \mid \mu x.c' \rangle \rightarrow \{ \mu x.c' / x \} c$, while CBV-reduction would allow the reduction $\langle \mu\alpha.c \mid \mu x.c' \rangle \rightarrow \{ \mu x.c' / x \} c$ and disallow $\langle \mu\alpha.c \mid \mu x.c' \rangle \rightarrow \{ \mu\alpha.c / \alpha \} c'$. In other words, CBV-reduction gives priority to the right while CBN-reduction gives priority to the left.

We can formalise this as the following definition:

DEFINITION 20 (CBN and CBV for System L -first attempt)

We identify the following notions of term values and continuation values:

$$\begin{array}{ll} \text{Term values} & V ::= x \mid \lambda x.t \mid (t_1, t_2) \mid \text{inj}_i(t) \\ \text{Continuation values} & E ::= \alpha \mid t::e \mid (e_1, e_2) \mid \text{inj}_i(e) \end{array}$$

The reduction relations $\rightarrow_{\text{CBN}}$ and $\rightarrow_{\text{CBV}}$ are defined as the contextual closures of the (groups of) rules in Fig. 9. ※

$(\rightarrow)$	$\langle \lambda x.t_1 \mid t_2::e \rangle$	$\rightarrow$	$\langle t_2 \mid \mu x.\langle t_1 \mid e \rangle \rangle$
$(\wedge)$	$\langle (t_1, t_2) \mid \text{inj}_i(e) \rangle$	$\rightarrow$	$\langle t_i \mid e \rangle$
$(\vee)$	$\langle \text{inj}_i(t) \mid (e_1, e_2) \rangle$	$\rightarrow$	$\langle t \mid e_i \rangle$
$(\overleftarrow{\mu}_{\text{N}})$	$\langle \mu\beta.c \mid E \rangle \rightarrow \{E/\beta\}c$	$(\overleftarrow{\mu})$	$\langle \mu\beta.c \mid e \rangle \rightarrow \{e/\beta\}c$
$(\overrightarrow{\mu})$	$\langle t \mid \mu x.c \rangle \rightarrow \{t/x\}c$	$(\overrightarrow{\mu}_{\text{V}})$	$\langle V \mid \mu x.c \rangle \rightarrow \{V/x\}c$
CBN		CBV	

Figure 8: CBN and CBV reduction in System L (first attempt)

The fact that CBV-reduction keeps $(\overleftarrow{\mu})$ and restricts $(\overrightarrow{\mu})$ into $(\overrightarrow{\mu}_{\text{V}})$ (and vice versa for CBN), is the formalisation of what we described before.

Doing this “works” in the sense that both CBN and CBV reductions are confluent systems (as higher-order orthogonal rewrite systems).

Unfortunately, these restrictions are not sufficient to build the denotational semantics of those systems according to methodology of CPS-translations, at least if we are to re-use the CPS-translations of types that we have seen in Section 1.6: Indeed, the simulation property that would be, for System L, the equivalent of Theorem 14, fails.

This was noticed in an erratum of [CH00], which also notices that the simulation does work on two specific fragments of System L: Concentrating on the implicational fragment,

- CBN-reduction can be simulated in the λ -calculus (according to Hofmann and Streicher’s translation of types [HS97]) when every continuation of the form $t::e$ is such that t is a term value;
- CBV-reduction can be simulated in the λ -calculus (according to Reynold’s or Fischer’s translation of types [Rey72, Fis72]) when every continuation of the form $t::e$ is such that e is a continuation value.

Note that this makes sense because the former and the latter fragments are stable under CBN and CBV reduction, respectively.

This also suggests how to refine the notions of term values and continuation values as follows: instead of these notions concerning only the top-level construct of a term or a continuation, our new and more appropriate notions of values will be recursively defined.

This *strong* notion of value is taken primarily from [Wad03]:¹²

¹²Inspired by [MM09], we make a change about implication in the case of CBV, for which we *also* restrict continuation values, since this will make CBV normal forms correspond to proofs in LKQ [DJS95, DJS97], as we shall see in Chapter 3.

DEFINITION 21 (CBN and CBV for System L)

In CBN, we identify the following notion of *continuation values*:

$$\text{CBN continuation values } E ::= \alpha \mid t :: E \mid (E_1, E_2) \mid \text{inj}_i(E)$$

In CBV, we identify the following notion of *term values* and *continuation values*:

$$\text{CBV term values } V ::= x \mid \lambda x.t \mid (V_1, V_2) \mid \text{inj}_i(V)$$

$$\text{CBV continuation values } F ::= \alpha \mid V :: F \mid (F_1, F_2) \mid \text{inj}_i(F) \mid \mu x.c$$

The reduction relations $\rightarrow_{\text{CBN}}$ and $\rightarrow_{\text{CBV}}$ are the contextual closures of the rules in Fig. 9. ※

$(\overleftarrow{\mu}_{\text{N}})$	$\langle \mu\beta.c \mid E \rangle$	$\rightarrow \{E/\beta\}c$	$(\overleftarrow{\mu})$	$\langle \mu\beta.c \mid e \rangle$	$\rightarrow \{e/\beta\}c$
$(\overrightarrow{\mu})$	$\langle t \mid \mu x.c \rangle$	$\rightarrow \{t/x\}c$	$(\overrightarrow{\mu}_{\text{V}})$	$\langle V \mid \mu x.c \rangle$	$\rightarrow \{V/x\}c$
(ζ_{N})	$R[e]$	$\rightarrow \langle \mu\alpha.R[\alpha] \mid e \rangle$	(ζ_{V})	$S[t]$	$\rightarrow \langle t \mid \mu x.S[x] \rangle$
			(ζ_{V})	$T[e]$	$\rightarrow T[\mu x.\langle x \mid e \rangle]$
$(\rightarrow_{\text{N}})$	$\langle \lambda x.t_1 \mid t_2 :: E \rangle$	$\rightarrow \langle \{t_2/x\}t_1 \mid E \rangle$	$(\rightarrow_{\text{V}})$	$\langle \lambda x.t_1 \mid V :: F \rangle$	$\rightarrow \langle \{V/x\}t_1 \mid F \rangle$
$(\wedge_{\text{N}})$	$\langle (t_1, t_2) \mid \text{inj}_i(E) \rangle$	$\rightarrow \langle t_i \mid E \rangle$	$(\wedge_{\text{V}})$	$\langle (V_1, V_2) \mid \text{inj}_i(F) \rangle$	$\rightarrow \langle V_i \mid F \rangle$
$(\vee_{\text{N}})$	$\langle \text{inj}_i(t) \mid (E_1, E_2) \rangle$	$\rightarrow \langle t \mid E_i \rangle$	$(\vee_{\text{V}})$	$\langle \text{inj}_i(V) \mid (F_1, F_2) \rangle$	$\rightarrow \langle V \mid F_i \rangle$
CBN			CBV		

where R , S , and T range over contexts of the following grammar:

$$\text{CBN continuation contexts } R ::= \langle t \mid t' :: [] \rangle \mid \langle t \mid ([], e) \rangle \mid \langle t \mid (E, []) \rangle \mid \langle t \mid \text{inj}_i([]) \rangle$$

$$\text{CBV term contexts } S ::= \langle V \mid [] :: e \rangle \mid \langle ([], t) \mid e \rangle \mid \langle (V, []) \mid e \rangle \mid \langle \text{inj}_i([]) \mid e \rangle$$

$$\text{CBV continuation contexts } T ::= \langle V \mid V' :: [] \rangle \mid \langle V \mid ([], e) \rangle \mid \langle V \mid (F, []) \rangle \mid \langle V \mid \text{inj}_i([]) \rangle$$

the (ζ_{N}) only applies under the condition that e is not a (CBN-) continuation value,

the (ζ_{V}) -rules only apply under the condition that t is not a (CBV-) term value and e is not a (CBV-) continuation value.

Figure 9: CBN and CBV reduction in System L

In this version, we kept the CBV-rules $(\overleftarrow{\mu})$ and $(\overrightarrow{\mu}_{\text{V}})$, and the CBN-rules $(\overleftarrow{\mu}_{\text{N}})$ and $(\overrightarrow{\mu})$.

The (ζ_{V}) -rules (resp. the (ζ_{N}) -rule) are new: they were introduced in a slightly more general version in [Wad03], while the version we take here more closely follows [MM09]. These rules are due to our strong restriction on term values (resp. continuation values): the fact that a term is a value is not just the fact that it is not of the form $\mu\alpha.c$, as term values are recursively defined. Therefore if a term t is not of the form $\mu\alpha.c$ but one of its (say direct) subterms is, then t is not a value and there is no CBV-rule to reduce $\langle t \mid \mu x.c \rangle$. Progress then fails if we do not add the (ζ_{V}) -rules to pull the first subterm of t that is not a value to the top-level.

These ζ rules also impact rules $(\rightarrow)$, $(\wedge)$, and $(\vee)$, which now have to be restricted in order to preserve confluence: for instance in CBV, the fact that $(\mu\alpha.c, t)$ is not a term value means that, when facing a continuation e , $\mu\alpha.c$ will be extracted from the pair and will have the control of computation

$$\langle (\mu\alpha.c, t) \mid e \rangle \rightarrow_{\zeta_{\text{V}}} \langle \mu\alpha.c \mid \mu x.\langle (x, t) \mid e \rangle \rangle \rightarrow_{\overleftarrow{\mu}} \left\{ \mu x.\langle (x, t) \mid e \rangle / \alpha \right\} c$$

and therefore it is clear that, should e be of the form $\text{inj}_2(e')$, the original application of rule $(\wedge)$ would have a totally different semantics:

$$\langle (\mu\alpha.c, t) \mid e \rangle \longrightarrow_{\wedge} \langle t \mid e' \rangle$$

Hence the restriction of $(\rightarrow)$, $(\wedge)$, and $(\vee)$ to $(\rightarrow_N)$, $(\wedge_N)$, and $(\vee_N)$ in the CBN case, and to $(\rightarrow_V)$, $(\wedge_V)$, and $(\vee_V)$ in the CBV case.

Note that in CBN, we decided to make $(\rightarrow_N)$ collapse the two reduction steps

$$\langle \lambda x.t_1 \mid t_2 :: E \rangle \longrightarrow_{\rightarrow} \langle t_2 \mid \mu x.\langle t_1 \mid E \rangle \rangle \longrightarrow_{\mu} \langle \{t_2/x\} t_1 \mid E \rangle$$

into one step, because (μ) has priority anyway.¹³ The rule $(\rightarrow_V)$ is designed by symmetry, collapsing the two steps

$$\langle \lambda x.t_1 \mid V :: F \rangle \longrightarrow_{\rightarrow} \langle V \mid \mu x.\langle t_1 \mid F \rangle \rangle \longrightarrow_{\mu} \langle \{V/x\} t_1 \mid F \rangle$$

and noticing that if t_2 is not a term value, then the application of the original rule $(\rightarrow)$ is recovered as follows:

$$\langle \lambda x.t_1 \mid t_2 :: F \rangle \longrightarrow_{\zeta_V} \langle t_2 \mid \mu y.\langle \lambda x.t_1 \mid y :: F \rangle \rangle \longrightarrow_{\rightarrow_V} \langle t_2 \mid \mu x.\langle t_1 \mid F \rangle \rangle$$

Of course the extra rules satisfy Subject Reduction, so that we have:

THEOREM 17 (Subject reduction for System L: CBN & CBV)

If $c : (\Gamma \vdash \Delta)$ and either $c \longrightarrow_{\text{CBN}} c'$ or $c \longrightarrow_{\text{CBV}} c'$ then $c' : (\Gamma \vdash \Delta)$.

And similarly for terms and continuations. ※

Proof: Straightforward induction on the rewrite derivation. □

THEOREM 18 (Confluence) $\longrightarrow_{\text{CBN}}$ and $\longrightarrow_{\text{CBV}}$ are confluent. ※

Proof: They are orthogonal higher-order rewrite systems.¹⁴ □

Finally, one could be puzzled by what seems like an asymmetry between CBN and CBV, the latter having more rules and requiring a notion of continuation value while the former does not need a notion of term value. This asymmetry is not due to CBN vs. CBV, but is due to the implication: its main continuation construct $t :: e$ has a term as a direct sub-term, while no term construct has a continuation as a direct sub-term. It would be the case if we considered the De Morgan dual of implication, namely *subtraction* (see for instance [Cro04]), which would make CBN completely symmetric to CBV.

1.7.2 Two stable fragments

Now it is easy to connect the $\longrightarrow_{\text{CBN}}$ and $\longrightarrow_{\text{CBV}}$ reduction relations of Definition 21 to those of our first attempt in Definition 20, if we concentrate on the two fragments:¹⁵

DEFINITION 22 (LK^N and LK^V) Let LK^N and LK^V be the fragments of System L consisting of $\longrightarrow_{\zeta_N}$ -normal forms and $\longrightarrow_{\zeta_V}$ -normal forms, respectively. ※

¹³We shall see that it makes the $\mu x.$ construct superfluous (in the sense that the fragment without this construct is logically complete, and stable under reduction).

¹⁴The rewrite system presented in Fig. 9 is a standard (higher-order) rewrite system: we did use a non-standard formulation for the ζ rules based on a grammar for continuation contexts and term contexts as well as on side-conditions (“ t is not a term value and e is not a continuation value”), but we could equally have formulated all the cases as standard (but numerous!) rewrite rules.

¹⁵Namely, those fragments where the reductions of Definition 20 are actually simulated by (the adaptation to System L of) the CPS-translations.

REMARK 19

1. Concerning implication only, LK^N is the fragment where every continuation of the form $t::e$ is such that e is a continuation value.
2. Concerning implication only, LK^V is the fragment where every continuation of the form $t::e$ is such that t is a term value;
3. Also, $\rightarrow_{\zeta_N}$ and $\rightarrow_{\zeta_V}$ are terminating reduction relations, so it is easy to normalise a command into one of these fragments, using cuts.
4. Moreover, LK^N and LK^V are respectively stable under $\rightarrow_{CBN}$ and $\rightarrow_{CBV}$, so the cuts can be reduced while staying in the fragments.
5. Furthermore, in LK^N and LK^V , $\rightarrow_{CBN}$ and $\rightarrow_{CBV}$ of Definition 21 respectively coincide with those of Definition 20.¹⁶
6. Notice that the encoding of λ -calculus in System L in Definition 14, actually only reaches the fragment LK^N , and the simulation lemma (Lemma 9) only involves CBN-reduction.

※

1.7.3 Denotational semantics of CBN and CBV

As anticipated, it is now possible to define CPS-translations of terms, continuations, and commands, respectively denoted $\underline{t}$, $\underline{e}$, $\underline{c}$ for CBN, and $\bar{t}$, $\bar{e}$, $\bar{c}$ for CBV, in a way that preserves reductions:

THEOREM 20 (Preservation of reduction)

CBN: If $c_1 \rightarrow_{CBN} c_2$ then $\underline{c_1} \rightarrow_{\beta}^* \underline{c_2}$

CBV: If $c_1 \rightarrow_{CBV} c_2$ then $\bar{c_1} \rightarrow_{\beta}^* \bar{c_2}$

※

We do not give the details here, which are just technical, but they can be found in e.g. [Wad03].

And these encodings also preserve typing, if Hofmann and Streicher's encoding of types for CBN, and Fischer's encoding of types for CBV, are considered not just for $\rightarrow$ (i.e. $\Rightarrow$) but also $\times$ (i.e. $\wedge$) and $+$ (i.e. $\vee$):

THEOREM 21 (Preservation of typing)

Assume

$$\Gamma \vdash t : A ; \Delta$$

$$\Gamma ; e : A \vdash \Delta$$

$$c : (\Gamma \vdash \Delta)$$

¹⁶Almost, since in Definition 21 the rule $(\rightarrow_N)$ (resp. $(\rightarrow_V)$) collapses the two steps $\rightarrow_{\rightarrow} \cdot \rightarrow_{\mu}^{\rightarrow}$ (resp. $\rightarrow_{\rightarrow} \cdot \rightarrow_{\mu_V}^{\rightarrow}$) of $\rightarrow_{CBN}$ (resp. $\rightarrow_{CBV}$) from Definition 20: but in LK^N (resp. LK^V), there is no other choice than $\rightarrow_{\mu}^{\rightarrow}$ (resp. $\rightarrow_{\mu_V}^{\rightarrow}$) for a top-level reduction that can follow $\rightarrow_{\rightarrow}$ in CBN-reduction (resp. CBV-reduction).

Then

$$\begin{array}{ll} \frac{\Gamma \rightarrow R, \underline{\Delta} \vdash \underline{t} : \underline{A} \rightarrow R}{\frac{\Gamma \rightarrow R, \underline{\Delta} \vdash \underline{e} : (\underline{A} \rightarrow R) \rightarrow R}{\Gamma \rightarrow R, \underline{\Delta} \vdash \underline{c} : R}} & \frac{\bar{\Gamma}, \bar{\Delta} \rightarrow R \vdash \bar{t} : (\bar{A} \rightarrow R) \rightarrow R}{\frac{\bar{\Gamma}, \bar{\Delta} \rightarrow R \vdash \bar{e} : \bar{A} \rightarrow R}{\bar{\Gamma}, \bar{\Delta} \rightarrow R \vdash \bar{c} : R}} \end{array}$$

Using Hofmann & Streicher's
translation of types [HS97]

Using Fischer's
translation of types [Fis72]

※

As already mentioned, we can now use these CPS-translations to define categorical semantics for classical proofs:

DEFINITION 23 (Semantics of System L in a response category)

Assume $c : (x_1 : A_1, \dots, x_n : A_n \vdash \alpha_1 : B_1, \dots, \alpha_m : B_m)$.

Define the semantics $\llbracket c \rrbracket_N^r := \llbracket \underline{c} \rrbracket$ and $\llbracket c \rrbracket_V^r := \llbracket \bar{c} \rrbracket$, where $\llbracket t \rrbracket$ is the semantics, in a response category, of a λ -term t in the CPS-fragment, as defined by the rules of Fig. 3.

Writing K_A for the object corresponding to $\underline{A}$, and C_A for R^{K_A} , we have

$$\llbracket c \rrbracket_N^r : (C_{A_1} \times \dots \times C_{A_n} \times K_{B_1} \times \dots \times K_{B_m}) \longrightarrow R$$

Writing V_A for the object corresponding to $\bar{A}$, K_A and for R^{V_A} , we have

$$\llbracket c \rrbracket_V^r : (V_{A_1} \times \dots \times V_{A_n} \times K_{B_1} \times \dots \times K_{B_m}) \longrightarrow R$$

※

Now, remember that a *control category* is the sub-category of a response category $\mathcal{C}$ whose objects are in $\{R^A \mid A \in \mathcal{C}\}$, and that $R^A \wp R^B$ denotes $R^{A \times B}$.

DEFINITION 24 (Semantics of System L in control and co-control categories)

Assume $c : (x_1 : A_1, \dots, x_n : A_n \vdash \alpha_1 : B_1, \dots, \alpha_m : B_m)$.

CBN The semantics $\llbracket c \rrbracket_N^r$ in a response category gives rise, by curryfication, to a morphism

$$\llbracket c \rrbracket_N : C_{A_1} \times \dots \times C_{A_n} \longrightarrow C_{B_1} \wp \dots \wp C_{B_m}$$

in a control category.

CBV The semantics $\llbracket c \rrbracket_V^r$ in a response category gives rise, by curryfication, to a morphism

$$\llbracket c \rrbracket_V : K_{B_1} \times \dots \times K_{B_m} \longrightarrow K_{A_1} \wp \dots \wp K_{A_n}$$

in a control category.

As the arrow looks “reversed”, from the original typing of c , it is more natural to interpret c as the corresponding morphism

$$\llbracket c \rrbracket_V : K_{A_1} \otimes \dots \otimes K_{A_n} \longrightarrow K_{B_1} + \dots + K_{B_m}$$

in a *co-control category*, the dual of a control category where $\otimes$ denotes the dual of $\wp$ (and the co-product $+$ denotes the dual of the product $\times$).

※

This formalises the idea that CBN-reduction corresponds to a denotational semantics in control categories, while CBV-reduction corresponds to a denotational semantics in co-control categories.

Indeed, from Theorem 20 we get that the semantics validate the reductions:

THEOREM 22 (Soundness of CBN and CBV in control and co-control categories)

$$\left| \begin{array}{l} \text{If } c \longrightarrow_{\text{CBN}} c' \text{ then } \llbracket c \rrbracket_{\text{N}} = \llbracket c' \rrbracket_{\text{N}} \\ \text{If } c \longrightarrow_{\text{CBV}} c' \text{ then } \llbracket c \rrbracket_{\text{V}} = \llbracket c' \rrbracket_{\text{V}} \end{array} \right. \quad \ast$$

And we did this by breaking the symmetry between $\wedge$ and $\vee$: Indeed in a control category, $\wp$ is not the dual of $\times$ (equivalently in a co-control category, $+$ is not the dual of $\otimes$).

Conclusion

The work on control and co-control categories is due to Selinger [Sel01] for Parigot's $\lambda\mu$ -calculus, showing a duality between CBN and CBV in the categorical sense of duality, and even before a syntactic duality between CBN and CBV was displayed by System L and its variants. It follows preliminary works by Hofmann, Streicher, Reus [HS97, SR98], etc, on the semantics of continuations (where the question of duality between CBV and CBN -in $\lambda\mu$ - is conjectured).

In conclusion, many variants of classical proof calculi have been studied; in particular,

- the De Morgan dual of implication in classical logic, namely *subtraction*, can also be given a computational interpretation, as shown for instance by [Cro04];
- variants of Parigot's $\lambda\mu$ have different properties with respect to observational equivalence, separation and η -conversion, etc... [Sau05, Sau08, HZ09, Sau10c, Sau10b, Sau12];
- *control delimiters* can be used to limit the scope of the context that can be captured by a term via control operators, allowing for instance the capture of the *shift* and *reset* operators of [DF89, DF90]; these give rise to *delimited continuations* and can be given a proof-theoretic interpretation [AHS09, HG08, Sau10a];
- other reduction strategies than CBV and CBN have been investigated under the light of the duality of computation, such as Call-by-Need [AF97, AHS11, ADH⁺12].

Chapter 2

Orthogonality: models for normalisation and witness extraction

Contents

2.1 Revisiting Proofs of Strong Normalisation for System F	52
2.1.1 Orthogonality models and the Adequacy Lemma	53
2.1.2 Applicative orthogonality models and Strong Normalisation	54
2.2 Adapting the approach to classical calculi	55
2.2.1 The case of a confluent calculus	56
2.2.2 The case of a non-confluent calculus	57
2.3 Orthogonality models for extracting witnesses from classical proofs	60
Conclusion	64

In this chapter we present the concept of *orthogonality* and two applications of it that are useful for classical proof-term calculi: strong normalisation and witness extraction.

Orthogonality was used by Girard in the context of *linear logic* [Gir87] to prove normalisation of cut-elimination and it lies at the heart of its proof semantics based on *coherent spaces*.

The concept of orthogonality has also proved a key concept in the proof theory of classical logic, as it features, just like linear logic does, a duality that is most immediately seen in the form of De Morgan's laws. Indeed, orthogonality is the basis of *classical realisability* [DK00, Kri01], which can be seen as a semantical approach to the Curry-Howard correspondence for classical logic. It also provided a new tool for models of classical proofs, and for proving properties of programs [Par97, MV05, LM08], most notoriously the strong normalisation property. Furthermore, orthogonality shed a new light on the theory of *polarisation* and *focussing* for classical logic, as revealed for instance in [MM09] and explored in Chapter 3.

In this chapter we start by illustrating how proofs of strong normalisation relate to orthogonality. Summarising [BL11b], Section 2.1 rephrases and modularises, with the notion of *orthogonality models*, the well-known techniques by Tait [Tai67, Tai75], Reynolds [Rey72] and Girard [Gir72] for proving the strong normalisation of the simply-typed λ -calculus and

System F . Section 2.2 shows how such models allow the adaptation, to classical proof-term calculi, of strong normalisation proofs, both in the case of a confluent calculus and non-confluent calculus. Section 2.3 then shows how orthogonality models can be used for classical witness extraction, using a technique due to Miquel [Miq09, Miq11].

2.1 Revisiting Proofs of Strong Normalisation for System F

This section presents concepts and a methodology developed in [BL11b]: in particular, we approach the strong normalisation of System F with the notion of *orthogonality model*, adapting the Tait-Girard methodology [Tai75, Gir72]. Although System F is an intuitionistic system, this approach will form the starting point from which the adaptation to classical logic will be explored.

DEFINITION 25 (System F)

The types of System F are given by the following grammar:

$$A, B, \dots ::= \alpha \mid A \rightarrow B \mid \forall \alpha A$$

where α ranges over a denumerable set of elements called *type variables*, and the construct $\forall \alpha A$ binds α in A .

Typing contexts are defined as in Definition 6, using System F types instead of simple types; they will be denoted Γ, Δ , etc.

The *free type variables* of a type A (resp. a typing context Γ) will be denoted $ftv(A)$ (resp. $ftv(\Gamma)$).

The typing system of System F is given in Fig. 10. Derivability of a sequent in System F is denoted $\Gamma \vdash_F M : A$. *

$$\begin{array}{c}
 \frac{}{\Gamma, x : A \vdash x : A} \quad \frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x. M : A \rightarrow B} \\
 \\
 \frac{\Gamma \vdash M : A \rightarrow B \quad \Gamma \vdash N : A}{\Gamma \vdash M N : B} \\
 \\
 \frac{\Gamma \vdash M : A}{\Gamma \vdash M : \forall \alpha A} \alpha \notin ftv(\Gamma) \quad \frac{\Gamma \vdash M : \forall \alpha A}{\Gamma \vdash M : \left\{ \frac{B}{\alpha} \right\} A}
 \end{array}$$

Figure 10: System F

The method to prove strong normalisation is to build a model with

- an interpretation for terms as elements of a set $\mathcal{E}$
- an interpretation for types as (interesting) subsets of $\mathcal{E}$
- such that
 - the interpretation of a term of type A is in the interpretation of A
 - if the interpretation of a term is in there, the term is strongly normalising.

2.1.1 Orthogonality models and the Adequacy Lemma

DEFINITION 26 (Orthogonality models)

An *orthogonality model* is a 4-tuple $(\mathcal{D}, \perp, \mathcal{E}, \llbracket _ \rrbracket)$ where

- $\mathcal{D}$ is a set of elements called *values*;
- $\perp$ is a relation between values and lists of values called the *orthogonality relation*;
- $\mathcal{E}$ is a superset of $\mathcal{D}$;
- $\llbracket _ \rrbracket$ is a function mapping every λ -term M (typed or untyped) to an element $\llbracket M \rrbracket_\rho$ of $\mathcal{E}$, where ρ is a parameter called *semantic context* mapping term variables to values.
- the following axioms are satisfied:
 - (A1) For all $\rho, \vec{v}, x$, if $\rho(x) \perp \vec{v}$ then $\llbracket x \rrbracket_\rho \perp \vec{v}$.
 - (A2) For all $\rho, \vec{v}, M_1, M_2$, if $\llbracket M_1 \rrbracket_\rho \perp (\llbracket M_2 \rrbracket_\rho :: \vec{v})$ then $\llbracket M_1 M_2 \rrbracket_\rho \perp \vec{v}$.
 - (A3) For all $\rho, \vec{v}, x, M$ and for all values u , if $\llbracket M \rrbracket_{\rho, x \mapsto u} \perp \vec{v}$ then $\llbracket \lambda x. M \rrbracket_\rho \perp (u :: \vec{v})$.

※

In fact, $\mathcal{D}$ and $\perp$ are already sufficient to interpret any System F type A as a set $\llbracket A \rrbracket^+$ of values (see Definition 28 below): if types are seen as logical formulae, we can see this construction as a way of building some of their realisability / set-theoretical models.

There is no notion of computation pertaining to values, but the interplay between the interpretation of terms and the orthogonality relation is imposed by the axioms so that the Adequacy Lemma (which relates typing to semantics) holds:

$$\text{If } \vdash_F M : A \text{ then } \llbracket M \rrbracket \in \llbracket A \rrbracket^+$$

We now assume that we are given an orthogonality model $(\mathcal{D}, \perp, \mathcal{E}, \llbracket _ \rrbracket)$.

NOTATION 27 By $\mathcal{D}^*$ we denote the set of lists of values.

If $X \subseteq \mathcal{D}$ and $Y \subseteq \mathcal{D}^*$ let

$$\begin{aligned} X :: Y &:= \{u :: \vec{v} \mid u \in X, \vec{v} \in Y\} \\ X^\perp &:= \{\vec{v} \in \mathcal{D}^* \mid \forall u \in X, u \perp \vec{v}\} \\ Y^\perp &:= \{u \in \mathcal{D} \mid \forall \vec{v} \in Y, u \perp \vec{v}\} \end{aligned}$$

※

REMARK 23 The usual properties of orthogonality hold:

$$X \subseteq X^{\perp\perp}, X^{\perp\perp\perp} = X^\perp, \text{ and if } X \subseteq X' \text{ then } X'^\perp \subseteq X^\perp$$

※

We now define the interpretation of types. The intuition is the same as that of Krivine's classical realisability [DK00, Kri01]:

- we first interpret a formula A as a set of “counter-proofs”, with the basic constructs that we expect to use in order to refute the formula: for instance the basic way to refute $A_1 \rightarrow A_2$ is to provide a “proof” of A_1 and a “counter-proof” of A_2 ; similarly, the basic way to refute $\forall \alpha A_1$ is to find a suitable interpretation of α and produce a “counter-proof” of A_1 under this interpretation; the set of counter-proofs for atomic formulae is then naturally given by a *valuation*;
- we then define the “proofs” (“realisers” would be a better term) of a formula A as any value that is able to “face all counter-proofs”, this latter concept being what the orthogonality relation is precisely there to specify.

DEFINITION 28 (Interpretation of types)

A *valuation* is a function, denoted $\sigma, \sigma', \dots$, from type variables to subsets of $\mathcal{D}^*$.

Two interpretation of types are defined by simultaneous induction of types, a *positive interpretation* and a *negative interpretation*:

$$\begin{aligned} \llbracket A \rrbracket_\sigma^+ &:= \llbracket A \rrbracket_\sigma^{-\perp} & [\alpha]_\sigma^- &:= \sigma(\alpha) \\ \llbracket A \rightarrow B \rrbracket_\sigma^- &:= \llbracket A \rrbracket_\sigma^+ :: \llbracket B \rrbracket_\sigma^- & [\forall \alpha A]_\sigma^- &:= \bigcup_{Y \subseteq \mathcal{D}^*} \llbracket A \rrbracket_{\sigma, \alpha \mapsto Y}^- \end{aligned}$$

We then define the interpretation of typing contexts:

$$\llbracket \Gamma \rrbracket_\sigma := \{ \rho \mid \forall (x : A) \in \Gamma, \rho(x) \in \llbracket A \rrbracket_\sigma^+ \}$$

※

This approach begs the question: why is it the case that counter-proofs are defined more primitively than proofs? As described for instance in [MM09] (and in the rest of this thesis), this is simply a coincidence about the two type constructs we use in System F : they both have a “negative polarity”, if we see them in the more general context of *polarised logic*, as we shall discuss in Chapter 3. Second-order quantification is often used in the field of realisability (and elsewhere) to encode other logical connectives, which made the negative approach prevalent in that field, with counter-proofs being defined first and proofs being defined by orthogonality. With primitive connectives that have a “positive polarity”, such as intuitionistic disjunction, proofs would be defined first and counter-proofs would be defined by orthogonality. We shall come back to that discussion in the next chapters.

REMARK 24 We have the usual properties of substitutions:

$$\left[\left\{ \frac{B}{\alpha} \right\} A \right]_\sigma^- = \llbracket A \rrbracket_{\sigma, \alpha \mapsto \llbracket B \rrbracket_\sigma^-}^- \quad \text{and} \quad \left[\left\{ \frac{B}{\alpha} \right\} A \right]_\sigma^+ = \llbracket A \rrbracket_{\sigma, \alpha \mapsto \llbracket B \rrbracket_\sigma^-}^+$$

Also notice that the *for all* quantifier is interpreted as an intersection:

$$\llbracket \forall \alpha A \rrbracket_\sigma^+ = \bigcap_{Y \subseteq \mathcal{D}^*} \llbracket A \rrbracket_{\sigma, \alpha \mapsto Y}^+$$

※

With these definitions we can prove the Adequacy Lemma:

LEMMA 25 (Adequacy Lemma)

If $\Gamma \vdash_F M : A$, then for all valuations σ and for all mappings $\rho \in \llbracket \Gamma \rrbracket_\sigma$ we have $\llbracket M \rrbracket_\rho \in \llbracket A \rrbracket_\sigma^+$.

※

Proof: By induction on the derivation of $\Gamma \vdash M : A$, using axioms (A1), (A2) and (A3). See [BL11b]. $\square$

2.1.2 Applicative orthogonality models and Strong Normalisation**DEFINITION 29 (Applicative orthogonality model)**

An *applicative orthogonality model* is a 4-tuple $(\mathcal{D}, \mathcal{E}, @, \llbracket _ \rrbracket__)$ where:

- $\mathcal{D}$ is a set, $\mathcal{E}$ is a superset of $\mathcal{D}$, $@$ is a (total) function from $\mathcal{E} \times \mathcal{E}$ to $\mathcal{E}$, and $\llbracket _ \rrbracket__$ is a function (parameterised by a semantic context) from λ -terms to $\mathcal{E}$.
- $(\mathcal{E}, \mathcal{D}, \perp, \llbracket _ \rrbracket__)$ is an orthogonality model, where the relation $u \perp \vec{v}$ is defined as $(u @ \vec{v}) \in \mathcal{D}$ (writing $u @ \vec{v}$ for $(\dots (u @ v_1) @ \dots @ v_n)$ if $\vec{v} = v_1 :: \dots v_n :: []$).

※

REMARK 26 Axioms (A1) and (A2) are ensured provided that $\llbracket M N \rrbracket_\rho = \llbracket M \rrbracket_\rho @ \llbracket N \rrbracket_\rho$ and $\llbracket x \rrbracket_\rho = \rho(x)$. These conditions can hold by definition (as in the following example), or can be proved. $\ast$

We now give an applicative orthogonality model to conclude strong normalisation of System F ; this will capture, in essence, the Tait-Girard proof methodology [Tai75, Gir72]. The model is here a *term model*, in that $\mathcal{E}$ is the set of all λ -terms and a λ -term is interpreted as itself.

EXAMPLE 6 (A term-model for Strong Normalisation)

Let $\mathcal{D}$ be the set of strongly-normalising λ -terms, and let $\mathcal{E}$ be set of all λ -terms. We define $u \perp \vec{v}$ as $(u @ \vec{v}) \in \text{SN}$, and the interpretation of terms as follows:

$$\begin{aligned} \llbracket x \rrbracket_\rho &:= \rho(x) \\ \llbracket M_1 M_2 \rrbracket_\rho &:= \llbracket M_1 \rrbracket_\rho @ \llbracket M_2 \rrbracket_\rho \\ \llbracket \lambda x. M \rrbracket_\rho &:= \lambda x. \llbracket M \rrbracket_{\rho, x \mapsto x} \end{aligned}$$

Requirement 3 is a consequence of anti-reduction:

If $\left\{ \frac{P}{x} \right\} M \vec{N} \in \text{SN}$ and $P \in \text{SN}$ then $(\lambda x. M) P \vec{N} \in \text{SN}$.

Note that for all $\vec{N} \in \text{SN}^*$ and all term variables x , $x \perp \vec{N}$.

Hence, for all valuations σ and all types A , $x \in \llbracket A \rrbracket_\sigma^+$.

We apply the Adequacy Lemma (Lemma 25):

If $\Gamma \vdash M : A$, then for all valuation σ and all mapping $\rho \in \llbracket \Gamma \rrbracket_\sigma$ we have $\llbracket M \rrbracket_\rho \in \text{SN}$.

Hence, $M \in \text{SN}$. $\ast$

In summary, we have defined a family of models for the (polymorphically) typed λ -calculus, and presented one instance with which strong normalisation could be inferred. In [BL11b] we presented other instances of orthogonality models, based for instance on intersection types. Unlike usual models (e.g. CCC), orthogonality models do not necessarily equate terms up to β -reduction (if $M \rightarrow_\beta N$, we do not necessarily have $\llbracket M \rrbracket = \llbracket N \rrbracket$). This allows us to build a model where $\llbracket M \rrbracket = M$, from which we can infer strong normalisation of typed terms (an instance of CCC would be useless for this).

2.2 Adapting the approach to classical calculi

Orthogonality was used by Parigot to prove strong normalisation of CBN $\lambda\mu$ -calculus [Par97]. For their non-confluent calculus, Barbanera & Berardi [BB96] adapted the Tait-Girard reducibility technique with “symmetric reducibility candidates”. The key idea in both cases is still that a type A is interpreted as a pair of two orthogonal sets:¹

- a set $\llbracket A \rrbracket^+$ of proof(-terms)
- a set $\llbracket A \rrbracket^-$ of counter-proof(-terms)

...satisfying some saturation property (like reducibility candidates do).

In the proof of strong normalisation of System F that we presented in the previous section, the notion of saturation that holds for $\llbracket A \rrbracket^+$ is that it is closed under bi-orthogonal ($\llbracket A \rrbracket^{+\perp\perp} = \llbracket A \rrbracket^+$). In particular, in an applicative term model, the fact that $\llbracket A \rrbracket^+$ is closed under bi-orthogonal allows to derive, from axiom (A3) of the orthogonality relation, the property that

¹Two sets $\mathcal{U}$ and $\mathcal{V}$ are *orthogonal* if $\forall t \in \mathcal{U}, \forall u \in \mathcal{V}, t \perp u$.

if $(\{N/x\} M) N_1 \dots N_n \in \llbracket A \rrbracket^+$ and $N, N_1, \dots, N_n \in \mathcal{D}$, then $(\lambda x.M) N N_1 \dots N_n \in \llbracket A \rrbracket^+$. Although not explicitly used in our proof of strong normalisation (Example 6), this property lies in the background and is often explicitly used in more traditional presentations of the reducibility technique [Tai75, Gir72]. In brief, the technique works because the interpretation of a type is closed under “head-anti-reduction”.

This is also the approach for classical proof-term calculi, in particular for a confluent calculus such as the Parigot’s $\lambda\mu$ [Par97].

2.2.1 The case of a confluent calculus

In this section we take the example of the LK^N fragment of System L (Definition 22), with $\Rightarrow$ as the only connective, and considering the reduction relation CBN. We can also prove strong normalisation by building a term model based on orthogonality:

Three rewrite rules apply:

$$\begin{array}{lll} (\rightarrow) & \langle \lambda x.t_1 \mid t_2 :: E \rangle & \longrightarrow \langle \{t_2/x\} t_1 \mid E \rangle \\ (\mu_N) & \langle \mu\beta.c \mid E \rangle & \longrightarrow \langle \{E/\beta\} c \rangle \\ (\mu) & \langle t \mid \mu x.c \rangle & \longrightarrow \langle t/x \rangle c \end{array}$$

Therefore we can adapt the axiom (A3) of Definition 26 as follows:

DEFINITION 30 (Orthogonality model for System LK^N)

An *orthogonality model* for System LK^N is given by $(\mathcal{D}_t, \mathcal{D}_e, \perp)$ where $\mathcal{D}_t$ is a set of terms, $\mathcal{D}_e$ is a set of continuations, and $\perp$ is a relation between $\mathcal{D}_t$ and $\mathcal{D}_e$, which can be seen as a set of commands, and satisfying the following *saturation requirements*:

If $\langle t_1 \mid e \rangle (\rho, x \mapsto t_2) \in \perp$ and $t_2 \in \mathcal{D}_t$ then $\langle \lambda x.t_1 \mid t_2 :: e \rangle \rho \in \perp$

If $c(\rho, \beta \mapsto E) \in \perp$ and $E \in \mathcal{D}_e$ then $\langle \mu\beta.c \mid E \rangle \rho \in \perp$

If $c(\rho, x \mapsto t) \in \perp$ and $t \in \mathcal{D}_t$ then $\langle t \mid \mu x.c \rangle \rho \in \perp$

where ρ is a *semantic context*, i.e. a substitution mapping term variables to terms in $\mathcal{D}_t$ and continuation variables to continuations in $\mathcal{D}_e$, and $c\rho$ denotes the capture-avoiding application, to c , of the substitution ρ . ※

DEFINITION 31 (Interpretation of types for System LK^N)

A *valuation* is a function, denoted $\sigma, \sigma', \dots$, from type variables to subsets of $\mathcal{D}_e$ that only contain value continuations.

Two interpretation of types are defined by simultaneous induction of types, a *positive interpretation* and a *negative interpretation*:

$$\begin{array}{ll} [\alpha]_\sigma^- & := \sigma(\alpha) \\ [A \rightarrow B]_\sigma^- & := \llbracket A \rrbracket_\sigma^+ :: \llbracket B \rrbracket_\sigma^- \\ \llbracket A \rrbracket_\sigma^+ & := [A]_\sigma^{-\perp} \quad \llbracket A \rrbracket_\sigma^- := [A]_\sigma^{-\perp\perp} \end{array}$$

where $X :: Y$ denotes $\{t :: E \mid t \in X, E \in Y\}$ for any $X \subseteq \mathcal{D}_t$ and $Y \subseteq \mathcal{D}_e$.

We then define the interpretation of a typing context (i.e. a pair of a typing context for term variables and a typing context for continuation variables):

$$\llbracket \Gamma, \Delta \rrbracket_\sigma := \{ \rho \mid \forall (x:A) \in \Gamma, \rho(x) \in \llbracket A \rrbracket_\sigma^+, \text{ and } \forall (\alpha:A) \in \Delta, \rho(\alpha) \in \llbracket A \rrbracket_\sigma^- \}$$

※

LEMMA 27 (Adequacy Lemma for System $\mathbf{LK}^N$)

1. If $\Gamma \vdash_{\mathbf{L}} t : A ; \Delta$, then for all valuations σ and for all $\rho \in \llbracket \Gamma, \Delta \rrbracket_{\sigma}$ we have $t\rho \in \llbracket A \rrbracket_{\sigma}^+$.
 2. If $\Gamma ; e : A \vdash_{\mathbf{L}} \Delta$, then for all valuations σ and for all $\rho \in \llbracket \Gamma, \Delta \rrbracket_{\sigma}$ we have $e\rho \in \llbracket A \rrbracket_{\sigma}^-$.
 3. If $c : (\Gamma \vdash_{\mathbf{L}} \Delta)$, then for all valuations σ and for all $\rho \in \llbracket \Gamma, \Delta \rrbracket_{\sigma}$ we have $c\rho \in \perp$.
- where $t\rho$, $e\rho$, and $c\rho$ denotes the capture-avoiding application of ρ , seen as a substitution, to t , e , and c , respectively. *

Proof: By simultaneous induction on the typing derivations, using the axioms about $\perp$. □

EXAMPLE 7 (Strong Normalisation of System $\mathbf{LK}^N$)

We define $\mathcal{D}_t$ to be the set of strongly normalising terms and $\mathcal{D}_e$ to be the set of strongly normalising continuations. We define the orthogonality relation $\perp$ between $\mathcal{D}_t$ and $\mathcal{D}_e$ as those commands that are strongly normalising.²

We can check that the saturation requirements are met by purely syntactical/rewriting reasoning, but it only works because there is at most one way to reduce the top-level command.

Take σ to map every type variable to $\mathcal{D}_e$. Notice that term variables are in every $\llbracket A \rrbracket_{\sigma}^+$ and continuation variables are in every $\llbracket A \rrbracket_{\sigma}^-$, and that the identity substitution ρ is in every $\llbracket \Gamma, \Delta \rrbracket_{\sigma}$.

We then apply the Adequacy Lemma with σ and ρ , and get that every typed term, continuation, and command is strongly normalising for $\rightarrow_{\text{CBN}}$. *

In this section, we have proved the strong normalisation of the confluent calculus $\mathbf{LK}^N$ for classical logic, a CBN-fragment of System L. We could have done it along the same lines for the full syntax of System L (but still with the confluent reduction $\rightarrow_{\text{CBN}}$), but dealing with the extra constructs and extra reductions (ζ_N) would have meant a heavier machinery (along the lines of [MM09, CMM10, MM13]). We aimed instead at simplicity, which emphasises the connection with the orthogonality models for System F, and those that we present in the next section.

In summary, in a confluent calculus such as the $\mathbf{LK}^N$, building the positive and negative interpretations of a type A can be described as follows:

	Sets of terms	Sets of continuations
Stage 1		$Y_0 := \llbracket A \rrbracket^-$
Stage 2	$X_1 := Y_0^{\perp}$	$Y_1 := Y_0^{\perp\perp}$
Finished	$\llbracket A \rrbracket^+ := X_1$	$\llbracket A \rrbracket^- := Y_1$

The construction is finished in 2 steps, because the sets X_1 and Y_1 , which are orthogonal, already have all of the saturation properties required to contain all the terms and continuations of type A , which is checked when proving the Adequacy Lemma.

In other words, closure under bi-orthogonality provides adequate saturation properties.

2.2.2 The case of a non-confluent calculus

Let us now consider the situation of a non-confluent calculus such as System L with its original reduction system

²The notion of strong normalisation in the definition of $\mathcal{D}_t$, $\mathcal{D}_e$ and $\perp$ is of course considered for $\rightarrow_{\text{CBN}}$.

$$\begin{array}{lll}
(\rightarrow) & \langle \lambda x.t_1 \mid t_2 :: e \rangle & \longrightarrow \langle t_2 \mid \mu x.\langle t_1 \mid e \rangle \rangle \\
(\overleftarrow{\mu}) & \langle \mu\beta.c \mid e \rangle & \longrightarrow \{e/\beta\}c \\
(\overrightarrow{\mu}) & \langle t \mid \mu x.c \rangle & \longrightarrow \{t/x\}c
\end{array}$$

The Adequacy Lemma might still work if we had the saturation requirements:

If $\langle t_2 \mid \mu x.\langle t_1 \mid e \rangle \rangle \rho \in \perp$ then $\langle \lambda x.t_1 \mid t_2 :: e \rangle \rho \in \perp$

If $c(\rho, \beta \mapsto e) \in \perp$ and $e \in \mathcal{D}_e$ then $\langle \mu\beta.c \mid e \rangle \rho \in \perp$

If $c(\rho, x \mapsto t) \in \perp$ and $t \in \mathcal{D}_t$ then $\langle t \mid \mu x.c \rangle \rho \in \perp$

But in any case, because of non-confluence, these requirements are not met if we define $\mathcal{D}_t$ to be the set of strongly normalising terms and $\mathcal{D}_e$ to be the set of strongly normalising continuations, and $\perp$ the set of strongly normalising commands.³

This means that because of non-confluence, we need to change our notion of saturation, so that $\llbracket A \rrbracket^+$ and $\llbracket A \rrbracket^-$ respectively contain enough terms and continuations for the Adequacy Lemma to hold, and because of that change, the pair $(\llbracket A \rrbracket^+, \llbracket A \rrbracket^-)$ will not be constructed in 2 steps as in the confluent case, but in infinitely many steps:

	Sets of terms		Sets of continuations	
Stage 1	X_0	$\perp$	Y_0	not saturated
Stage 2	X_1	$\perp$	Y_1	not saturated
Stage 3	X_2	$\perp$	Y_2	not saturated
	$\dots$	$\perp$	$\dots$	$\dots$
Stage ∞	X_∞	$\perp$	Y_∞	saturated
Finished	$\llbracket A \rrbracket^+$	$\perp$	$\llbracket A \rrbracket^-$	saturated

We get a saturated pair of sets in infinitely many steps (via a fixpoint construct). In [LM08], we showed that the fixpoint construct could not be captured by a bi-orthogonal completion step.

We now see the details of the technique. In the rest of this section, we fix $\perp$ to be the set of strongly normalising commands.

DEFINITION 32 (Orthogonality and saturation)

Let $\mathbf{Lab}_t$ denote the set of term variables and $\mathbf{Lab}_e$ denote the set of continuation variables. Given a set $\mathcal{U}$ of terms and a set $\mathcal{V}$ of continuations, the pair $(\mathcal{U}, \mathcal{V})$ is

- *orthogonal* if $\forall t \in \mathcal{U}, \forall u \in \mathcal{V}, t \perp u$
- *saturated* if the following two conditions hold
 1. $\mathbf{Lab}_t \subseteq \mathcal{U}$ and $\mathbf{Lab}_e \subseteq \mathcal{V}$
 2. $\{\mu\alpha.c \mid \forall e \in \mathcal{V}, \{e/\alpha\}c \in \perp\} \subseteq \mathcal{U}$ and $\{\mu x.c \mid \forall t \in \mathcal{U}, \{t/x\}c \in \perp\} \subseteq \mathcal{V}$.

A set of terms (resp. continuations) is said to be *simple* if it is non-empty and it contains no term of the form $\mu\alpha.c$ (resp. $\mu x.c$).

For every set X of terms (set Y of continuations), we define a function

$$\Phi_X(\mathcal{W}) := X \cup \mathbf{Lab}_t \cup \{\mu\alpha.c \mid \forall e \in \mathcal{W}, \{e/\alpha\}c \in \perp\}$$

resp.

$$\Phi_Y(\mathcal{W}) := Y \cup \mathbf{Lab}_e \cup \{\mu x.c \mid \forall t \in \mathcal{W}, \{t/x\}c \in \perp\}$$

※

³The notion of strong normalisation in the definition of $\mathcal{D}_t$, $\mathcal{D}_e$ and $\perp$ is now considered for the full reduction relation $\longrightarrow$.

LEMMA 28 Given a set of terms X_0 and a set of continuations Y_0 ,

1. Φ_{X_0} and Φ_{Y_0} are anti-monotonic.⁴
2. Hence, $\Phi_{X_0} \circ \Phi_{Y_0}$ is monotonic and admits a fixpoint X_∞ , with $\Phi_{X_0}(\Phi_{Y_0}(X_\infty)) = X_\infty$.
3. Writing Y_∞ for $\Phi_{Y_0}(X_\infty)$, we clearly have

$$X_\infty = X \cup \text{Lab}_t \cup \{\mu\alpha.c \mid \forall e \in Y_\infty, \{e/\alpha\} c \in \perp\}$$

$$Y_\infty = Y \cup \text{Lab}_e \cup \{\mu x.c \mid \forall t \in X_\infty, \{t/x\} c \in \perp\}$$
4. So (X_∞, Y_∞) is saturated, and a quick case analysis shows that it is orthogonal if X_0 and Y_0 are simple and orthogonal to each other.
5. Finally, $X_0 \subseteq X_\infty$ and $Y_0 \subseteq Y_\infty$.

We finally define $\text{satur}(X_0, Y_0)$ as (X_∞, Y_∞) . ※

DEFINITION 33 (Interpretation of types for System L)

A *valuation* is a function, denoted $\sigma, \sigma', \dots$, from type variables to orthogonal pairs of simple sets.

Two interpretation of types are defined by simultaneous induction of types, a *positive interpretation* and a *negative interpretation*:

$$\begin{aligned} ([a]_\sigma^+, [a]_\sigma^-) &:= \sigma(\alpha) \\ ([A \rightarrow B]_\sigma^+, [A \rightarrow B]_\sigma^-) &:= (\{t \in ([A]_\sigma^+ :: [B]_\sigma^-)^\perp \mid t \text{ not of the form } \mu\alpha.c, [A]_\sigma^+ :: [B]_\sigma^-\} \\ &\quad \text{where } X :: Y \text{ denotes } \{t :: e \mid t \in X, e \in Y\}) \\ ([A]_\sigma^+, [A]_\sigma^-) &:= \text{satur}([A]_\sigma^+, [A]_\sigma^-) \end{aligned}$$

Again, we define

$$[\Gamma, \Delta]_\sigma := \{\rho \mid \forall (x:A) \in \Gamma, \rho(x) \in [A]_\sigma^+, \text{ and } \forall (\alpha:A) \in \Delta, \rho(\alpha) \in [A]_\sigma^-\}$$
※

Now notice the difference with Definition 31: the definition of $[A \rightarrow B]_\sigma^-$ is the same but if we just took $[A \rightarrow B]_\sigma^+$ to be its orthogonal, the pair $([A \rightarrow B]_\sigma^+, [A \rightarrow B]_\sigma^-)$ would not be saturated, as we have already seen; so instead we restrict $[A \rightarrow B]_\sigma^-$ to those terms that are not of the form $\mu\alpha.c$, and thus form an orthogonal (but not saturated) pair of simple sets $([A \rightarrow B]_\sigma^+, [A \rightarrow B]_\sigma^-)$. Then we saturate that pair into $([A \rightarrow B]_\sigma^+, [A \rightarrow B]_\sigma^-)$, which is orthogonal and saturated:

LEMMA 29 (Interpretations of types are orthogonal and saturated)

For all valuations σ and all types A , $([A]_\sigma^+, [A]_\sigma^-)$ is orthogonal and saturated. ※

The rest is now just as in the CBN case:

LEMMA 30 (Adequacy Lemma for System L)

1. If $\Gamma \vdash_L t : A ; \Delta$, then for all valuations σ and for all $\rho \in [\Gamma, \Delta]_\sigma$ we have $t\rho \in [A]_\sigma^+$.
2. If $\Gamma ; e : A \vdash_L \Delta$, then for all valuations σ and for all $\rho \in [\Gamma, \Delta]_\sigma$ we have $e\rho \in [A]_\sigma^-$.
3. If $c : (\Gamma \vdash_L \Delta)$, then for all valuations σ and for all $\rho \in [\Gamma, \Delta]_\sigma$ we have $c\rho \in \perp$.

where $t\rho$, $e\rho$, and $c\rho$ denotes the capture-avoiding application of ρ , seen as a substitution, to t , e , and c , respectively. ※

Proof: By simultaneous induction on the typing derivations, using the Lemma 29. □

⁴In other words for Φ_{X_0} , if $\mathcal{W} \subseteq \mathcal{W}'$ then $\Phi_{X_0}(\mathcal{W}) \supseteq \Phi_{X_0}(\mathcal{W}')$. And similarly for Φ_{Y_0} .

THEOREM 31 (Strong Normalisation of System L)

Take σ to map every type variable to the orthogonal pair $(\text{Lab}_t, \text{Lab}_e)$ of simple sets. Notice again that the identity substitution ρ is in every $\llbracket \Gamma, \Delta \rrbracket_\sigma$.

We then apply the Adequacy Lemma with σ and ρ , and get that every typed term, continuation, and command is strongly normalising for $\longrightarrow$. *

The points to remember are

- As for System F , we have proved strong normalisation by building a term model
 - which does not equate terms up to reduction (non-confluence would make that very problematic)
 - where axiom (A3) is replaced by a saturation property.
- Because of non-confluence,
 - saturation has to be a property of pairs $(\llbracket A \rrbracket_\sigma^+, \llbracket A \rrbracket_\sigma^-)$, not a property of each component separately;
 - saturating is difficult (adding terms in one component of the pair affects the other component), and obtained by a fixpoint construction.

As shown in [LM08], the saturation process is not just a bi-orthogonality completion: if $(\mathcal{U}, \mathcal{V})$ is orthogonal, then $(\mathcal{U}^{\perp\perp}, \mathcal{V}^{\perp\perp})$ is orthogonal but not necessarily saturated.

2.3 Orthogonality models for extracting witnesses from classical proofs

We now show how to extract a witness from a classical proof of a Σ_1^0 -formula, i.e. a closed formula of the form $\exists a A(a)$ where $A(a)$ is a quantifier-free formula of arithmetics.

The technique is due to Miquel [Miq09, Miq11], we simply adapted it to our proof-term calculus for classical logic, and somewhat simplified it using the concepts and notations of the previous sections.

We work in a particular setting where such a formula is expressed in the shape of $\neg \forall a \neg \text{isnull}(e(a))$, the grammar of formulae being defined as follows:

DEFINITION 34 (Expressions and Formulae)

$$\begin{array}{ll} \textbf{Expressions} & u, u', \dots ::= a \mid 0 \mid s(u) \mid u + u' \mid u \times u' \mid u \leq u' \\ \textbf{Formulae} & A, B, \dots ::= \text{isnull}(u) \mid A \rightarrow B \mid \forall a A \end{array}$$

We represent integers as expressions: let $\bar{0} := 0$ and, for all integers n , let $\overline{n+1} := s(\bar{n})$.

We define $\neg A := A \rightarrow \text{isnull}(\bar{1})$. *

This shape for a Σ_1^0 -formula brings no loss of generality. Moreover, such an expression as $u(a)$, with one free variable a , expresses a primitive recursive function from $\mathbb{N}$ to $\mathbb{N}$.

We will now build an orthogonality model that we will use for witness extraction. As in the previous sections, each formula A will be interpreted as a set $\llbracket A \rrbracket_\sigma^+$ of terms and a set $\llbracket A \rrbracket_\sigma^-$ of continuations, terms and continuations being those of LK^N .

The extraction mechanism itself will be given by the reductions of LK^N , and more precisely by **root** CBN-reduction, which we denote $\longrightarrow_{\text{CBNr}}$.⁵

In other words, from a proof of $\neg\forall a.\neg\text{isnull}(u(a))$ in (the extension to arithmetic of) System $\mathbf{L}$, we will perform $\longrightarrow_{\text{CBNr}}$ -reduction until we reach (in a provably finite number of steps) a command where we can directly read a witness.

For this we need to express numerals as proof-terms. We simply use Church's numerals in λ -calculus (see e.g. [Bar84]) and encode them in LK^N with Definition 14:

DEFINITION 35 (Church's numerals as terms)

$$\begin{aligned} c_0 &:= \langle x \mid \alpha \rangle \\ c_{n+1} &:= \langle f \mid (\mu\alpha.c_n)::\alpha \rangle \\ \underline{n} &:= \lambda x.\lambda f.\mu\alpha.c_n \end{aligned}$$

※

REMARK 32 Doing the same thing with the λ -terms for the successor function and the recursion function, we get two LK^N terms **s** and **rec** such that, for all t , u_0 , u_1 , for all continuation values E , and all integer n ,

$$\begin{aligned} \langle \mathbf{s} \mid \underline{n}::t::E \rangle &\longrightarrow_{\text{CBNr}}^* \langle t \mid \underline{n+1}::E \rangle \\ \langle \mathbf{rec} \mid u_0::u_1::\underline{0}::E \rangle &\longrightarrow_{\text{CBNr}}^* \langle u_0 \mid E \rangle \\ \langle \mathbf{rec} \mid u_0::u_1::\underline{n+1}::E \rangle &\longrightarrow_{\text{CBNr}}^* \langle u_1 \mid \underline{n}::(\mu\alpha.\langle \mathbf{rec} \mid u_0::u_1::\underline{n}::\alpha \rangle)::E \rangle \end{aligned}$$

using the simulation of β -reduction by $\longrightarrow_{\text{CBN}}$.⁶

Let **ifz** := $\lambda n x_0 x_1.\mu\alpha.\langle \mathbf{rec} \mid x_0::(\lambda y_0 y_1.x_1)::n::\alpha \rangle$, so that

$$\begin{aligned} \langle \mathbf{ifz} \mid \underline{0}::u_0::u_1::E \rangle &\longrightarrow_{\text{CBNr}}^* \langle u_0 \mid E \rangle \\ \langle \mathbf{ifz} \mid \underline{n+1}::u_0::u_1::E \rangle &\longrightarrow_{\text{CBNr}}^* \langle u_1 \mid E \rangle \end{aligned}$$

※

DEFINITION 36 (Orthogonality semantics)

Let $\perp$ be an arbitrary set of commands, stable under anti-reduction (if $c \longrightarrow_{\text{CBNr}} c'$ and $c' \in \perp$ then $c \in \perp$).

A valuation σ is a mapping from expression variables (a , etc) to integers.

Given a valuation σ , Fig. 11 defines the interpretation of an expression u as an integer $\llbracket u \rrbracket_\sigma$ and a formula A as a set $\llbracket A \rrbracket_\sigma^+$ of terms and a set $\llbracket A \rrbracket_\sigma^-$ of continuations.

※

REMARK 33

1. Clearly, $\llbracket \bar{n} \rrbracket_\sigma = n$ for all n and σ .
2. By induction on u we get $\llbracket \{\bar{n}/a\} u \rrbracket_\sigma = \llbracket u \rrbracket_{\sigma, a \mapsto n}$, and by induction on A we get $\llbracket \{\bar{n}/a\} A \rrbracket_\sigma^- = \llbracket A \rrbracket_{\sigma, a \mapsto n}^-$ and $\llbracket \{\bar{n}/a\} A \rrbracket_\sigma^+ = \llbracket A \rrbracket_{\sigma, a \mapsto n}^+$.

※

Now, for simplicity we do not specify the exact proof system for arithmetic, nor do we give a typing system corresponding to it through the Curry-Howard correspondence. We assume that it could be built as an extension of Fig. 6, and that the Adequacy Lemma can be proved (along the lines of Lemma 27 for LK^N):

⁵The fact that we use CBN-reduction is important to make sure that reduction **can** produce a witness; the fact that we only use root reduction is not, but in order to implement the extraction mechanism deterministically, it is convenient to never have to choose the next redex to reduce.

⁶And the fact that we can do this with root-reduction only is rather clear.

$$\begin{aligned}
\llbracket a \rrbracket_\sigma &:= \sigma(a) \\
\llbracket 0 \rrbracket_\sigma &:= 0 \\
\llbracket \mathbf{s}(u) \rrbracket_\sigma &:= \llbracket u \rrbracket_\sigma + 1 \\
\llbracket u_1 + u_2 \rrbracket_\sigma &:= \llbracket u_1 \rrbracket_\sigma + \llbracket u_2 \rrbracket_\sigma \\
\llbracket u_1 \times u_2 \rrbracket_\sigma &:= \llbracket u_1 \rrbracket_\sigma \times \llbracket u_2 \rrbracket_\sigma \\
\llbracket u_1 \leq u_2 \rrbracket_\sigma &:= 1 && \text{if } \llbracket u_1 \rrbracket_\sigma \leq \llbracket u_2 \rrbracket_\sigma \\
\llbracket u_1 \leq u_2 \rrbracket_\sigma &:= 0 && \text{if } \llbracket u_1 \rrbracket_\sigma > \llbracket u_2 \rrbracket_\sigma \\
\llbracket \mathbf{isnull}(u) \rrbracket_\sigma^- &:= \mathcal{E} && \text{if } \llbracket u \rrbracket_\sigma \neq 0 \\
\llbracket \mathbf{isnull}(u) \rrbracket_\sigma^- &:= \emptyset && \text{if } \llbracket u \rrbracket_\sigma = 0 \\
\llbracket A \rightarrow B \rrbracket_\sigma^- &:= \llbracket A \rrbracket_\sigma^+ :: \llbracket B \rrbracket_\sigma^- \\
\llbracket \forall a A \rrbracket_\sigma^- &:= \bigcup_{n \in \mathbb{N}} (\{ \underline{n} \} :: \llbracket A \rrbracket_{\sigma, a \rightarrow n}^-)
\end{aligned}$$

$$\llbracket A \rrbracket_\sigma^+ := \llbracket A \rrbracket_\sigma^{-\perp} \quad \llbracket A \rrbracket_\sigma^- := \llbracket A \rrbracket_\sigma^{-\perp\perp}$$

where $\mathcal{E}$ is the set of all continuation values and $X :: Y$ denotes $\{t :: E \mid t \in X, E \in Y\}$.

Figure 11: Semantics of expressions and formulae

A closed proof t_0 of a formula $\neg \forall a \neg \mathbf{isnull}(u(a))$ is such that,
for all possible $\perp$ closed under “anti-reduction” (the inverse of $\longrightarrow_{\text{CBNr}}$),
 $t_0 \in \llbracket \neg \forall a \neg \mathbf{isnull}(u(a)) \rrbracket$.

We thus start with such a term t_0 .

We now define a term that can check whether an integer is a witness of the property and, depending on this check, continue with one term or another:

DEFINITION 37 (Witness checker)

Let f be the primitive recursive function defined by: for any integer n , $f(n) := \llbracket u(a) \rrbracket_{a \rightarrow n}$.
Let $\underline{f}$ be a term representing f in the sense that, for any integer n , and term t and any continuation value E ,

$$\langle \underline{f} \mid \underline{n} :: t :: E \rangle \longrightarrow_{\text{CBNr}}^* \langle t \mid \underline{f}(n) :: E \rangle$$

Such a term can be constructed from $\mathbf{s}$ and $\mathbf{rec}$, as the projections, composition, etc are all available in System L.

We define the *witness checker* as follows:

$$d_f := \lambda nxy. \mu \alpha. \langle \underline{f} \mid n :: (\lambda p. \mu \alpha_1. \langle \mathbf{ifz} \mid p :: x :: y :: \alpha_1 \rangle) :: \alpha \rangle$$

※

LEMMA 34 (Witness checker property)

For any integer n , any u_0 and u_1 and E , we have

$$\langle d_f \mid \underline{n} :: u_0 :: u_1 :: E \rangle \longrightarrow_{\text{CBNr}}^* \langle u_0 \mid E \rangle \text{ if } f(n) = 0$$

$$\langle d_f \mid \underline{n} :: u_0 :: u_1 :: E \rangle \longrightarrow_{\text{CBNr}}^* \langle u_1 \mid E \rangle \text{ if } f(n) \neq 0$$

※

Proof:

$$\begin{aligned}
\langle d_f \mid \underline{n}::u_0::u_1::E \rangle &\longrightarrow_{\text{CBNr}}^* \langle \mu\alpha.\langle \underline{f} \mid \underline{n}::(\lambda p.\mu\alpha_1.\langle \mathbf{ifz} \mid p::u_0::u_1::\alpha_1 \rangle)::\alpha \rangle \mid E \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle \underline{f} \mid \underline{n}::(\lambda p.\mu\alpha_1.\langle \mathbf{ifz} \mid p::u_0::u_1::\alpha_1 \rangle)::E \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle \lambda p.\mu\alpha_1.\langle \mathbf{ifz} \mid p::u_0::u_1::\alpha_1 \rangle \mid \underline{f(n)}::E \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle \mu\alpha_1.\langle \mathbf{ifz} \mid \underline{f(n)}::u_0::u_1::\alpha_1 \rangle \mid E \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle \mathbf{ifz} \mid \underline{f(n)}::u_0::u_1::E \rangle
\end{aligned}$$

If $f(n) = 0$, this reduces to $\langle u_0 \mid E \rangle$. Otherwise, this reduces to $\langle u_1 \mid E \rangle$. $\square$

DEFINITION 38 (Orthogonality and contradicter)

Let **stop** be an arbitrary term and **go** be an arbitrary continuation.

We now take a particular orthogonality set defined by

$$\perp := \{c \mid \text{there exists } n \text{ such that } f(n) = 0 \text{ and } c \longrightarrow_{\text{CBNr}}^* \langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle\}$$

It is closed under anti-CBNr-reduction.

We now define a “contradicter”:⁷ Let $t_1 := \lambda n x. \mu\alpha. \langle d_f \mid n::(\mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle)::x::\alpha \rangle$. $\ast$

LEMMA 35 (Behaviour of the contradicter)

For all integers n , and all continuation values E in $\llbracket \neg \text{isnull}(u(\bar{n})) \rrbracket^-$, we have $t_1 \perp \underline{n}::E$. $\ast$

Proof: We have

$$\langle t_1 \mid \underline{n}::E \rangle \longrightarrow_{\text{CBNr}}^* \langle \lambda x. \mu\alpha. \langle d_f \mid \underline{n}::(\mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle)::x::\alpha \rangle \mid E \rangle$$

To prove that this is an orthogonal command, we only have to show, as $E \in \llbracket \neg \text{isnull}(u(\bar{n})) \rrbracket^{-\perp\perp}$, that the left-hand side term is orthogonal to every continuation in $\llbracket \neg \text{isnull}(u(\bar{n})) \rrbracket^-$.

Consider such a continuation; it is of the form $t::E'$ with $t \in \llbracket \text{isnull}(u(\bar{n})) \rrbracket^+$.

If $f(n) \neq 0$ then $[u(\bar{n})] \neq 0$, so $\llbracket \text{isnull}(u(\bar{n})) \rrbracket^- = \mathcal{E}$ and t is orthogonal to every continuation, in particular E' . So we have

$$\begin{aligned}
&\langle \lambda x. \mu\alpha. \langle d_f \mid \underline{n}::(\mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle)::x::\alpha \rangle \mid t::E' \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle d_f \mid \underline{n}::(\mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle)::t::E' \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle t \mid E' \rangle
\end{aligned}$$

which is in $\perp$.

If $f(n) = 0$ then we have

$$\begin{aligned}
&\langle \lambda x. \mu\alpha. \langle d_f \mid \underline{n}::(\mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle)::x::\alpha \rangle \mid t::E' \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle d_f \mid \underline{n}::(\mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle)::t::E' \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle \mu\alpha_0.\langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle \mid E' \rangle \\
&\longrightarrow_{\text{CBNr}}^* \langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle
\end{aligned}$$

$\square$

COROLLARY 36 (Classical witness extraction)

$\langle t_0 \mid t_1::\mathbf{go} \rangle \longrightarrow_{\text{CBNr}}^* \langle \mathbf{stop} \mid \underline{n}::\mathbf{go} \rangle$ for some integer n such that $f(n) = 0$. $\ast$

Proof:

From Lemma 35 we get that $t_1 \in \llbracket \forall a. \neg \text{isnull}(u(a)) \rrbracket^+$ and therefore $t_1::\mathbf{go} \in \llbracket \neg \forall a. \neg \text{isnull}(u(a)) \rrbracket^-$.

Since we have assumed $t_0 \in \llbracket \neg \forall a. \neg \text{isnull}(u(a)) \rrbracket^+$, we have $t_0 \perp t_1::\mathbf{go}$, from which we conclude. $\square$

⁷In the sense that it will contradict what the proof of t_0 claims.

In other words, once given a classical proof, we match it against the continuation $t_1 :: \mathbf{go}$ and we are certain that $\mathbf{CBNr}$ will produce $\langle \mathbf{stop} \mid \underline{n} :: \mathbf{go} \rangle$ in a finite number of steps, with n being a witness.

For a comparison with other techniques of classical witness extraction, see [Miq09, Miq11].

Conclusion

In summary, we have seen in this chapter a fundamental concept for model construction, namely orthogonality. We built several orthogonality models for various purposes: rephrase strong normalisation proofs for System F , prove the strong normalisation of a confluent proof-term calculus for classical logic as well as a non-confluent calculus (thereby proving cut-elimination), and finally extract witnesses from classical proofs of Σ_1^0 -formulae.

In each of those model constructions, we have interpreted formulae first with basic inhabitants (terms or continuations), and then closed their interpretation by a completion process that could simply be taking the bi-orthogonal, in the case of confluent calculi, or a more complex fixpoint, in the case of a non-confluent calculus.

Whether in those constructions we first define the negative interpretation of a formula (as a set of “counter-proofs”) or its positive interpretation (as a set of “proofs”), is a question that depends on the formula’s *polarity*. This is the topic of the next chapter.

Chapter 3

Polarisation and focussing

Contents

3.1 Recovering confluence by polarisation	66
3.1.1 Symmetry, asymmetry, and η -expansions	66
3.1.2 Towards polarised System L	68
3.1.3 Focussing	71
3.1.4 Weak η -conversion	72
3.1.5 Related works	73
3.2 Computational interpretation of a focussed calculus	74
3.2.1 Informal relation to System L	76
3.2.2 Identifying phases as atomic steps	77
3.2.3 Functional interpretation as pattern-matching	82
Conclusion	84

In the previous chapters, we have seen that

- A proof-term syntax, together with a typing system, can be used to represent classical proofs (e.g. System L), such that the symmetry of classical logic is reflected by the symmetry between programs and continuations. The use of classical reasoning corresponds to letting a program capture its continuation.
- A rewrite system on proof-terms can be given to represent cut-elimination, following the intuitions of continuations and control. This gives a non-confluent calculus because (unrestricted) cut-elimination is non-confluent in classical logic, reflected by the fact that programs and continuations compete for the control of computation.
- Still, the rewrite system is strongly normalising on typed proof-terms (i.e. those representing real proofs), showing that cuts are admissible. The proof of strong normalisation was the occasion to introduce orthogonality techniques, although non-confluence requires more, namely specific saturation properties.
- The semantics of classical proofs, or typed proof-terms, is problematic until confluence is recovered in some way.

Back to the main issue, a CCC with $\neg\neg A \simeq A$ collapses, and out of the 3 natural ways to avoid the collapse, namely

1. break the symmetry between $\wedge$ and $\vee$,

2. break the cartesian product,
3. break the curryfication,

we investigate the breaking of symmetry between $\wedge$ and $\vee$.

In Chapter 1 we saw how to break the $\wedge\vee$ -symmetry by the CBV/CBN approach. In this chapter, we break the $\wedge\vee$ symmetry by *polarisation*.

3.1 Recovering confluence by polarisation

3.1.1 Symmetry, asymmetry, and η -expansions

We start this section by coming back to a fundamental question: What is symmetrical about Classical Logic? There is definitely a symmetry based on the duality of negation / De Morgan's duality. It can be seen in the truth semantics of formulae, in e.g. truth tables or more generally in a boolean algebra: meet / join and top / bottom are swapped when flipping the order upside-down, and all the axioms of a boolean algebra are preserved.

At the level of proofs, there is also a symmetry that can be seen for instance in the two-sided sequent calculus: the left-introduction rule of a connective is symmetric to the right-introduction rule of its dual connective (in other words, the rules are preserved under duality flipping).

Cut-elimination is symmetrical (e.g. the rewrite system of Fig. 6), but to make semantical sense of it, one breaks the symmetry by making a choice between CBN and CBV that is completely arbitrary.

More interestingly, the following example reveals something asymmetric between the left-introduction of $\vee$ and the right-introduction of $\vee$:

$$\frac{\Gamma, A \vdash \Delta \quad \Gamma, B \vdash \Delta}{\Gamma, A \vee B \vdash \Delta} \quad \frac{\Gamma \vdash A, \Delta}{\Gamma \vdash A \vee B, \Delta} \quad \frac{\Gamma \vdash B, \Delta}{\Gamma \vdash A \vee B, \Delta}$$

Of course, we have never claimed that there is a symmetry between the left-introduction of $\vee$ and the right-introduction of $\vee$, but an interesting question is raised by the following situation: it is known (see e.g. [TS00]) that in the sequent calculus, the axiom rule

$$\frac{}{A \vdash A}$$

(say in a context-splitting setting) can be restricted, without losing logical completeness, to the *atomic* axiom rule

$$\frac{}{a \vdash a}$$

Every instance of the general instance can be replaced by a small proof only using atomic axioms, which is proved by induction on A : in particular, transforming the axiom $\frac{}{A \vee B \vdash A \vee B}$ into a proof with atomic axioms, we produce

$$\frac{\frac{\dots}{A \vdash A} \quad \frac{\dots}{B \vdash B}}{A \vdash A \vee B \quad B \vdash A \vee B} \quad \frac{}{A \vee B \vdash A \vee B}$$

and then recursively transform $A \vdash A$ and $B \vdash B$ (until all of the used axioms are atomic).

So the interesting question is whether there is a fundamental reason why $\vee$ is decomposed on the left before being decomposed on the right (looking at the bottom-up construction of the proof). Starting the decomposition on the right would have failed.

A related situation occurs with η -expansion in λ -calculus:

In λ -calculus, the use of an axiom corresponds to a variable in the proof-term. Typing the term

$$\lambda z^{A \rightarrow B}.z$$

(where we indicate the types of variables as superscripts) uses an axiom on $A \rightarrow B$. Typing its η -expansion

$$\lambda z^{A \rightarrow B}.\lambda y^A.z y$$

uses, strictly speaking, an axiom on $A \rightarrow B$ and an axiom on A , but as z is immediately applied and its type $A \rightarrow B$ immediately destructed, the η -expansion only uses, “morally” speaking, axioms on the smaller formulae A and B . Turning this moral intuition into something formal can be done by taking a proof-term calculus for sequent calculus (rather than natural deduction), as we shall see below, but still: we first have the λ -abstraction, and underneath it the application. Why again do they have to appear in that order?

Indeed, in a classical calculus such a System L, the axiom on $A \rightarrow B$ is represented as

$$\frac{}{z : A \rightarrow B \vdash z : A \rightarrow B ;}$$

The η -expansion of z is:

$$\frac{\frac{}{z : A \rightarrow B, y : A \vdash z : A \rightarrow B ; \alpha : B} \quad \frac{\frac{}{z : A \rightarrow B, y : A \vdash y : A ; \alpha : B} \quad \frac{}{z : A \rightarrow B, y : A ; \alpha : B \vdash \alpha : B}}{z : A \rightarrow B, y : A ; (y :: \alpha) : A \rightarrow B \vdash \alpha : B}}{\frac{\langle z \mid y :: \alpha \rangle : (z : A \rightarrow B, y : A \vdash \alpha : B)}{z : A \rightarrow B, y : A \vdash \mu \alpha. \langle z \mid y :: \alpha \rangle : B ;} \quad z : A \rightarrow B \vdash \lambda y. \mu \alpha. \langle z \mid y :: \alpha \rangle : A \rightarrow B ;}$$

and then we can recursively transform the axioms on $y : A$ and $\alpha : B$ (until axioms are atomic). Of course, this η -expansion still features the use of $z : A \rightarrow B$, but only to implement a contraction (or even more precisely to implement the placing of the formula $A \rightarrow B$ where it can be decomposed), not to implement a proper axiom.

Now the η -expansion we used in the λ -calculus to illustrate our point is only one particular instance of η -expansion: the general form

$$M \longrightarrow_{\eta} \lambda y^A.M y \quad y \notin \text{FV}(M)$$

can be recovered, by the capture avoiding substitution of M for z , from the axiomatic η -expansion on axiom z :

$$z \longrightarrow_{\eta} \lambda y^A.z y$$

And in λ -calculus, no matter M (where y is not free), M and $\lambda y.M y$ have the same computational behaviour (with respect to β -reduction). In technical terms, M and $\lambda y.M y$ cannot be separated (even using untyped terms) [Bar84].

In System L, the η -expansion on axiom z

$$z \longrightarrow_{\eta} \lambda y. \mu \alpha. \langle z \mid y :: \alpha \rangle$$

also provides, after instantiation of z by t (where y and α are not free), a general form of η -expansion:

$$t \longrightarrow_{\eta} \lambda y. \mu\alpha. \langle t \mid y::\alpha \rangle$$

which transforms

$$\Gamma \vdash t : A \rightarrow B ; \Delta$$

into

$$\frac{\frac{\frac{\Gamma, y : A \vdash y : A ; \alpha : B, \Delta}{\Gamma, y : A \vdash y : A ; (y::\alpha) : A \rightarrow B \vdash \alpha : B, \Delta}}{\Gamma \vdash t : A \rightarrow B ; \Delta} \quad \frac{\frac{\frac{\langle t \mid y::\alpha \rangle : (\Gamma, y : A \vdash \alpha : B, \Delta)}{\Gamma, y : A \vdash \mu\alpha. \langle t \mid y::\alpha \rangle : B ; \Delta}}{\Gamma \vdash \lambda y. \mu\alpha. \langle t \mid y::\alpha \rangle : A \rightarrow B ; \Delta}}$$

3.1.2 Towards polarised System L

Now the above general η -expansion can be instantiated with $t = \mu\beta.c$:

$$\mu\beta.c \longrightarrow_{\eta} \lambda y. \mu\alpha. \langle \mu\beta.c \mid y::\alpha \rangle$$

If we put those two terms in context, e.g. facing a continuation $\mu x.c'$, we get that

- $\langle \mu\beta.c \mid \mu x.c' \rangle$ rewrites to

$$\left\{ \mu\beta.c /_x \right\} c' \quad \text{or} \quad \left\{ \mu x.c' /_{\beta} \right\} c$$

- **but** $\langle \lambda y. \mu\alpha. \langle \mu\beta.c \mid y::\alpha \rangle \mid \mu x.c' \rangle$ rewrites only to

$$\left\{ \lambda y. \mu\alpha. \langle \mu\beta.c \mid y::\alpha \rangle /_x \right\} c'$$

If η -convertible terms should have undistinguishable computational behaviour, we must forbid $\langle \mu\beta^{A_1 \rightarrow A_2}.c \mid \mu x^{A_1 \rightarrow A_2}.c' \rangle \longrightarrow \left\{ \mu x^{A_1 \rightarrow A_2}.c' /_{\beta} \right\} c$

The grounds for breaking the symmetry in such a way is that $\mu\beta^{A_1 \rightarrow A_2}.c$ can be η -expanded, but $\mu x^{A_1 \rightarrow A_2}.c'$ cannot, which reflects the fact that the right-introduction rule for $A_1 \rightarrow A_2$ is invertible while its left-introduction rule is not.

In short, when encountering

$$\frac{\frac{c : (\Gamma \vdash \beta : A_1 \rightarrow A_2, \Delta)}{\Gamma \vdash \mu\beta.c : A_1 \rightarrow A_2 ; \Delta} \quad \frac{c' : (\Gamma, x : A_1 \rightarrow A_2 \vdash \Delta)}{\Gamma ; \mu x.c' : A_1 \rightarrow A_2 \vdash \Delta}}{\langle \mu\beta.c \mid \mu x.c' \rangle : (\Gamma \vdash \Delta)}$$

we could consider that the term $\mu\beta.c$ is a “cheater” in the sense that its type $A_1 \rightarrow A_2$ could be proved or inhabited in another way (e.g. with the η -expansion of $\mu\beta.c$), avoiding the critical pair, and solving the non-confluence problem.

In particular, if β is used 0 times in c , or more than once, we can understand the typing tree

$$\frac{c : (\Gamma \vdash \beta : A_1 \rightarrow A_2, \Delta)}{\Gamma \vdash \mu\beta.c : A_1 \rightarrow A_2 ; \Delta}$$

as finishing with a weakening or a contraction. What η -expansion proves is that the proof can be transformed into a proof that finishes with a proper introduction of the implication.

In our earlier example about the connective $\vee$, it is the contrary: its left-introduction rule is invertible while its right-introduction rules are not.

This leads to considering a notion that arose from linear logic [Gir87]: *polarities*.

The intuition for *positive connectives* is that we expect no particular property of their right-introduction rules. These rules are called *synchronous*. In particular for goal-directed proof-search, applying such a rule bottom-up is *a priori* a choice which we may have to backtrack to if we fail to finish the proof. For logical completeness, (right-)weakenings or (right-)contractions may be necessary on a formula with a positive connective at its root.

The intuition for *negative connectives* is that their right-introduction rules *are* invertible. These rules are called *asynchronous*. In goal-directed proof-search we may apply such rules without loss of generality and therefore without creating backtrack points. Also, (right-)weakenings and (right-)contractions (on formulae that have a negative connective at their roots) are superfluous as far as logical completeness is concerned. On the other hand, the right-introduction rules “must interact well with the left-introduction rules” (or the right-introduction rules of the dual connective), in cut-elimination as well as in the expansion of axioms that we described in this section.

Just as in λ -calculus you can *always* inhabit a (non-empty) function type with a λ -abstraction, you can always η -expand an inhabitant of a type whose main connective is negative. Considering the η -expansion rules that we can apply in System L, we can derive the polarities of the three connectives we considered:

$$\begin{array}{llll} \text{negative} & A \rightarrow B & t & \longrightarrow \lambda y. \mu\alpha.\langle t \mid y::\alpha \rangle \\ \text{negative} & A \wedge B & t & \longrightarrow (\mu\alpha.\langle t \mid \text{inj}_1(\alpha) \rangle, \mu\gamma.\langle t \mid \text{inj}_2(\gamma) \rangle) \\ \text{positive} & A \vee B & e & \longrightarrow (\mu x.\langle \text{inj}_1(x) \mid e \rangle, \mu z.\langle \text{inj}_2(z) \mid e \rangle) \end{array}$$

Now in order to solve the confluence problem, we also need to determine how to reduce $\langle \mu\beta.c \mid \mu x.c' \rangle$ when the cut-formula is atomic. This leads to splitting the set of the atomic formulae into positives and negatives as well. Unlike non-atomic formulae, the choice of polarity for each atom is arbitrary, and sometimes called the *bias* [LM09].

Now we can use these ideas to layer System L with polarities:

DEFINITION 39 (Polarised System L) The polarised syntax of formulae is defined as

$$\begin{array}{ll} P, P', \dots & ::= a^+ \mid A \vee B \\ N, N', \dots & ::= a^- \mid A \wedge B \mid A \rightarrow B \\ A, B, \dots & ::= P \mid N \end{array}$$

The syntax for proof-terms, together with their associated forms of typing judgements, is given below:

$$\begin{array}{llll} \text{-terms} & t^- & ::= x^- \mid \lambda x.t \mid (t_1, t_2) & \mid \mu\beta^-.c & \Gamma \vdash t^- : N ; \Delta \\ \text{+terms} & t^+ & ::= x^+ & \mid \text{inj}_i(t) \mid \mu\beta^+.c & \Gamma \vdash t^+ : P ; \Delta \\ \text{terms} & t & ::= t^+ \mid t^- & & \Gamma \vdash t : A ; \Delta \\ \\ \text{-continuations} & e^- & ::= \alpha^- \mid t::e & \mid \text{inj}_i(e) \mid \mu x^-.c & \Gamma ; e^- : N \vdash \Delta \\ \text{+continuations} & e^+ & ::= \alpha^+ & \mid (e_1, e_2) & \mid \mu x^+.c & \Gamma ; e^+ : P \vdash \Delta \\ \text{continuations} & t & ::= e^+ \mid e^- & & \Gamma ; e : A \vdash \Delta \\ \\ \text{commands} & c & ::= \langle t^+ \mid e^+ \rangle \mid \langle t^- \mid e^- \rangle & & c : (\Gamma \vdash \Delta) \end{array}$$

writing x for either x^+ or x^- .

※

Now that polarities explicitly appear in the syntax of proof-terms, it is easy to reduce $\langle \mu\alpha.c \mid \mu x.c' \rangle$:

$$\langle \mu\alpha^+.c \mid \mu x^+.c' \rangle \longrightarrow \left\{ \mu x^+.c' /_{\alpha^+} \right\} c \quad \text{and} \quad \langle \mu\alpha^-.c \mid \mu x^-.c' \rangle \longrightarrow \left\{ \mu\alpha^-.c /_{x^-} \right\} c'$$

This turns into the following rewrite system:

DEFINITION 40 (Reductions in the polarised System L) Again, we define values:

$$\begin{array}{ll} \text{term values} & V ::= x \mid \text{inj}_i(V) \mid t^- \\ \text{continuation values} & E ::= \alpha \mid V :: E \mid \text{inj}_i(E) \mid e^+ \end{array}$$

The reduction relation $\rightarrow_F$ is defined as the contextual closure of the rules in Fig. 12. $\ast$

$$\begin{array}{lll} (\rightarrow) & \langle \lambda x. t \mid V :: E \rangle & \longrightarrow \langle \{V/x\} t \mid E \rangle \\ (\wedge) & \langle (t_1, t_2) \mid \text{inj}_i(E) \rangle & \longrightarrow \langle t_i \mid E \rangle \\ (\vee) & \langle \text{inj}_i(V) \mid (e_1, e_2) \rangle & \longrightarrow \langle V \mid e_i \rangle \\ \\ (\overleftarrow{\mu}_-) & \langle \mu \beta^-. c \mid E \rangle & \longrightarrow \langle \{E/\beta^+\} c \\ (\overrightarrow{\mu}) & \langle t^- \mid \mu x^-. c \rangle & \longrightarrow \langle \{t^-/x^-\} c \\ (\overleftarrow{\mu}) & \langle \mu \beta^+. c \mid e^+ \rangle & \longrightarrow \langle \{e^+/\beta^+\} c \\ (\overrightarrow{\mu}_+) & \langle V \mid \mu x^+. c \rangle & \longrightarrow \langle \{V/x^+\} c \\ \\ (\zeta_F) & \langle t^- \mid t^+ :: e \rangle & \longrightarrow \langle t^+ \mid \mu x^+. (t^- \mid x^+ :: e) \rangle \\ (\zeta_F) & \langle t^- \mid V :: e^- \rangle & \longrightarrow \langle \mu \alpha. \langle t^- \mid V :: \alpha \rangle \mid e^- \rangle \\ (\zeta_F) & \langle t^- \mid \text{inj}_i(e^-) \rangle & \longrightarrow \langle \mu \alpha. \langle t^- \mid \text{inj}_i(\alpha) \rangle \mid e^- \rangle \\ (\zeta_F) & \langle \text{inj}_i(t^+) \mid e^+ \rangle & \longrightarrow \langle \mu x^+. \langle \text{inj}_i(x^+) \mid e^+ \rangle \mid t^+ \rangle \end{array}$$

where the (ζ_F) -rules apply only under the condition that t^+ and e^- are not values.

Figure 12: Rewrite system for polarised System L

As in the CBN and CBV cases, we have:

THEOREM 37 (Confluence and Subject Reduction)

$\rightarrow_F$ is confluent and satisfies Subject Reduction. $\ast$

Notice that the notion of value is slightly different from that of Definition 21: Indeed, if $\wedge$ is to be taken to be negative, as the dual of (the obviously positive) $\vee$, we can take every pair to be a value (in Definition 21 we stuck to Wadler's presentation [Wad03]); this also removes the need for ζ -rules for pairs. On the other hand, for a continuation $t :: e$ to be a value, we require it to be of the form $V :: E$, as we no longer recover confluence by opposing left vs. right (terms vs. continuations) but by opposing positives vs. negatives.

Precisely because we now no longer make any distinction based on the left vs. right opposition (terms vs. continuations opposition), this system could equally be given as a one-sided system, merging the syntaxes of terms and continuations, but keeping of course the distinction between positive terms and negative terms.¹ At the level of formulae, we would get 4 connectives $\vee^+$ and $\wedge^+$ of positive polarity, and $\vee^-$ and $\wedge^-$ of negative polarity:

$$\begin{array}{ll} A \vee B & \text{becomes } A \vee^+ B \\ A \wedge B & \text{becomes } A \wedge^- B \end{array} \quad \begin{array}{ll} A \rightarrow B & \text{becomes } A^\perp \vee^- B \\ (A \rightarrow B)^\perp & \text{becomes } A \wedge^+ B^\perp \end{array}$$

where $(A \rightarrow B)^\perp$ represents the dual of implication: *subtraction* (see e.g. [Cro04]).

¹Otherwise we would get back to Barbanera and Berardi's symmetric (and non-confluent) λ -calculus, with unclear denotational semantics.

This is what we will do in Section 3.2.

3.1.3 Focussing

Now, the polarised System **L** presented above, which has been studied at length by Munch-Maccagnoni [MM09, CMM10, MM13], solves non-confluence, not by giving systematic priority to the left (CBV) or to the right (CBN), but by giving priority to the non-invertible side (depending on the connective).

So the system takes advantage of the invertibility properties of the asynchronous rules (right-introduction of negative connectives, left-introduction of positive connectives). Invertibility entails that, in terms of proof-search, you can *chain* the decomposition of every formula of the sequent that has an asynchronous introduction rule, before doing anything else, without loss of generality (i.e. without losing logical completeness).

Now in [AP89, And92], Andreoli proved a more surprising result: *focussing*, that says that, once you have chosen to decompose by a synchronous rule a particular formula in the sequent,² you can also chain without loss of generality (i.e. without losing logical completeness) the recursive decomposition of its subformulae by synchronous rules until you reveal a subformula of the opposite polarity (whose decomposition can then be done by asynchronous rules again).

This was in the context of linear logic, whence polarities have come, but it is now understood in other polarised logics (classical or intuitionistic). This result can be expressed as the completeness of a sequent calculus with a *focus* device, which syntactically highlights a formula in the sequent and forces the next proof-search step to decompose it with a synchronous rule, keeping the focus on its newly revealed subformulae. In terms of proof-search, focussing considerably reduces the search space, otherwise heavily redundant when Gentzen-style inference rules are used.

Focussed proofs are proofs that implement such a chaining of synchronous decompositions. The main idea is that focussed proofs are those whose proof-terms systematically use term values and continuation values, in other words, the normal forms for ζ -rules. Of course, such normal forms may feature cuts (ζ -rules introduce cuts), but one should notice the following properties:

REMARK 38

Just like $\longrightarrow_{\zeta_N}$ and $\longrightarrow_{\zeta_V}$ (from Definition 21), the relation $\longrightarrow_{\zeta_F}$ is terminating. ※

DEFINITION 41 (LK^F)

Let LK^F be the fragment of System **L** consisting of $\longrightarrow_{\zeta_F}$ -normal forms. ※

REMARK 39 Just like LK^N and LK^V are stable under $\longrightarrow_{CBN}$ and $\longrightarrow_{CBV}$, the fragment LK^F is stable under $\longrightarrow_F$. ※

REMARK 40 These normal form fragments relate to calculi of the literature:

1. LK^N is exactly the calculus called LKT [DJS95, DJS97];
2. LK^V is exactly the calculus called LKQ [DJS95, DJS97];
3. LK^F relates to Liang and Miller's LKF [LM09], and this will be the object of Section 3.2. ※

²Positive formula on the right-hand side of a sequent, or a negative formula on its left-hand side

We therefore have 3 versions of the focussing result in classical logic, an unfocussed proof c can be turned into a focussed proof c' (in the sense of LK^N , LK^V , or LK^F) by normalising it with respectively $\longrightarrow_{\zeta_N}$, $\longrightarrow_{\zeta_V}$, or $\longrightarrow_{\zeta_F}$, and normalising it by respectively $\longrightarrow_{\text{CBN}}$, $\longrightarrow_{\text{CBV}}$, or $\longrightarrow_F$ to eliminate cuts and finally obtain a cut-free focussed proof.

3.1.4 Weak η -conversion

Now, the notion of η -conversion that we used as an introduction to polarities in Section 3.1.2, is a *strong* notion of η -conversion:

Inspired by the way we can reduce an axiom on a non-atomic formula into a proof using axioms on smaller formulae, we considered the η -expansion of variables x and α :

$$\begin{aligned} x &\longrightarrow \lambda y. \mu \alpha. \langle x \mid y :: \alpha \rangle \\ x &\longrightarrow (\mu \alpha. \langle x \mid \text{inj}_1(\alpha) \rangle, \mu \gamma. \langle x \mid \text{inj}_2(\gamma) \rangle) \\ \alpha &\longrightarrow (\mu x. \langle \text{inj}_1(x) \mid \alpha \rangle, \mu z. \langle \text{inj}_2(z) \mid \alpha \rangle) \end{aligned}$$

which we sought to generalise to

$$\begin{aligned} t &\longrightarrow \lambda y. \mu \alpha. \langle t \mid y :: \alpha \rangle \\ t &\longrightarrow (\mu \alpha. \langle t \mid \text{inj}_1(\alpha) \rangle, \mu \gamma. \langle t \mid \text{inj}_2(\gamma) \rangle) \\ e &\longrightarrow (\mu x. \langle \text{inj}_1(x) \mid e \rangle, \mu z. \langle \text{inj}_2(z) \mid e \rangle) \end{aligned}$$

for *any* term t and any continuation e .

This led to a polarity-based reduction relation that contrasts with the CBN and CBV reduction relations from Chapter 1. But that does not mean that CBN and CBV are incompatible with the concept of η -conversion: they just require *weaker* notions of η -conversion than that discussed above.

The notion of η -conversion that is suitable for CBN are

$$\begin{aligned} t &\longrightarrow \lambda y. \mu \alpha. \langle t \mid y :: \alpha \rangle \\ t &\longrightarrow (\mu \alpha. \langle t \mid \text{inj}_1(\alpha) \rangle, \mu \gamma. \langle t \mid \text{inj}_2(\gamma) \rangle) \\ E &\longrightarrow (\mu x. \langle \text{inj}_1(x) \mid E \rangle, \mu z. \langle \text{inj}_2(z) \mid E \rangle) \end{aligned}$$

where α has not been substituted by any continuation e but *only* by a continuation value E .

The notion of η -conversion that is suitable for CBV are

$$\begin{aligned} V &\longrightarrow \lambda y. \mu \alpha. \langle V \mid y :: \alpha \rangle \\ V &\longrightarrow (\mu \alpha. \langle V \mid \text{inj}_1(\alpha) \rangle, \mu \gamma. \langle V \mid \text{inj}_2(\gamma) \rangle) \\ e &\longrightarrow (\mu x. \langle \text{inj}_1(x) \mid e \rangle, \mu z. \langle \text{inj}_2(z) \mid e \rangle) \end{aligned}$$

where x has not been substituted by any term t but *only* by a term value V .

Including these notions of CBN- η -conversion and CBV- η -conversion in the CBN and CBV notions of reduction, is actually necessary if these are to capture the semantics of classical proofs in control and co-control categories, respectively: just like in Theorem 4 we needed η -conversion to make the simply-typed λ -calculus sound and complete with respect to the semantics given by CCC, here we would need the above notions of CBN- η -conversion and CBV- η -conversion in order to turn the implications of Theorem 22 (soundness) into equivalences (soundness and completeness). This is actually what Selinger proved [Sel01] in the context of the $\lambda\mu$ -calculus.

3.1.5 Related works

The role of polarities and focussing in classical proof theory has been investigated by a substantial literature, inspired by Girard’s linear logic [Gir87]. Following this work and Andreoli’s on focussing [AP89, And92], Girard developed in [Gir91] a sequent calculus LC for classical logic with more structure than Gentzen’s LK [Gen35], based on an assignment of polarities to classical formulae. Danos, Joinet and Schellinx [DJS95, DJS97] studied semantically meaningful ways to make cut-elimination confluent in the classical sequent calculus, introducing

- the calculi LKT and LKQ mentioned above
- a version of the sequent calculus called LK^{tq} where a *colour* t or q on each formula indicates whether a cut on that formula should be pushed to the right or to the left,
- more restricted versions thereof,

all inspired by the various translations of classical logic into linear and intuitionistic logics. Out of that field, which includes the duality between CBN and CBV,³ *polarised classical logic* emerged, developed as such by Laurent et al. [Lau02, LQdF05]. It develops and enriches Girard’s work on LC, in particular by explaining the proof theory of classical formulae as given by LC as a combination of

- an encoding from classical formulae to polarised classical formulae
- a proof theory for polarised classical logic.

Closer to Andreoli’s original line of research, which was motivated by logic programming, Liang and Miller then formalised LKF as a more strongly focussed calculus than that called LK^F above; we will study it in the next section.

A useful introduction to that literature can be found in Chapter 2 of Farooque’s thesis [Far13].

More recently, Munch-Maccagnoni approached the concept of focussing via orthogonality models [MM09]. He built for the polarised version of System L the same kind of orthogonality model as the one we presented in Section 2.2.1 for the LK^N -case, with an interpretation $\llbracket A \rrbracket^+$ of a formula A built as the orthogonal or bi-orthogonal of a more basic set of (counter-)proof-terms. He essentially shows that the interpretation $\llbracket A \rrbracket^+$ of a formula is generated from its values, in the sense that $\llbracket A \rrbracket^+ = (\llbracket A \rrbracket^+ \cap \mathcal{V})^{\perp\perp}$ where $\mathcal{V}$ denotes the set of term values (and symmetrically $\llbracket A \rrbracket^- = (\llbracket A \rrbracket^- \cap \mathcal{E})^{\perp\perp}$ where $\mathcal{E}$ denotes the set of continuation values).

This sheds an interesting light on our definition of

$$[A \rightarrow B]_\sigma^- := \llbracket A \rrbracket_\sigma^+ :: \llbracket B \rrbracket_\sigma^-$$

in our orthogonality models of LK^N (Definitions 31 and 36): While in LK^N the above construct only considers those inhabitants of $\llbracket B \rrbracket_\sigma^-$ that are continuation values anyway, there is in the general case of System L a question of whether we want continuations of the form $t :: \mu x.c$,⁴ which are not focussed (in the sense of LK^N or LK^F), in the interpretation $[A \rightarrow B]_\sigma^-$. The result that $\llbracket A \rightarrow B \rrbracket^- = (\llbracket A \rightarrow B \rrbracket^- \cap \mathcal{E})^{\perp\perp}$ means that an “unfocussed” counter-proof such as $t :: \mu x.c$ would be accepted after the bi-orthogonal completion (i.e. in $\llbracket A \rightarrow B \rrbracket_\sigma^- = (\llbracket A \rightarrow B \rrbracket^-)^{\perp\perp}$).

So far we have taken advantage of focussing, i.e. the chaining of synchronous rules, to identify complete fragments LK^N , LK^V , and LK^F of classical sequent calculus proofs.

³As revealed by Curien and Herbelin’s System L [CH00] and Selinger’s control categories [Sel01].

⁴(with $t \in \llbracket A \rrbracket_\sigma^+$ and $\mu x.c \in \llbracket B \rrbracket_\sigma^-$)

Although we have discussed invertibility of asynchronous rules, in order to introduce the notion of polarity, we have not forced our proofs to apply asynchronous rules eagerly, before applying other rules (in LK^F , $\mu\alpha^-.c$ is still an accepted proof of a negative formula such as $A \rightarrow B$).

This is the main difference with the focussed proof systems in the style of e.g. Liang and Miller [LM09], where e.g. all proofs of $A \rightarrow B$ finish with the right-introduction of $\rightarrow$. In terms of proof-terms, it means that all proof-terms are in η -long normal forms (we can transform every proof-term into a proof-term in that form by a series of η -expansions, but it is always tricky to control the termination of η -expansion without having the types explicitly in the terms).

Coming back to the purely logical level, focussed proofs in the tradition of Miller et al. can be described in terms of “big-step focussing”: going up a branch of the proof is an alternation of synchronous and asynchronous phases, which we may consider to be atomic. The next section shows a computational interpretation of that strongly focussed formalism.

3.2 Computational interpretation of a focussed calculus

The starting point of this section is Liang and Miller’s LKF [LM09], a variant of the system LK^F described in the previous section that forces the asynchronous decomposition of formulae. It is described in purely logical terms and we will see how, by formalising the concept of *big-step focussing*, a Curry-Howard interpretation can be given to LKF, following Zeilberger’s work [Zei08a, Zei08b].

We start with the formulae of *polarised classical logic*.

DEFINITION 42 (Polarised formulae)

The syntax of formulae is given by the following grammar

$$\begin{array}{lll} \text{Positive formulae } P & ::= & a \mid A_1 \wedge^+ A_2 \mid A_1 \vee^+ A_2 \\ \text{Negative formulae } N & ::= & a^\perp \mid A_1 \wedge^- A_2 \mid A_1 \vee^- A_2 \\ \text{Formulae } A & ::= & P \mid N \end{array}$$

where a ranges over a fixed set of elements called *positive atoms*, and $a^\perp$ ranges over a bijective copy of that set ($a \mapsto a^\perp$ is the bijection), whose elements are called *negative atoms*.

We extend the bijection between positive and negative atoms into an involutive bijection, called *negation*, between positive and negative formulae:

$(a)^\perp$	$:= a^\perp$	$(a^\perp)^\perp$	$:= a$
$(A_1 \wedge^+ A_2)^\perp$	$:= A_1^\perp \vee^- A_2^\perp$	$(A_1 \wedge^- A_2)^\perp$	$:= A_1^\perp \vee^+ A_2^\perp$
$(A_1 \vee^+ A_2)^\perp$	$:= A_1^\perp \wedge^- A_2^\perp$	$(A_1 \vee^- A_2)^\perp$	$:= A_1^\perp \wedge^+ A_2^\perp$

※

Following the suggestion made in the previous section, we fold LKF into a 1-sided sequent calculus (hence our interest for the involutive negation), as it is traditional in the field arising from linear logic [Gir87].

EXAMPLE 8 For instance, Peirce’s law, which in the previous chapters and sections we wrote as $((a \rightarrow b) \rightarrow a) \rightarrow a$, is now $((a^\perp \vee^- b) \wedge^+ a^\perp) \vee^- a$. ※

DEFINITION 43 (Liang-Miller's LKF)

The rules of LKF are given in Fig. 13 for two kinds of sequents:

- $\vdash \Theta \Downarrow A$ focussed sequent
- $\vdash \Theta \Uparrow \Gamma$ unfocussed sequent

where A is an arbitrary formula, Θ is a multiset of either negative atoms or positive formulae and Γ is a multiset of arbitrary formulae.

Derivability in LKF of the sequents $\vdash \Theta \Downarrow A$ and $\vdash \Theta \Uparrow \Gamma$ is respectively denoted $\vdash_{\text{LKF}} \Theta \Downarrow A$ and $\vdash_{\text{LKF}} \Theta \Uparrow \Gamma$. *

Synchronous phase	$\frac{\vdash \Theta \Downarrow A_1 \quad \vdash \Theta \Downarrow A_2}{\vdash \Theta \Downarrow A_1 \wedge^+ A_2}$	$\frac{\vdash \Theta \Downarrow A_i}{\vdash \Theta \Downarrow A_1 \vee^+ A_2}$
End of synchronous phase	$\frac{\vdash \Theta \Uparrow N}{\vdash \Theta \Downarrow N}$	$\frac{}{\vdash \Theta \Downarrow a} a^\perp \in \Theta$
Asynchronous phase	$\frac{\vdash \Theta \Uparrow A_1, \Gamma \quad \vdash \Theta \Uparrow A_2, \Gamma}{\vdash \Theta \Uparrow A_1 \wedge^- A_2, \Gamma}$	$\frac{\vdash \Theta \Uparrow A_1, A_2, \Gamma}{\vdash \Theta \Uparrow A_1 \vee^- A_2, \Gamma}$
End of asynchronous phase	$\frac{\vdash \Theta, P \Uparrow \Gamma}{\vdash \Theta \Uparrow P, \Gamma}$	$\frac{\vdash \Theta, p^\perp \Uparrow \Gamma}{\vdash \Theta \Uparrow p^\perp, \Gamma} \quad \frac{\vdash \Theta, P \Downarrow P}{\vdash \Theta, P \Uparrow}$

Figure 13: LKF

Liang and Miller showed in [LM09] that

- various cut-rules are admissible;
- the polarities of atoms and connectives do not change the provability of an unfocussed sequent, but they change the shape of its proofs;
- with the admissibility of cut-rules, the system is complete for classical logic, no matter which polarities are placed on connectives and literals.

To prove cut-admissibility, they do not explicitly formalise a cut-elimination procedure, but it could probably be inferred from the proof.

Note that soundness of the system with respect to classical logic is trivially checked, rule by rule, forgetting about polarities and the structure of sequents.

Polarities in classical logic raise interesting questions: $A \wedge^+ B$ and $A \wedge^- B$ are equiprovable, and so are $A \vee^+ B$ and $A \vee^- B$. But, while the difference between the (direct) proofs of $A \vee^+ B$ and (direct) proofs of $A \vee^- B$ is clear,⁵ one may wonder what the real difference is between (direct) proofs of $A \wedge^+ B$ and (direct) proofs of $A \wedge^- B$, given that the two rules look very much alike.

The difference lies not in their structure, but in the way they will behave in cut-elimination:

- from a proof of $A \wedge^- B$ (facing a proof of $A^\perp \vee^+ B^\perp$), only one sub-proof is used while the other is thrown away,
- from a proof of $A \wedge^+ B$ (facing a proof of $A^\perp \vee^- B^\perp$), both sub-proofs are used.

⁵Direct proofs of $A \vee^+ B$ choose one side and throw away the other, while direct proofs of $A \vee^- B$ keep the two sides.

Metaphorically, proving either conjunction is like picking 1 boy name and 1 girl name, when your couple is pregnant: proving a negative conjunction is picking the two names when you are expecting one baby (not knowing whether it is a boy or a girl), while proving the positive conjunction is picking the two names when expecting twins (a boy and a girl). On the paper, you have the same job to do, but you will probably approach the problem very differently.

3.2.1 Informal relation to System L

It may not be obvious, but system LKF roughly expresses, without proof-terms, some derivations of System L in a 1-sided format.⁶

Indeed, think of

- a focussed sequent $\vdash \Theta \Downarrow A$ as a typing judgement for a term value V : $\vdash V : A ; \Theta$.
- an unfocussed sequent $\vdash \Theta \Uparrow \Gamma$ as a typing judgement for a command c : $c : (\vdash ; \Theta, \Gamma)$.

with Γ being the part of the typing context for negative continuation variables α^- with non-atomic types (which can be asynchronously decomposed), and Θ the rest of it (typing negative continuation variables α^- with atomic types, and typing positive continuation variables α^+).

To derive a focussed sequent $\vdash \Theta \Downarrow A$:

- the two rules of the group ‘Synchronous phase’ correspond to the typing rules for $V :: V'$ and for $\text{inj}_i(V')$;
- the first rule of the group ‘End of synchronous phase’ does not correspond to a rule of System L but simply the realisation that the value V is a negative term t^- ;
- the second rule of that group is when V is a variable.

To derive an unfocussed sequent $\vdash \Theta \Uparrow \Gamma$:

- the two rules of the group ‘Asynchronous phase’ correspond to the typing of (t_1, t_2) and $\lambda \alpha. t$;
- the first two rules of the group ‘End of asynchronous phase’ are not reflected in System L (they just move formulae that cannot be asynchronously decomposed from Γ to Θ)
- the third rule of that phase corresponds to the typing of $\langle V \mid \alpha^+ \rangle$.

Roughly speaking, we should think of LKF as typing those proof-terms of (a 1-sided version of) System L that are η -long $\rightarrow_F$ -normal forms. These are described by the following grammar:

$$\begin{aligned} V, V' &::= \alpha^+ \mid V :: V' \mid \text{inj}_i(V) \mid t^- \mid \mu \alpha^-. c \\ t^-, \dots &::= (t_1, t_2) \mid \lambda \alpha. t \\ t, \dots &::= \mu \alpha^+. c \mid t^- \mid \mu \alpha^-. c \\ c, \dots &::= \langle V \mid \alpha^+ \rangle \mid \langle t^- \mid \alpha^- \rangle \end{aligned}$$

where the type of every $\mu \alpha^-. c$ is atomic.

It is only “roughly speaking”, because in Liang-Miller’s LKF:

1. when a formula is asynchronously decomposed, no copy of the formula is kept in the sequent (which means in the above grammar that in every command $\langle t^- \mid \alpha^- \rangle$ we impose $\alpha^- \notin \text{FV}(t^-)$),

⁶A 1-sided version of System L would merge terms and continuations into terms, so that $V :: V'$ is a term value, and merge term variables and continuation variables into continuation variables.

2. when the focus is placed on a formula, all the formulae that could be asynchronously decomposed have already been asynchronously decomposed (which means in the above grammar that in every command $\langle V \mid \alpha^+ \rangle$, V has no free variable of the form α^+).
3. finally, the order in which formulae are decomposed in the asynchronous phase, is less deterministic than that imposed by the above grammar.

This is because, in Liang and Miller's view of focussing, and more generally in the tradition of linear logic, "what happens in the asynchronous phase stays in the asynchronous phase", in the sense that the details of the asynchronous phase (e.g. the order in which formulae are decomposed) are meaningless and should not impact the semantics of the proof.

This is difficult to reflect at the level of System L's proof-terms.

Therefore, we will now develop a Curry-Howard interpretation for that particular view of focussing, along the lines of Zeilberger's work [Zei08a, Zei08b, Zei10]: in order to forget about the inner details of the asynchronous phase, we formalise the idea of compacting each phase (asynchronous and even synchronous) into one atomic inference. This is called big-step focussing.

3.2.2 Identifying phases as atomic steps

We start by showing an example of how positive connectives are decomposed.

EXAMPLE 9 Trying to prove $\vdash \Theta \Downarrow N_1 \wedge^+(a \vee^+ N_2)$ we can build:

either

$$\begin{array}{c}
 \text{End of synch phase} \qquad \text{End of synch phase} \\
 \hline
 \vdash \Theta \Uparrow N_1 \qquad \frac{a^\perp \in \Theta}{\vdash \Theta \Downarrow a} \\
 \hline
 \vdash \Theta \Downarrow N_1 \qquad \vdash \Theta \Downarrow a \vee^+ N_2 \\
 \hline
 \vdash \Theta \Downarrow N_1 \wedge^+(a \vee^+ N_2)
 \end{array}$$

or

$$\begin{array}{c}
 \text{End of synch phase} \qquad \text{End of synch phase} \\
 \hline
 \vdash \Theta \Uparrow N_1 \qquad \vdash \Theta \Uparrow N_2 \\
 \hline
 \vdash \Theta \Downarrow N_1 \qquad \vdash \Theta \Downarrow N_2 \\
 \hline
 \vdash \Theta \Downarrow a \vee^+ N_2 \\
 \hline
 \vdash \Theta \Downarrow N_1 \wedge^+(a \vee^+ N_2)
 \end{array}$$

The whole synchronous phase can be expressed in just one step:

$$\frac{\vdash \Theta \Uparrow N_1 \quad a^\perp \in \Theta}{\vdash \Theta \Downarrow N_1 \wedge^+(a \vee^+ N_2)} \quad \text{or} \quad \frac{\vdash \Theta \Uparrow N_1 \quad \vdash \Theta \Uparrow N_2}{\vdash \Theta \Downarrow N_1 \wedge^+(a \vee^+ N_2)}$$

In other words:

$$\frac{\forall N \in \Gamma, \quad \vdash \Theta \Uparrow N \quad \forall a \in \Gamma, \quad a^\perp \in \Theta}{\vdash \Theta \Downarrow N_1 \wedge^+(a \vee^+ N_2)}$$

with $\Gamma = N_1, a$ or $\Gamma = N_1, N_2$

In either case, we say that Γ "is a positive decomposition of" $N_1 \wedge^+(a \vee^+ N_2)$, which we denote: $N_1, a \Vdash^+ N_1 \wedge^+(a \vee^+ N_2)$ and $N_1, N_2 \Vdash^+ N_1 \wedge^+(a \vee^+ N_2)$. *

Now we generalise this example into a formal definition:

DEFINITION 44 (Decomposition of positive connectives)

The *positive decomposition relation* is the binary relation, defined by the rules of Fig. 14, where $\Gamma, \Gamma_1, \Gamma_2$ are sets of positive atoms or negative formulae.

The *one-step synchronous phase* is the rule:

$$\frac{\Gamma \Vdash^+ A \quad \forall N \in \Gamma, \vdash \Theta \uparrow N \quad \forall a \in \Gamma, \quad a^\perp \in \Theta}{\vdash \Theta \Downarrow A} \text{synch}$$

※

$$\frac{\frac{\overline{N \Vdash^+ N} \quad \overline{a \Vdash^+ a}}{\Gamma_1 \Vdash^+ A_1 \quad \Gamma_2 \Vdash^+ A_2} \quad \frac{\Gamma \Vdash^+ A_i}{\Gamma \Vdash^+ A_1 \vee^+ A_2}}{\Gamma_1, \Gamma_2 \Vdash^+ A_1 \wedge^+ A_2} \quad \Gamma \Vdash^+ A_1 \vee^+ A_2$$

Figure 14: Positive decomposition relation

Notice the syntax we use for the **synch** rule: the symbols $\forall$ and $\in$ are **meta-level** symbols: the number of premisses is the cardinal of Γ (plus one if you count $\Gamma \Vdash^+ A$).

We now show an example of how negative connectives are decomposed.

EXAMPLE 10

$$\frac{\frac{\text{End of asynch phase}}{\vdash \Theta, P_1, a^\perp \uparrow} \quad \frac{\text{End of asynch phase}}{\vdash \Theta, P_1, P_2 \uparrow}}{\vdash \Theta \uparrow P_1, a^\perp} \quad \frac{\vdash \Theta \uparrow P_1, (a^\perp \wedge^- P_2)}{\vdash \Theta \uparrow P_1 \vee^- (a^\perp \wedge^- P_2)}$$

The whole asynchronous phase can be expressed in just one step:

$$\frac{\vdash \Theta, P_1, a^\perp \uparrow \quad \vdash \Theta, P_1, P_2 \uparrow}{\vdash \Theta \uparrow P_1 \vee^- (a^\perp \wedge^- P_2)} \text{asynch}$$

In other words

$$\frac{\forall \Delta, \quad \vdash \Theta, \Delta \uparrow}{\vdash \Theta \uparrow P_1 \vee^- (a^\perp \wedge^- P_2)}$$

where Δ ranges over $\{ \{P_1, a^\perp\}, \{P_1, P_2\} \}$

In either case, we say that Δ “is a negative decomposition of” $P_1 \vee^- (a^\perp \wedge^- P_2)$, which we denote $P_1, a^\perp \Vdash^- P_1 \vee^- (a^\perp \wedge^- P_2)$ and $P_1, P_2 \Vdash^- P_1 \vee^- (a^\perp \wedge^- P_2)$. ※

Now we generalise this example into a formal definition:

DEFINITION 45 (Decomposition of negative connectives)

The *negative decomposition relation* is the binary relation, defined by the rules of Fig. 15, where where $\Delta, \Delta_1, \Delta_2$ are sets of negative atoms or positive formulae.

The *one-step asynchronous phase* is the rule:

$$\frac{\forall \Delta, (\Delta \Vdash^- A) \Rightarrow (\vdash \Theta, \Delta \Uparrow)}{\vdash \Theta \Uparrow A}$$

※

$$\frac{\overline{P \Vdash^- P} \quad \overline{a^\perp \Vdash^- a^\perp}}{\frac{\Delta \Vdash^- A_i}{\Delta \Vdash^- A_1 \wedge^- A_2} \quad \frac{\Delta_1 \Vdash^- A_1 \quad \Delta_2 \Vdash^- A_2}{\Delta_1, \Delta_2 \Vdash^- A_1 \vee^- A_2}}$$

Figure 15: Negative decomposition relation

Again, notice that the syntax we use for the **asynch** rule uses the **meta-level** symbols $\forall$ and $\Rightarrow$: the number of premisses is the number of Δ satisfying $\Delta \Vdash^- A$ for the given A .

We now put everything together in the style of Zeilberger [Zei08a, Zei08b, Zei10].

DEFINITION 46 (Big-step LKF, v1)

The big-step LKF system is given in Fig. 16, where Θ, Δ are sets of negative atoms or positive formulae and Γ is a set of positive atoms or negative formulae.

※

$$\frac{\Gamma \Vdash^+ A \quad \forall N \in \Gamma, \vdash \Theta \Downarrow N \quad \forall a \in \Gamma, a^\perp \in \Theta}{\vdash \Theta \Downarrow A} \text{synch}$$

$$\frac{\vdash \Theta, P \Downarrow P}{\vdash \Theta, P \Uparrow} \text{focus} \quad \frac{\forall \Delta, (\Delta \Vdash^- A) \Rightarrow (\vdash \Theta, \Delta \Uparrow)}{\vdash \Theta \Uparrow A} \text{asynch}$$

Figure 16: Big-step LKF, v1

REMARK 41

Sequents of the form $\vdash \Theta \Downarrow N$ and sequents of the form $\vdash \Theta \Uparrow P$ are never present in the premisses of the rules. Such sequents can only appear as the very conclusion of a whole proof-tree.

Hence, we can equivalently present the big-step LKF system as the system of Fig. 17, and declare $\vdash \Theta \Uparrow P$ as syntactic sugar for $\vdash \Theta, P \Uparrow$, and $\vdash \Theta \Downarrow N$ as syntactic sugar for $\vdash \Theta \Downarrow N$.

※

$$\frac{\Gamma \Vdash^+ P \quad \forall N \in \Gamma, \vdash \Theta \Uparrow N \quad \forall a \in \Gamma, a^\perp \in \Theta}{\vdash \Theta \Downarrow P}$$

$$\frac{\vdash \Theta, P \Downarrow P}{\vdash \Theta, P \Uparrow} \quad \frac{\forall \Delta, (\Delta \Vdash^- N) \Rightarrow (\vdash \Theta, \Delta \Uparrow)}{\vdash \Theta \Uparrow N}$$

Figure 17: Big-step LKF, v2

Now we should notice a complete symmetry, and therefore some redundancy, in the two decomposition relations that we have defined:

REMARK 42 $\Gamma \Vdash^+ P$ if and only if $\Gamma^\perp \Vdash^- P^\perp$. ※

Hence, we can define in Fig. 18 a simplified version of big-step LKF, where this redundancy is eliminated.

$$\begin{array}{c}
 \hline
 \Gamma \Vdash P \quad \forall N \in \Gamma, \vdash \Theta \uparrow N \quad \forall a \in \Gamma, a^\perp \in \Theta \\
 \hline
 \vdash \Theta \Downarrow P \\
 \\
 \frac{\vdash \Theta, P \Downarrow P}{\vdash \Theta, P \uparrow} \quad \frac{\forall \Gamma, (\Gamma \Vdash N^\perp) \Rightarrow (\vdash \Theta, \Gamma^\perp \uparrow)}{\vdash \Theta \uparrow N}
 \end{array}$$

where $\Gamma \Vdash P$ is $\Gamma \Vdash^+ P$.

Figure 18: Big-step LKF, v3

Now notice that the rules for negative connectives are never used in the system! Due to the duality in the syntax, given by the involutive negation, we should be able to remove negative connectives altogether. We just need to introduce a marker in the syntax of a formula, to denote every change of polarity.

Let us write $\neg$ for this marker.

DEFINITION 47 (Syntax with positive connectives only)

Formulae are now defined by the following syntax:

$$\begin{aligned}
 P &::= a \mid A_1 \wedge^+ A_2 \mid A_1 \vee^+ A_2 \\
 A &::= P \mid \neg P
 \end{aligned}$$

with the following involutive negation:

$$\begin{aligned}
 P^\perp &::= \neg P \\
 (\neg P)^\perp &::= P
 \end{aligned}$$

※

REMARK 43 The previous grammar can be encoded into that one:

$$\begin{aligned}
 \bar{a} &::= a \\
 \overline{A \wedge^+ B} &::= \bar{A} \wedge^+ \bar{B} & \bar{N} &::= \neg(\bar{N}^\perp) \\
 \overline{A \vee^+ B} &::= \bar{A} \vee^+ \bar{B}
 \end{aligned}$$

※

In Fig. 19 we reformulate big-step LKF with this syntax for formulae.

$$\begin{array}{c}
 \hline
 \Gamma \Vdash P \quad \forall \neg P' \in \Gamma, \vdash \Theta \uparrow \neg P' \quad \forall a \in \Gamma, \neg a \in \Theta \\
 \hline
 \vdash \Theta \Downarrow P \\
 \\
 \frac{\vdash \Theta, P \Downarrow P}{\vdash \Theta, P \uparrow} \quad \frac{\forall \Gamma, (\Gamma \Vdash P) \Rightarrow (\vdash \Theta, \Gamma^\perp \uparrow)}{\vdash \Theta \uparrow \neg P}
 \end{array}$$

where $\Gamma \Vdash P$ is $\Gamma \Vdash^+ P$.

Figure 19: Big-step LKF, v4

Finally, we notice that it is more natural to write Γ on the left-hand side of a sequent:

DEFINITION 48 (Big-step LKF, v5)

The big-step LKF system v5 is given in Fig. 20, where Γ is a set of atoms a or formulae of the form $\neg P$ and Θ is a set of negated atoms $\neg a$ or formulae of the form P . ※

$$\begin{array}{c}
 \hline
 \frac{\Gamma \Vdash P \quad \forall \neg P' \in \Gamma, \Gamma_0 \vdash \Uparrow \neg P' \quad \forall a \in \Gamma, a \in \Gamma_0}{\Gamma_0 \vdash \Downarrow P} \\
 \\
 \frac{\Gamma_0, \neg P \vdash \Downarrow P}{\Gamma_0, \neg P \vdash \Uparrow} \qquad \frac{\forall \Gamma, (\Gamma \Vdash P) \Rightarrow (\Gamma_0, \Gamma \vdash \Uparrow)}{\Gamma_0 \vdash \Uparrow \neg P} \\
 \hline
 \end{array}$$

where $\Gamma \Vdash P$ is $\Gamma \Vdash^+ P$.

Figure 20: Big-step LKF, v5

The lesson to be remembered from this formulation of big-step LKF, is that (the big-step version of) asynchronous rules happens to **coincide** with a rule inferred from (the big-step version of) synchronous rules. This will make cut-elimination work, and it formalises (at least in classical logic) the concept known in philosophical logic as *harmony* [Ten78, Rea00, Rea10] (expressed originally between the introduction rules and elimination rules of Natural Deduction, or between left-introduction rules and right-introduction rules of Sequent Calculus).

Now we can consider that both synchronous and asynchronous rules are defined primitively, and notice the somewhat “miraculous” coincidence, or we can adopt the view that only synchronous rules are defined primitively; asynchronous phases then work by duality from the way synchronous phases work.

In other words,

- positive connectives are “defined” by their introduction rules;
- negative connectives are “defined” by duality, from the introduction rules of their positive duals.

We may not even need to bother representing their rules.

In this view, and via the Curry-Howard correspondence, we should define how to inhabit a type $A \rightarrow B$ with (proof-)terms, from the way we inhabit A with terms and B with continuations (with $\rightarrow$ being a negative connective). Writing $\lambda x.M$ with a variable $x:A$ and a body $M:B$, would then only be a mere representation for (or a mere even implementation of) an inhabitant of $A \rightarrow B$ that pre-exists the syntactical notation.

This is of course expressed *semantically* in orthogonality models (say in Definitions 28, 31 and 36) by the fact that we first define an interpretation for positive formulae (mentioning the syntax of their basic inhabitants, such as the construct $t::e$), and in a second step we define the interpretation of negative formulae simply as the orthogonal of the interpretation of their dual formula. The definition for negative formulae does not even mention the syntax of their inhabitants (such as $\lambda x.M$), but if we have a syntax for them, they “happen to live” (somewhat miraculously) in the interpretation.

We now formalise a way to express this *syntactically*, as a proof-term calculus for big-step LKF.

3.2.3 Functional interpretation as pattern-matching

Earlier we wrote that “the proofs of negatives *must interact well with* the proofs of the positive dual”. The intuition we formalise is that

- the proofs of a positive connective (i.e. of some $\Gamma_0 \vdash \Downarrow P$) are some data that can be *pattern-matched*;
- the proofs of a negative connective (i.e. of some $\Gamma_0 \vdash \Uparrow \neg P$) are *functions* that consume data by *pattern-matching*.

The fact that the proofs of a negative are determined by duality from the proofs of the positive dual, is reflected by the fact that the shape of a pattern-matching function is indeed completely determined by the data-type of its argument.

So the Curry-Howard interpretation of big-step LKF is an abstract system of pattern-matching.

The “proof-terms” for the decomposition of (positive) connectives are *patterns*. For instance for the connectives $\wedge^+, \vee^+$:

DEFINITION 49 (Patterns for $\wedge^+, \vee^+$) Patterns are defined by the following syntax:

$$p ::= x^+ \mid x^- \mid (p_1, p_2) \mid \text{inj}_i(p)$$

Their typing rules are presented in Fig. 21. ※

$$\frac{}{x^- : \neg P \Vdash x^- : \neg P} \quad \frac{}{x^+ : a \Vdash x^+ : a}$$

$$\frac{\Gamma_1 \Vdash p_1 : A_1 \quad \Gamma_2 \Vdash p_2 : A_2}{\Gamma_1, \Gamma_2 \Vdash (p_1, p_2) : A_1 \wedge^+ A_2} \quad \frac{\Gamma \Vdash p : A_i}{\Gamma \Vdash \text{inj}_i(p) : A_1 \vee^+ A_2}$$

Figure 21: Decomposition with patterns

We now give the proof-terms for big-step LKF:

DEFINITION 50 (Pattern-matching calculus)

Let Pat be a set of elements called *patterns*, and denoted $p, p', \dots$

The syntax of proof-terms is given by the following grammar:

$$\begin{array}{lll} \textbf{Positive terms} & t^+ & ::= p.\sigma \\ \textbf{Negative terms} & t^- & ::= f \\ \textbf{Commands} & c & ::= \langle x^- \mid t^+ \rangle \mid \langle f \mid t^+ \rangle \end{array}$$

where

- σ is a substitution from negative variables such as x^- to negative terms, and from positive variables such as x^+ to positive terms;
- f is a function from patterns to commands.

Let $\mathbb{A}$ and $\mathbb{M}$ be two sets of elements called *atoms* and *molecules* and denoted p and P , respectively.

Let *typing contexts* be functions mapping negative variables to molecules (written $x^- : \neg P$) and positive variables to atoms (written $x^+ : a$).

Let $\Gamma \Vdash p : P$ be a typing relation where p is a pattern, P is a molecule, and Γ is a typing context.

The typing rules for proof-terms are presented in Fig. 22.

There is just one cut-elimination rule:

$$(\text{pat-match}) \quad \langle f \mid p.\sigma \rangle \longrightarrow (f(p))\sigma$$

where $c\sigma$ denotes the application of substitution σ to the command c . ※

$$\frac{\Gamma \Vdash p:P \quad \forall(x^-:\neg P) \in \Gamma, \quad \Gamma_0 \vdash \uparrow \sigma(x^-):\neg P \quad \forall(x^+:a) \in \Gamma, \quad (\sigma(x^+):a) \in \Gamma_0}{\Gamma_0 \vdash \Downarrow p.\sigma:P}$$

$$\frac{\forall \Gamma, (\Gamma \Vdash p:P) \Rightarrow f(p):(\Gamma_0, \Gamma \vdash \uparrow)}{\Gamma_0 \vdash \uparrow f:\neg P}$$

$$\frac{\Gamma_0, x^-:\neg P \vdash \Downarrow p.\sigma:P}{\langle x^- \mid p.\sigma \rangle:(\Gamma_0, x^-:\neg P \vdash \uparrow)} \quad \frac{\Gamma_0 \vdash \uparrow f:\neg P \quad \Gamma_0 \vdash \Downarrow t^+:P}{\langle f \mid t^+ \rangle:(\Gamma_0 \vdash \uparrow)}$$

where $\Gamma_0, \Gamma, \dots$ are typing contexts.

Figure 22: Typing for the pattern-matching calculus

The one cut-elimination rule is the very standard mechanism of pattern-matching, with the command $\langle f \mid t^+ \rangle$ representing what we could informally write as:

“match t^+ with $\underbrace{\dots \mapsto \dots}_f$ ”

REMARK 44

1. Notice how negative terms are not really terms, but functions of the meta-level (or meta-level functions that are reified in the term syntax); this is a higher-order definition, and we do not give any concret syntax for such functions.

Strictly speaking, our definition depends on the notion of function space that we take for the definition of negative terms (We could for instance restrict it to computable functions, but so far we do not specify such things).

Also, with such a definition, it may not be clear exactly what the contextual closure of the rule (pat-match) is. By $\longrightarrow_{(\text{pat-match})}$ we therefore denote the reduction relation where (pat-match) is applied at the top-level of a given command (no contextual closure).

2. Also notice how we emphasised that the definition of the proof-term calculus is *independent* from the syntax of patterns and the typing system for them.

Definition 49 and Fig. 21 give one example (where atoms are positive atomic formulae and molecules are positive formulae).

But the construction of the Curry-Howard interpretation for big-step focussing is modular in those notions.

This is a gain of genericity / abstraction that we will further develop in the next Chapters. ※

THEOREM 45 (Subject Reduction)

If $c:(\Gamma \vdash \uparrow)$ and $c \longrightarrow_{(\text{pat-match})} c'$ then $c':(\Gamma \vdash \uparrow)$. ※

Reviewing the various properties that are desirable for an instance of the Curry-Howard correspondence, we find that Progress (in this case, cut-elimination), depends on how we may reduce functions (so far, $\longrightarrow_{(\text{pat-match})}$ only applies at the root); Confluence does not make sense here because, until we define how to reduce functions, there is at most one redex to reduce; and for Normalisation, we need to explore models (e.g. orthogonality models) of such a calculus.

Conclusion

The study of orthogonality models for the pattern-matching calculus should be particularly interesting:

In [MM09], Munch-Maccagnoni already explored the construction of orthogonality models for polarised System L with an emphasis on focussing properties. Big-step LKF and its underlying pattern-matching calculus seems to be an even more appropriate framework to look at the connection between focussing and orthogonality models, since this framework reflects at the syntactical level what orthogonality models describe at the semantical level, namely the fact that we first declare what the “inhabitants of positive formulae” are, and then we define the “inhabitants of negative formulae” by duality as those inhabitants that “interact well with” the inhabitants of the dual (positive) formula. In case of an orthogonality model, to “interact well with” means to “be orthogonal to”; in the case of pattern-matching, it means to “be able to consume”. To be more precise:

- Inhabitants of positive types have *structure*: in an orthogonality model we need an algebraic structure to interpret positive constructs such as $_::_$ or $\text{inj}_(_)$; in big-step LKF, these inhabitants come as the combination of a pattern (e.g. $_::_$ or $\text{inj}_(_)$) and a substitution that fills its holes.
- Inhabitants of negative formulae may lack any structure, but they come with a *behaviour*: in an orthogonality model, they can range over any abstract set for which the orthogonality relation with positive inhabitants is defined; in big-step LKF, they range over any abstract set of functions (we do not specify which) that can consume patterns.

So, in order to formalise the connections that are informally described above, the second part of this dissertation explores orthogonality models for big-step focussing systems. We shall strip anything that is not essential off the constructions we make, systematically seeking the greatest generality, and aiming at the cores of orthogonality models and focussing systems. Doing so reveals the essential difference between realisability and typing:

- in realisability, checking whether a given negative inhabitant “interacts well with” an arbitrary inhabitant of a positive formula, requires the computation of an interaction that explores the positive inhabitant’s structure to an **arbitrary depth** (as nothing restricts the criterion given by orthogonality);
- in typing, checking whether a given negative inhabitant “interacts well with” an arbitrary inhabitant of a positive formula, only requires the computation of an interaction that explores the positive inhabitant’s structure to a **bounded depth** (as the negative inhabitant is a function that performs a case analysis on the positive inhabitant’s top-level pattern and the interaction has to uniformly treat the rest of the inhabitant’s structure).

In case each positive formula comes with a finite number of patterns for it, the above distinction is what makes typing decidable and realisability undecidable (in general).

Part II

Abstract focussing

Introduction

The second part of this dissertation presents unpublished material, on the theme of *abstract focussing*.

In the previous chapters we have seen the use of polarities and focussing in the proof theory of classical logic, where a focussed proof is a tree that alternates *synchronous phases* with *asynchronous phases*.

A level of abstraction is reached by *big-step focussing*, which compacts each phase into one inference step and thus allows the inner details of phases to be “forgotten”. As revealed by Zeilberger’s formulation of big-step focussing [Zei08a, Zei08b], the computational interpretation of this is pattern-matching.

In parallel to this, Munch-Maccagnoni [MM09] formalised the connection between focussing and the orthogonality techniques, which were presented in Chapter 2 for strong normalisation proofs and witness extraction.

The origin of the material presented in this second part of this dissertation is the idea that this deep connection could be revealed at a more abstract level if a Zeilberger-style system was used: For instance, the fact that, in such a system, the inhabitants of negative formulae live in an abstract function space and are not made of any syntax reflects the fact that, in orthogonality models, inhabitants of negative formulae can range over an abstract set and have no algebraic structure. The second part of this dissertation therefore started as a formalisation, in the proof-assistant Coq [Coq], of orthogonality models for a Zeilberger-style system, culminating with the Adequacy Lemma that connects the big-step focussing proof system with the orthogonality approach.

Doing this formalises the connection at a level of abstraction that forgets about the syntax or structure not only of the inhabitants of negative formulae (as suggested above) but also of positive formulae, abstracting over the logical connectives and the very syntax of formulae.

In the abstract framework that we present here, called LAF, an extra step of abstraction is also reached (compared to [Zei08a, Zei08b]) over the construction of (typing) contexts, which allows the same framework to capture both classical and intuitionistic systems. More substantially, the treatment of quantifiers is also new.

As the material developed and expanded, it also appeared that our framework, together with its machine-checked formalisation, could be directly implemented and serve as the theoretical foundations for a new version of the PSYCHE system, discussed in Part III of this dissertation. It could perhaps even serve as the basis for a formal proof of the system’s correctness. Thinking along those lines oriented the design of the LAF framework with implementation issues in mind (e.g. using De Bruijn’s indices or De Bruijn’s levels), and resulted in the formalisation of mathematical structures behind which the OCaml modules can clearly

be seen.

This Coq formalisation and the implementability concern also resulted in a presentation of the material that is admittedly technical, with e.g. numerous parameters and long specifications, which was also fuelled by the desire to identify the connection between focussing and orthogonality at the “purest” level: every design choice or ingredient of the framework that was not essential to establishing the connection was systematically turned into a parameter of the framework, with an axiomatisation for it that we sought to be as weak as possible for the theory to hold.

Chapter 4 presents a description of the proof-term system for big-step focussing that is more formal than that with which we concluded Part I of this dissertation. This formalisation, called LAF, is essentially a reformulation of the ideas in [Zei08a, Zei08b], with no substantial difference but the modular description of typing contexts. This allows classical and intuitionistic systems to be instances of the same parameterised system LAF, as we describe at the end of the chapter.

Chapter 5, on the other hand, presents a substantial extension: the LAF system with quantifiers. It therefore subsumes Chapter 4, but giving the version of LAF with quantifiers straight away would be a bit harsh on the reader.

Chapter 6 explores realisability models for LAF, based on orthogonality, and its contents was the original motivation for the development of this material, as the Adequacy Lemma connects big-step focussing with orthogonality, typing with realisability, syntax with semantics. We apply this methodology to derive the consistency of LAF systems.

Chapter 7 then investigates the operational semantics of LAF, which interprets the proof-terms for big-step focussing as a pattern-matching calculus. We first present a small-step semantics by means of an abstract machine for head-reduction. Adapting the methodology of Chapter 2, we apply the orthogonality models of Chapter 6 to prove the normalisation of typed terms with respect to this abstract machine. Then we develop the abstract machine into a big-step operational semantics, for which a new application of orthogonality models provides the cut-elimination result for LAF.

Chapter 4

An abstract focussed sequent calculus - without quantifiers

Contents

4.1	Presentation of the system	90
4.1.1	Atoms, molecules, typing decompositions and typing contexts	90
4.1.2	Logical connectives	92
4.1.3	Definition of the system	92
4.2	Capturing existing systems	93
4.3	Examples in propositional logic	95
4.3.1	Polarised classical logic - one-sided	95
4.3.2	Polarised classical logic - two-sided	99
4.3.3	Polarised intuitionistic logic	100
4.4	Examples of labels implementation: De Bruijn's indices and levels	104
4.4.1	Labels for classical logic	104
4.4.2	Labels for intuitionistic logic	105

In this chapter, we show how Zeilberger's ideas [Zei08a, Zei08b], as presented in Chapter 3, can be developed into an *abstract focussed sequent calculus* called LAF, and whose instances express the big-step versions of standard focussed sequent calculi.

The system of Chapter 3 is already abstract in the relation $\Vdash$ that decomposes a positive formula into a collection of positive atoms and negative formulae. Correspondingly, it is also abstract in the notion of *pattern* whose typing judgement is given by the relation $\Vdash$.

We push this abstraction further:

- Since this decomposition relation $\Vdash$ was the only ingredient of the system that used the syntax of formulae, we do not even have to assume that formulae are syntax, i.e. have an inductive structure, nor do we have to assume that “positive atoms” are particular kinds of formulae; positive atoms and formulae could literally be two arbitrary sets. We shall now respectively call them *atoms* and *molecules*.
- Moreover, a typing context Γ could be extended in an asynchronous step into Γ, Δ , where Δ is the result of decomposing some positive formula according to some pattern p and the

decomposition relation $\Vdash$. We have in fact no reason to assume that Γ and Δ are of the same nature and that Γ, Δ corresponds to set union (or whatever standard combination of typing contexts one usually considers). Therefore, “typing contexts” such as Γ will form an abstract notion, namely an algebra equipped with specific functions among which an arbitrary asymmetric construction $\Gamma; \Delta$ that replaces the above.¹

Something that is difficult to treat formally at this abstract level is the use of a non-deterministic way of naming variables, and then having to deal with α -conversion, in particular when we formalise our framework **LAF** and its meta-theory in the proof-assistant **Coq**. Therefore we adopt a deterministic way of naming variables (now called *labels* since they are not subject to α -conversion), but we remain abstract in the exact system that we use for naming them: this approach will capture for instance De Bruijn’s indices as well as De Bruijn’s levels.

Section 4.1 presents **LAF**. Section 4.3 describes how to tune (i.e. instantiate) the abstract parameters so as to capture different logics (or logical systems). Section 4.4 provide instances illustrating different implementations of labels corresponding to De Bruijn’s indices and De Bruijn’s levels.

4.1 Presentation of the system

This section presents the quantifier-free version of system **LAF**, a highly modular / parameterised sequent calculus for big-step focussing.

An instance of **LAF** is given by a tuple of parameters

$$(\mathbb{A}, \mathbb{M}, \text{Lab}_+, \text{Lab}_-, \text{Co}, \text{Pat}, \Vdash)$$

where each parameter is described below.

4.1.1 Atoms, molecules, typing decompositions and typing contexts

The first group of parameters $(\mathbb{A}, \mathbb{M})$ specifies what the instance of **LAF**, as a logical system, talks about. A typical example is when $\mathbb{A}$ and $\mathbb{M}$ are respectively the sets of (positive) atoms and the set of formulae from a polarised logic. We will see in the next sections how our level of abstraction allows for some interesting variants. In the Curry-Howard view, $\mathbb{A}$ and $\mathbb{M}$ are our sets of types.

DEFINITION 51 (Atoms & molecules)

LAF is parameterised by two sets $\mathbb{A}$ and $\mathbb{M}$, whose elements are respectively called *atoms* (denoted $a, a', \dots$), and *molecules* (denoted $M, M', \dots$). ✱

We then aim at defining *typing contexts*, those structures denoted Γ in a typing judgement of the form $\Gamma \vdash \dots$

Intuitively, we expect Γ to “contain” atoms and molecules, or more precisely to declare some variables as having atoms and molecules as their types.

For this it will be useful (e.g. to build models of **LAF**) to define contexts more generically, mapping variables to elements of two sets $\mathcal{A}$ and $\mathcal{B}$.

¹Following Zeilberger’s style, Δ itself will not be a typing context but will have a tree structure that may reflect the way a positive formula is decomposed into it.

Contexts will be extendable (in the case of typing contexts, we may want to extend Γ with a new type declaration for a fresh variable), and the following data-structure formalises what generic contexts will be extended with.

DEFINITION 52 (Generic decomposition algebras)

Given two sets $\mathcal{A}$ and $\mathcal{B}$, the $(\mathcal{A}, \mathcal{B})$ -decomposition algebra $\mathbb{D}_{\mathcal{A}, \mathcal{B}}$, whose elements are called $(\mathcal{A}, \mathcal{B})$ -decompositions, is the free algebra defined by the following grammar:

$$\Delta, \Delta_1, \dots ::= a \mid \sim b \mid \bullet \mid \Delta_1, \Delta_2$$

where a (resp. b) ranges over $\mathcal{A}$ (resp. $\mathcal{B}$).

Let $\mathbb{D}_{\text{st}}$ abbreviate $\mathbb{D}_{\text{unit}, \text{unit}}$, whose elements we call *decomposition structures*.

The *(decomposition) structure* of an $(\mathcal{A}, \mathcal{B})$ -decomposition Δ , denoted $|\Delta|$, is its obvious homomorphic projection in $\mathbb{D}_{\text{st}}$. ※

Intuitively, a $(\mathcal{A}, \mathcal{B})$ -decomposition Δ is simply the packaging of elements of $\mathcal{A}$ and elements of $\mathcal{B}$; we could flatten this packaging by seeing $\bullet$ as the empty set (or multiset), and Δ_1, Δ_2 as the union of the two sets (or multisets) Δ_1 and Δ_2 .

Note that the coercion from $\mathcal{B}$ into $\mathbb{D}_{\mathcal{A}, \mathcal{B}}$ is denoted with $\sim$. It helps distinguishing it from the coercion from $\mathcal{A}$ (e.g. when $\mathcal{A}$ and $\mathcal{B}$ intersect each other), and in many instances of LAF it will remind us of the presence of an otherwise implicit negation. But so far it has no logical meaning, and in particular $\mathcal{B}$ is not equipped with an operator $\sim$ of syntactical or semantical nature.

DEFINITION 53 (Generic contexts)

LAF is parameterised by two sets Lab_+ and Lab_- , of elements called *positive labels* and *negative labels*, respectively.

Given two sets $\mathcal{A}$ and $\mathcal{B}$, an $(\mathcal{A}, \mathcal{B})$ -context algebra is an algebra of the form

$$\left(\mathcal{G}, \left(\begin{array}{c} \mathcal{G} \times \text{Lab}_+ \rightarrow \mathcal{A} \\ (\Gamma, x^+) \mapsto \Gamma[x^+] \end{array} \right), \left(\begin{array}{c} \mathcal{G} \times \text{Lab}_- \rightarrow \mathcal{B} \\ (\Gamma, x^-) \mapsto \Gamma[x^-] \end{array} \right), \left(\begin{array}{c} \mathcal{G} \times \mathbb{D}_{\mathcal{A}, \mathcal{B}} \rightarrow \mathcal{G} \\ (\Gamma, \Delta) \mapsto \Gamma; \Delta \end{array} \right) \right)$$

whose elements are called $(\mathcal{A}, \mathcal{B})$ -contexts.

As $(\Gamma, x^+) \mapsto \Gamma[x^+]$ and $(\Gamma, x^-) \mapsto \Gamma[x^-]$ are partial functions, we denote by $\text{dom}^+(\Gamma)$ (resp. $\text{dom}^-(\Gamma)$) the subset of Lab_+ (resp. Lab_-) where $\Gamma[x^+]$ (resp. $\Gamma[x^-]$) is defined. ※

We choose to call elements of Lab_+ and Lab_- “labels”, rather than “variables”, because “variable” suggests an object identified by a name that “does not matter” and somewhere subject to α -conversion. For instance in the following typing rule for the (simply-typed) λ -calculus

$$\frac{\Gamma, x:A \vdash t:B}{\Gamma \vdash \lambda x.t: A \rightarrow B}$$

the α -convertibility of the variable x bound in $\lambda x.t$ relates to a non-deterministic choice of name for the variable used to extend the context Γ into $\Gamma, x:A$.² It turns out that such non-determinism in context extension is quite tricky to adapt (though probably not impossible) to the level of abstraction of LAF, and in practice would not be used in an implementation of proof-search, where a deterministic choice of name would be performed (“first fresh name” picking, etc).

²The fact that the non-deterministic choice does not matter, a.k.a. *equivariance*, is covered at length in nominal logic [Pit03] and other works formalising binding.

Therefore, we decide to present LAF without the non-determinism related to α -conversion, yet without committing to using De Bruijn's indices or De Bruijn's levels. Hence the use of “labels”, that will accommodate both systems (and others, as long as the concept of context extension $\Gamma; \Delta$ is a proper function, i.e. remains deterministic).

DEFINITION 54 (Typing decompositions and typing contexts)

The *typing decomposition algebra*, denoted $\mathbb{D}$, whose elements are called *typing decompositions*, is the $(\mathbb{A}, \mathbb{M})$ -decomposition algebra.

LAF is then parameterised by an $(\mathbb{A}, \mathbb{M})$ -context algebra Co , whose elements are called *typing contexts*. ※

4.1.2 Logical connectives

Finally, the last group of parameters $(\text{Pat}, \Vdash)$ specifies the structure of molecules. If $\mathbb{M}$ is a set of formulae featuring logical connectives, those parameters specify the introduction rules for the connectives.

DEFINITION 55 (Patterns & decomposition relation)

LAF is parameterised by a *pattern algebra*, an algebra of the form

$$\left(\text{Pat}, \left(\begin{array}{c} \text{Pat} \rightarrow \mathbb{D}_{\text{st}} \\ p \mapsto |p| \end{array} \right) \right)$$

whose elements are called *patterns*, and by a *decomposition relation*, i.e. a set of elements

$$(_ \Vdash _ : _) : (\mathbb{D} \times \text{Pat} \times \mathbb{M})$$

such that if $\Delta \Vdash p : M$ then the structure of Δ is $|p|$. ※

The intuition behind the terminology is that the decomposition relation $\Vdash$ decomposes a molecule, according to a pattern, into a typing decomposition which, as a first approximation, can be seen as a “collection of atoms and (hopefully smaller) molecules”.

4.1.3 Definition of the system

DEFINITION 56 (Proof-Terms)

Proof-terms are defined by the following syntax:

Positive terms	Terms^+	$t^+ ::= pd$
Decomposition terms	Terms^d	$d ::= x^+ \mid f \mid \bullet \mid d_1, d_2$
Commands	Terms	$c ::= \langle x^- \mid t^+ \rangle \mid \langle f \mid t^+ \rangle$

where p ranges over Pat , x^+ ranges over Lab_+ , x^- ranges over Lab_- , and f ranges over the partial function space $\text{Pat} \rightarrow \text{Terms}$. ※

We can finally present the typing system LAF:

DEFINITION 57 (LAF) LAF is the inference system of Fig. 23 defining the derivability of three kinds of sequents

$$\begin{array}{ll} (_ \vdash [_ : _]) & : (\text{Co} \times \text{Terms}^+ \times \mathbb{M}) \\ (_ \vdash _ : _) & : (\text{Co} \times \text{Terms}^d \times \mathbb{D}) \\ (_ \vdash _) & : (\text{Co} \times \text{Terms}) \end{array}$$

We further impose in rule **async** that the domain of function f be exactly those patterns that can decompose M (if $p \in \text{Dom}(f)$ then there exists Δ such that $\Delta \Vdash p : M$).

LAF^{cf} is the inference system LAF without the cut -rule.

※

$$\begin{array}{c}
 \frac{\Delta \Vdash p:M \quad \Gamma \vdash d:\Delta}{\Gamma \vdash [pd:M]} \text{sync} \\
 \\
 \frac{}{\Gamma \vdash \bullet:\bullet} \quad \frac{\Gamma \vdash d_1:\Delta_1 \quad \Gamma \vdash d_2:\Delta_2}{\Gamma \vdash d_1, d_2:\Delta_1, \Delta_2} \\
 \\
 \frac{\Gamma [x^+] = a}{\Gamma \vdash x^+:a} \text{init} \quad \frac{\forall p, \forall \Delta, \quad \Delta \Vdash p:M \Rightarrow \Gamma; \Delta \vdash f(p)}{\Gamma \vdash f:\sim M} \text{async} \\
 \\
 \frac{\Gamma \vdash [t^+:\Gamma [x^-]]}{\Gamma \vdash \langle x^- \mid t^+ \rangle} \text{select} \quad \frac{\Gamma \vdash f:\sim M \quad \Gamma \vdash [t^+:M]}{\Gamma \vdash \langle f \mid t^+ \rangle} \text{cut}
 \end{array}$$

Figure 23: LAF

An intuition of LAF can be given in terms of proof-search:

When we want to “prove” a molecule, we first need to decompose it into a collection of atoms and (refutations of) molecules (rule sync). Each of those atoms must be found in the current typing context (rule init). Each of those molecules must be refuted, and the way to do this is to consider all the possible ways that this molecule could be decomposed, and for each of those decompositions, prove the inconsistency of the current typing context extended with the decomposition (rule async). This can be done by proving one of the molecules refuted in the typing context (rule select) or refuted by a complex proof (rule cut). Then a new cycle begins again.

Typing decompositions and decomposition terms organise the packaging of the proofs of atoms and (refuted) molecules decomposed by rule sync . Typing decompositions could here be taken to be a multiset of atoms and (refuted) molecules, but keeping a dedicated structure for the packaging will be more convenient when we add quantifiers: giving decompositions an inductive structure allows a lossless modelling of quantifiers’ scopes.

4.2 Capturing existing systems

The above intuitions may become clearer when we instantiate the parameters of LAF with actual literals, formulae, etc in order to capture existing systems:

In the rest of this chapter we illustrate system LAF by specifying different instances, providing each time the long list of parameters, that capture different focussed sequent calculus systems.

By “capture”, we mean of course a stronger result than just the equivalence between the notions of provability. In order to strengthen such a weak property between two systems, it is relevant to consider the notions of *adequacy* as defined in [Nig09, NM10]:

The shallowest level of adequacy, *relative completeness*, or adequacy of level -1, requires that a sequent is provable in one system if and only if the sequent to which it is mapped is provable in the other system. Level -2 of adequacy, *full completeness of proofs*, requires that there be a one-to-one correspondence between their (complete) proofs. Level -3 of adequacy, *full completeness of derivations* (a word used in [Nig09, NM10] for incomplete proofs), requires a one-to-one correspondence between the derivations in one system and those of the other system.

Strictly speaking, level -2 adequacy does not say more than level -1 as soon as the sequent has infinitely and denumerably many proofs. With level -3 adequacy, we aim at capturing much more. The simplest way to formalise its informal description above, for a function ϕ that maps the sequents of system $\mathcal{A}$ into the sequents of system $\mathcal{B}$, is probably as follows:

For every sequent $\mathcal{S}$ and multiset $\{\mathcal{S}_1, \dots, \mathcal{S}_n\}$ of sequents in $\mathcal{A}$, there is a one-to-one correspondence $\phi_{\mathcal{S}, \{\mathcal{S}_1, \dots, \mathcal{S}_n\}}$ between

- *the partial proofs in $\mathcal{A}$ whose conclusion is $\mathcal{S}$ and whose multiset of open leaves is $\{\mathcal{S}_1, \dots, \mathcal{S}_n\}$*
- *the partial proofs in $\mathcal{B}$ whose conclusion is $\phi(\mathcal{S})$ and whose multiset of open leaves is $\{\phi(\mathcal{S}_1), \dots, \phi(\mathcal{S}_n)\}$*

The above is a symmetric property when ϕ is itself a one-to-one correspondence between sequents, but can also make sense if it is not. However, the above property needs to be adapted

- when in either of the two systems, we are interested not in each individual application of the inference rules but rather in groupings of rules: for instance in a focussed calculus, we may want to consider the grouping of a synchronous phase followed by an asynchronous phase (a.k.a. a *macro-rule* decomposing a *synthetic connective*) as a single step whose internal details should be ignored by the correspondence (this is what happens in [Nig09, NM10]);
- when either of the two systems features proof-terms, as the notion of incomplete proof is polluted by the presence, in the sequent, of a proof-term denoting a complete proof (unless we start considering incomplete proof-terms as well).

Both situations jeopardise the bijective aspect of each $\phi_{\mathcal{S}, \{\mathcal{S}_1, \dots, \mathcal{S}_n\}}$: in the former situation, we probably want to quotient proofs in some way so that the internal details of a rule grouping do not lead to multiple proofs that are not reflected in the other system ([NM10] mentions for instance “up to the permutation of asynchronous rules”); in the latter situation, proof-term annotations would provide for instance two proofs of $x:A, y:A \vdash ? : A$ while we only count one proof of $A, A \vdash A$ (whether A, A denotes a set or a multiset).

Another issue with the above notion of adequacy is that it fails to impose any notion of compositionality (when derivations are “plugged into” the open leaves of another derivation) about the family $(\phi_{\mathcal{S}, \{\mathcal{S}_1, \dots, \mathcal{S}_n\}})_{\mathcal{S}, \{\mathcal{S}_1, \dots, \mathcal{S}_n\}}$, something which we may have in mind when thinking about the “deepest level of adequacy”.

System LAF features both focussing and proof-terms. Rather than trying to adapt to these concepts, and strengthen with compositionality, the above formalisation of level -3 adequacy, we opt for a version of level -3 adequacy that drops the use of bijections and the quantitative aspects that they provide (we no longer try to count proofs). On the other hand, we retain from level -3 adequacy, and formalise, the fact that the **structure** of proofs in one system matches the structure of proofs in the other system.

DEFINITION 58 (Structural adequacy)

- Let $\mathcal{A}$ be an inference system providing a notion of proof-trees for elements called “sequents”, and let $\mathcal{P}$ be a set of sequents.

Given a proof-tree π in $\mathcal{A}$, the multiset of $\mathcal{P}$ -*immediate sequents* of π is defined recursively on π : it contains the conclusions that are in $\mathcal{P}$ of the direct sub-trees of π , as well as the $\mathcal{P}$ -immediate sequents of the direct sub-trees of π whose conclusions are not in $\mathcal{P}$.

- Let $\mathcal{A}$ and $\mathcal{B}$ be two inference systems as in the previous point, and $\mathcal{R}$ be a relation between the sequents of $\mathcal{A}$ and the sequents of $\mathcal{B}$, with domain $\mathcal{D}$ and co-domain $\mathcal{C}$.

$\mathcal{R}$ satisfies *structural adequacy* if, whenever $\mathcal{S}\mathcal{R}\mathcal{S}', \mathcal{S}_1\mathcal{R}\mathcal{S}'_1 \dots, \mathcal{S}_n\mathcal{R}\mathcal{S}'_n$,

there is in $\mathcal{A}$ a proof of $\mathcal{S}$ with $\mathcal{D}$ -immediate sequents $\{\mathcal{S}_1, \dots, \mathcal{S}_n\}$

if and only if

there is in $\mathcal{B}$ a proof of $\mathcal{S}'$ with $\mathcal{C}$ -immediate sequents $\{\mathcal{S}'_1, \dots, \mathcal{S}'_n\}$

※

Structural adequacy clearly entails level -1 adequacy (by induction on a proof in $\mathcal{A}$, recursively finding its $\mathcal{D}$ -immediate sequents, we recompose a proof in $\mathcal{B}$), but implies neither level -2 nor level -3 since we are not counting proofs. Also notice that we do not require anything about incomplete proofs that cannot be completed.

Every instance below relates to a traditional system, as we define an encoding satisfying structural adequacy. While LAF is defined as a typing system (in other words with proof-terms decorating proofs in the view of the Curry-Howard correspondence), most traditional systems that we capture below are purely logical, with no proof-term decorations. The encoding therefore needs to erase proof-term annotation, and for this it is useful to project the notion of typing context as follows:

DEFINITION 59 (Referable atoms and molecules)

Let $\text{Im}^+(\Gamma)$ (resp. $\text{Im}^-(\Gamma)$) be the image set of $x^+ \mapsto \Gamma[x^+]$ (resp. $x^+ \mapsto \Gamma[x^+]$), i.e. the set of atoms (resp. molecules) that can be referred to, in Γ , by the use of a positive (resp. negative) label.

※

4.3 Examples in propositional logic

The parameters of LAF will be specified so as to capture: the one-sided version of LKF [LM09, LM11], its two-sided version, and LJF [LM09].

4.3.1 Polarised classical logic - one-sided

In this sub-section we define the instance LAF_{K1} corresponding to the one-sided focused sequent calculus LKF for polarised classical logic [LM09, LM11].

DEFINITION 60 (Literals, formulae, patterns, decomposition)

Let $\mathbb{L}$ be a set of elements called *literals*, equipped with an involutive function called *negation* mapping every literal l to a literal $l^\perp$.

Let $\mathbb{A}$ be a *polarisation set*, i.e. a subset of $\mathbb{L}$ such that $l \in \mathbb{A}$ if and only if $l^\perp \notin \mathbb{A}$. Elements of $\mathbb{A}$ will be ranged over by $a, a', \dots$

$$\begin{array}{c}
\frac{}{\bullet \Vdash \bullet : \top^+} \quad \frac{}{\sim N^\perp \Vdash _ : N} \quad \frac{}{a \Vdash _ : a} \\
\\
\frac{\Delta_1 \Vdash p_1 : A_1 \quad \Delta_2 \Vdash p_2 : A_2}{\Delta_1, \Delta_2 \Vdash (p_1, p_2) : A_1 \wedge^+ A_2} \quad \frac{\Delta \Vdash p : A_i}{\Delta \Vdash \text{inj}_i(p) : A_1 \vee^+ A_2}
\end{array}$$

Figure 24: Decomposition relation for LAF_{K1}

Let $\mathbb{M}$ be the set defined by the first line of the following grammar for (polarised) formulae of classical logic:

$$\begin{array}{lll}
\text{Positive formulae} & P, \dots & ::= a \mid \top^+ \mid \perp^+ \mid A \wedge^+ B \mid A \vee^+ B \\
\text{Negative formulae} & N, \dots & ::= a^\perp \mid \top^- \mid \perp^- \mid A \wedge^- B \mid A \vee^- B \\
\text{Unspecified formulae} & A & ::= P \mid N
\end{array}$$

Negation is extended to formulae as follows:

$$\begin{array}{llll}
\top^{+\perp} & := \perp^- & \top^{-\perp} & := \perp^+ \\
\perp^{+\perp} & := \top^- & \perp^{-\perp} & := \top^+ \\
(A \wedge^+ B)^\perp & := A^\perp \vee^- B^\perp & (A \wedge^- B)^\perp & := A^\perp \vee^+ B^\perp \\
(A \vee^+ B)^\perp & := A^\perp \wedge^- B^\perp & (A \vee^- B)^\perp & := A^\perp \wedge^+ B^\perp
\end{array}$$

and we extend it to sets or multisets of formulae pointwise.

The set Pat of *pattern* is defined by the following grammar:

$$p, p_1, p_2, \dots ::= _+ \mid _- \mid \bullet \mid (p_1, p_2) \mid \text{inj}_i(p)$$

The decomposition relation $(_ \Vdash _ : _) : (\mathbb{D} \times \text{Pat} \times \mathbb{M})$ is the restriction to molecules of the relation of $\mathbb{D} \times \text{Pat} \times \mathbb{F}$ defined inductively for all formulae by the inference system of Fig. 24.

The map $p \mapsto |p|$ can be inferred from the decomposition relation. *

Keeping the `sync` rule of LAF_{K1} in mind, we can already see in Fig. 24 the traditional introduction rules of positive connectives in polarised classical logic. The rest of this sub-section formalises that intuition and explains how LAF_{K1} manages the introduction of negative connectives, etc.

But in order to finish the instantiation of LAF for propositional polarised classical logic (1-sided), we need to define typing contexts, i.e. give Lab_+ , Lab_- , and Co . In particular, we have to decide how to refer to elements of the typing context. To avoid getting into aspects that may be considered as implementation details, we will stay rather generic and only assume the following property:

DEFINITION 61 (Typing contexts) We assume

$$\begin{array}{llll}
\text{lm}^+(\Gamma; a) & = \text{lm}^+(\Gamma) \cup \{a\} & \text{lm}^-(\Gamma; a) & = \text{lm}^-(\Gamma) \\
\text{lm}^+(\Gamma; \sim M) & = \text{lm}^+(\Gamma) & \text{lm}^-(\Gamma; \sim M) & = \text{lm}^-(\Gamma) \cup \{M\} \\
\text{lm}^\pm(\Gamma; \bullet) & = \text{lm}^\pm(\Gamma) & \text{lm}^\pm(\Gamma; (\Delta_1, \Delta_2)) & = \text{lm}^\pm(\Gamma; \Delta_1; \Delta_2)
\end{array}$$

where $\pm$ stands for either $+$ or $-$. *

In section 4.4 we present several implementations satisfying the above.

We now relate (cut-free) $\text{LAF}_{K1}^{\text{cf}}$ and the LKF system of [LM09, LM11].

DEFINITION 62 (Flattening typing decompositions) Let $\bar{\Delta}$ be the flattening of a typing decomposition as a multiset of positive literals and negative formulae, i.e.

$$\begin{aligned} \bar{a} &:= \{\{a\}\} & \overline{\sim P} &:= \{\{P^\perp\}\} \\ \bar{\bullet} &:= \emptyset & \overline{\Delta_1, \Delta_2} &:= \overline{\Delta_1 \cup \Delta_2} \end{aligned}$$

※

REMARK 46

- Notice that, for all formulae A and typing decomposition Δ , there exists $p \in \text{Pat}$ such that $\Delta \Vdash p:A$ if and only if $A \downarrow \bar{\Delta}$ as defined in [LM11].
- Our assumption about typing contexts implies that, for all Γ and Δ ,

$$\text{Im}^+(\Gamma; \Delta) \cup \text{Im}^-(\Gamma; \Delta)^\perp = \text{Im}^+(\Gamma) \cup \text{Im}^-(\Gamma)^\perp \cup \bar{\Delta}$$

※

DEFINITION 63 (Mapping sequents)

We encode the sequents of LAF_{K1} (regardless of derivability) to those of LKF as follows:

$$\begin{aligned} \phi(\Gamma \vdash c) &:= \vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \uparrow \\ \phi(\Gamma \vdash x^+ : a) &:= \vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \Downarrow a \\ \phi(\Gamma \vdash f : \sim P) &:= \vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \Downarrow P^\perp \\ \phi(\Gamma \vdash [t^+ : P]) &:= \vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \Downarrow P \end{aligned}$$

※

THEOREM 47 (Adequacy between $\text{LAF}_{K1}^{\text{cf}}$ and LKF)

ϕ satisfies structural adequacy between $\text{LAF}_{K1}^{\text{cf}}$ and LKF.

※

Proof: The Lemmata 2 and 3 of [LM11] (for the particular case of LKF) provide the correspondence with the big-step rules of $\text{LAF}_{K1}^{\text{cf}}$:

async Clearly, a derivation in $\text{LAF}_{K1}^{\text{cf}}$ concludes $\Gamma \vdash f : \sim P$ for some term f if and only if it is of the form

$$\frac{\Gamma; \Delta_1 \vdash c_1 \quad \dots \quad \Gamma; \Delta_n \vdash c_n}{\Gamma \vdash f : \sim P}$$

for some commands $\{c_1, \dots, c_n\}$, and where $\{\Delta_1, \dots, \Delta_n\} = \{\Delta \mid \exists p, \Delta \Vdash p:P\}$.

Correspondingly, Lemma 2 of [LM11]³ entails that a derivation in LKF concludes $\vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \Downarrow P^\perp$ if and only if it is of the form

$$\frac{\frac{\vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma), \Phi_1^\perp \uparrow}{\vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma), \Phi_n^\perp \uparrow} \quad \vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma), \Phi_n^\perp \uparrow}{\vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \uparrow P^\perp} \quad \vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma) \Downarrow P^\perp$$

where $\{\Phi_1, \dots, \Phi_n\} = \{\Phi \mid P \downarrow \Phi\}$.

Writing ϕ for the bijection from $1, \dots, n$ to itself such that $\bar{\Delta}_i = \Phi_{\phi(i)}$, we notice that every sequent $\Gamma; \Delta_i \vdash c_i$ is mapped to the sequent $\vdash \text{Im}^+(\Gamma)^\perp, \text{Im}^-(\Gamma), \Phi_{\phi(i)}^\perp \uparrow$. Indeed, Remark 46.2 entails that

³slightly reworded using its Lemma 4 as well

$$\text{lm}^+(\Gamma; \Delta_i)^\perp \cup \text{lm}^-(\Gamma; \Delta_i) = \text{lm}^+(\Gamma)^\perp \cup \text{lm}^-(\Gamma) \cup \Phi_{\phi(i)}^\perp$$

sync

Clearly, a derivation in $\text{LAF}_{K1}^{\text{cf}}$ concludes $\Gamma \vdash [t^+ : P]$ for some term t^+ if and only if it is of the form

$$\frac{\Delta \Vdash p : P \quad \frac{\frac{\Gamma \vdash t_1^- : u_1}{\vdots} \quad \frac{\Gamma \vdash t_n^- : u_n}{\vdots}}{\Gamma \vdash \sigma : \Delta}}{\Gamma \vdash [p\sigma : P]}$$

for some $\Delta, p, \sigma, t_1^-, \dots, t_n^-$, and where $\overline{\Delta} = \{u_1, \dots, u_n\}$.

Correspondingly, Lemma 3 of [LM11] entails that a derivation in LKF concludes $\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow P$ if and only if it is of the form

$$\frac{\frac{\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow u_1}{\vdots} \quad \frac{\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow u_n}{\vdots}}{\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow P}$$

for some $P \Downarrow \{u_1, \dots, u_n\}$.

init Clearly, a derivation in $\text{LAF}_{K1}^{\text{cf}}$ concludes $\Gamma \vdash x^+ : a$ for some positive label x^+ if and only if it is of the form

$$\overline{\Gamma \vdash x^+ : a}$$

with $a \in \text{lm}^+(\Gamma)$.

Correspondingly, a derivation in LKF concludes $\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow a$ if and only if it is of the form

$$\overline{\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow a}$$

with $a \in \text{lm}^+(\Gamma)$.

select Clearly, a derivation in $\text{LAF}_{K1}^{\text{cf}}$ concludes $\Gamma \vdash c$ for some command c if and only if it is of the form

$$\frac{\Gamma \vdash [t^+ : P]}{\Gamma \vdash \langle x^- \mid t^+ \rangle}$$

with $P \in \text{lm}^-(\Gamma)$.

Correspondingly, a derivation in LKF concludes $\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Uparrow$ if and only if it is of the form

$$\frac{\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Downarrow P}{\vdash \text{lm}^+(\Gamma)^\perp, \text{lm}^-(\Gamma) \Uparrow}$$

with $P \in \text{lm}^-(\Gamma)$.

□

COROLLARY 48 (Equivalence of provability)

⌊ The provability of a sequent in $\text{LAF}_{K1}^{\text{cf}}$ is the same as that of its encoding in LKF.

※

The proof may raise the question of why, in the definition of LAF , we gave a structure to typing decompositions, instead of directly using a flattened version (e.g. multiset). The reason is to allow the parametrisation of the system so as to capture logics for which the structure of typing decomposition may be important; if only for first-order logic, the scope of eigenvariables is more easily managed with a structure; this is even more true in higher-order logic.

4.3.2 Polarised classical logic - two-sided

Having seen how an instance of LAF captures a one-sided sequent calculus, we could see LAF itself as a sequent calculus that is intrinsically one-sided, considering as a notational idiosyncrasy our writing the typing environments on the left of the turnstyle.

Here, we show that, by enriching the atoms and molecules with a “side information”, we can also capture a two-sided version of LKF .

DEFINITION 64 (Literals, formulae, patterns, decomposition)

Let $\mathbb{L}^+$ (resp. $\mathbb{L}^-$) be a set of elements called positive (resp. negative) literals, and ranged over by $l^+, l_1^+, l_2^+, \dots$ (resp. $l^-, l_1^-, l_2^-, \dots$).

Formulae are defined by the following grammar:

$$\begin{array}{lll} \text{Positive formulae} & P, \dots & ::= l^+ \mid \top^+ \mid \perp^+ \mid A \wedge^+ B \mid A \vee^+ B \mid \neg^+ A \\ \text{Negative formulae} & N, \dots & ::= l^- \mid \top^- \mid \perp^- \mid A \wedge^- B \mid A \vee^- B \mid \neg^- A \\ \text{Unspecified formulae} & A & ::= P \mid N \end{array}$$

We *position* a literal or a formula on the left-hand side or the right-hand side of a sequent by combining it with an element, called *side information*, of the set $\{l, r\}$: we define

$$\begin{array}{ll} \mathbb{A} & ::= \{(l^+, r) \mid l^+ \text{ positive literal}\} \cup \{(l^-, l) \mid l^- \text{ negative literal}\} \\ \mathbb{M} & ::= \{(P, r) \mid P \text{ positive formula}\} \cup \{(N, l) \mid N \text{ negative formula}\} \end{array}$$

The set Pat of *patterns* is defined by the following grammar:

$$\begin{array}{ll} p, p_1, p_2, \dots & ::= \frac{_r^+ \mid _r^-}{_l^+ \mid _l^-} \mid \bullet_r \mid (p_1, p_2) \mid \text{inj}_i(p) \mid \curvearrowright(p) \\ & \mid \frac{_r^+ \mid _r^-}{_l^+ \mid _l^-} \mid \bullet_l \mid [p_1, p_2] \mid \pi_i(p) \mid \curvearrowleft(p) \end{array}$$

The decomposition relation $(_ \Vdash _ : _) : (\mathbb{D} \times \text{Pat} \times \mathbb{M})$ is the restriction to molecules of the relation of $\mathbb{D} \times \text{Pat} \times (\mathbb{F} \times \{l, r\})$ defined inductively for all positioned formulae by the inference system of Fig. 25.

Again, since we want to capture classical logic, we assume the same property about $(\text{Lab}_+, \text{Lab}_-, \text{Co})$ as we did in Definition 61. ※

Keeping the sync rule of LAF_{K2} in mind, we see in Fig. 25 the traditional right-introduction rules of positive connectives and left-introduction rules of negative connectives.

A deeper intuition can be given by encoding LAF_{K2} sequents as two-sided sequents, just like we encoded LAF_{K1} sequents as one-sided LKF sequents:

$$\begin{array}{c}
\frac{}{\sim(N, l) \Vdash \neg_r : (N, r)} \quad \frac{}{(l^+, r) \Vdash \neg_r^+ : (l^+, r)} \\
\frac{}{\bullet \Vdash \bullet_r : (\top^+, r)} \quad \frac{\Delta \Vdash p : (A, l)}{\Delta \Vdash \neg(p) : (\neg^+ A, r)} \\
\frac{\Delta_1 \Vdash p_1 : (A_1, r) \quad \Delta_2 \Vdash p_2 : (A_2, r)}{\Delta_1, \Delta_2 \Vdash (p_1, p_2) : (A_1 \wedge^+ A_2, r)} \quad \frac{\Delta \Vdash p : (A_i, r)}{\Delta \Vdash \text{inj}_i(p) : (A_1 \vee^+ A_2, r)} \\
\frac{}{\sim(P, r) \Vdash \neg_l : (P, l)} \quad \frac{}{(l^-, l) \Vdash \neg_l^+ : (l^-, l)} \\
\frac{}{\bullet \Vdash \bullet_l : (\perp^-, l)} \quad \frac{\Delta \Vdash p : (A, r)}{\Delta \Vdash \neg(p) : (\neg^- A, l)} \\
\frac{\Delta_1 \Vdash p_1 : (A_1, l) \quad \Delta_2 \Vdash p_2 : (A_2, l)}{\Delta_1, \Delta_2 \Vdash [p_1, p_2] : (A_1 \vee^- A_2, l)} \quad \frac{\Delta \Vdash p : (A_i, l)}{\Delta \Vdash \pi_i(p) : (A_1 \wedge^- A_2, l)}
\end{array}$$

Figure 25: Decomposition relation for LAF_{K2} **DEFINITION 65 (LAF_{K2} sequents as two-sided sequents)**

1. First, when $\pm$ is either $+$ or $-$, we define

$$\begin{aligned}
\text{Im}^{\pm r}(\Gamma) &:= \{A \mid (A, r) \in \text{Im}^\pm(\Gamma)\} \\
\text{Im}^{\pm l}(\Gamma) &:= \{A \mid (A, l) \in \text{Im}^\pm(\Gamma)\}
\end{aligned}$$

2. Then we define the encoding:

$$\phi(\Gamma \vdash c) \quad := \quad \text{Im}^{+r}(\Gamma), \text{Im}^{-l}(\Gamma) \vdash \text{Im}^{+l}(\Gamma), \text{Im}^{-r}(\Gamma)$$

※

Keeping the above interpretation of sequents in mind, we should now see how to develop the details of the correspondence (similar to that expressed in Theorem 47) between $\text{LAF}_{K2}^{\text{cf}}$ and the two-sided version of LKF (which may actually not be written down in the literature).

As we can see, the decomposition relation, and the whole inference system described by LAF_{K2} , is completely symmetric.

4.3.3 Polarised intuitionistic logic**DEFINITION 66 (Literals, formulae, patterns, decomposition)**

Let $\mathbb{L}^+$ (resp. $\mathbb{L}^-$) be a set of elements called positive (resp. negative) literals, and ranged over by $l^+, l_1^+, l_2^+, \dots$ (resp. $l^-, l_1^-, l_2^-, \dots$).

Formulae are defined by the following grammar:

$$\begin{array}{lll}
\text{Positive formulae} & P, \dots & ::= l^+ \mid \top^+ \mid \perp^+ \mid A \wedge^+ B \mid A \vee B \\
\text{Negative formulae} & N, \dots & ::= l^- \mid \top^- \mid \perp^- \mid A \wedge^- B \mid A \Rightarrow B \mid \neg A \\
\text{Unspecified formulae} & A & ::= P \mid N
\end{array}$$

$$\begin{aligned} \mathbb{A} &:= \{(l^+, r) \mid l^+ \text{ positive literal}\} \cup \{(l^-, l) \mid l^- \text{ negative literal}\} \cup \{(\perp^-, l)\} \\ \mathbb{M} &:= \{(P, r) \mid P \text{ positive formula}\} \cup \{(N, l) \mid N \text{ negative formula}\} \end{aligned}$$

The set Pat of *pattern* is defined by the following grammar:

The decomposition relation $(_ \Vdash _ : _) : (\mathbb{D} \times \mathbf{Pat} \times \mathbb{M})$ is the restriction (to molecules) of the relation of $\mathbb{D} \times \mathbf{Pat} \times (\mathbb{F} \times \{\mathbf{l}, \mathbf{r}\})$ defined inductively for all positioned formulae by the inference system of Fig. 26. *

Figure 26: Decomposition relation for LAF_J

With implication as a primitive connective, we could actually remove the (negative) negation from the system, since it can in turn be seen as the combination of implication and absurdity ($\neg A$ can be seen as the abbreviation for $A \Rightarrow \perp^-$) and its decomposition rule reflects this. Notice that the decomposition rule for $\perp^-$ (and therefore that of $\neg$) are slightly modified compared to $K2$. To understand this, we should start by making the following remark:

REMARK 49

1. Whenever $\Delta \vdash p:(A, r)$, Δ contains no items of the form $\sim(P, r)$ or (v, l) .
2. Whenever $\Delta \vdash p:(A, l)$, Δ contains **exactly one** item of the form $\sim(P, r)$ or (v, l) .
(v stands for either a negative literal l^- or $\perp^-$).

※

The first point corresponds to the fact that, when we have a right-hand side focus in intuitionistic logic, the focus never switches to the left-hand side when looking at a proof-tree bottom-up. Notice that this would be false in presence of the positive negation, which would precisely switch the focus to the left-hand side as in $K2$.

Now the second point would not hold if we kept the negative disjunction from $K2$, since its decomposition rule would create a branching with two premisses of the form (v, l) . Hence its replacement with implication, whose decomposition rule has only one premiss of that form, so that, in every derivation of the above inference system, at most one branch keeps decomposing formulae on the left. And that would be true with the $K2$ rules for $\perp^-$ and $\neg^-$. The reason to tweak them is to get point 2 with *exactly one* rather than *at most one*, and it is for this tweak that we added $(\perp^-, l)$ to $\mathbb{A}$ (compared to the $K2$ version).

To see why Remark 49.2 is so important for intuitionistic logic, we should now interpret LAF_{K2} sequents as intuitionistic sequents (from e.g. LJF [LM09]):

DEFINITION 67 (LAF_J sequents as two-sided LJF sequents)

1. First, when $\pm$ is either $+$ or $-$, we define

$$\begin{aligned} \text{lm}^{\pm r}(\Gamma) &:= \{A \mid (A, r) \in \text{lm}^{\pm}(\Gamma)\} \\ \text{lm}^{+l}(\Gamma) &:= \{l^- \mid (l^-, l) \in \text{lm}^+(\Gamma)\} \\ \text{lm}^{-l}(\Gamma) &:= \{N \mid (N, l) \in \text{lm}^-(\Gamma)\} \end{aligned}$$

2. Then we define the encoding:

$$\begin{aligned} \phi(\Gamma \vdash c) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)] \longrightarrow [\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)] \\ \phi(\Gamma \vdash x^+:(l^-, l)) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)] \xrightarrow{l^-} [\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)] \\ \phi(\Gamma \vdash f:\sim(P, r)) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)] \xrightarrow{P} [\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)] \\ \phi(\Gamma \vdash [t^+:(N, l)]) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)] \xrightarrow{N} [\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)] \\ \phi(\Gamma \vdash x^+:(l^+, r)) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)]_{-l^+} \rightarrow \\ \phi(\Gamma \vdash f:\sim(N, l)) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)]_{-N} \rightarrow \\ \phi(\Gamma \vdash [t^+:(P, r)]) &:= [\text{lm}^{+r}(\Gamma), \text{lm}^{-l}(\Gamma)]_{-P} \rightarrow \end{aligned}$$

In the first four cases, we require $\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)$ to be a singleton (or be empty).

※

The first line of the encoding is the same as for LAF_{K2} (Definition 65), but for the fact that we require $\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)$ to be a singleton (or be empty), since we are to capture an intuitionistic system such as LJF. We also see in the last three cases (when there is a right-hand side focus), that the encoding forgets $\text{lm}^{+l}(\Gamma), \text{lm}^{-r}(\Gamma)$ altogether. If it is not empty, then it should definitely play no further role in the proof of the LAF_J sequent.

The issue arises in particular when analysing the **select** rule:

In LJF, placing the focus on a formula on the left-hand side does not affect the formula stored on the right-hand side; on the contrary, placing the focus on the right-hand side formula removes it from the right-hand side (no backup copy is made).

This is an important feature of intuitionistic logic: if a backup copy of the formula was kept, we could place again the focus on it further up in the proof, and we could thus prove formulae such as $A \vee \neg A$ or the drinker's theorem; indeed this amounts to authorising contraction on the right.

Now looking at the **select** rule of LAF_J , notice that the typing context Γ is **unchanged** by the rule: placing the focus on a right-hand side formula (i.e. a formula from $\text{lm}^+(\Gamma), \text{lm}^-(\Gamma)$) does not remove it from the typing context.

We could therefore fear that, because of this feature, **LAF** forces the presence of contraction (left and right) and is therefore intrinsically classical. Fortunately, this is not the case:

After selecting a right-hand side formula for focus, it is decomposed according to the rules of the decomposition relation. As mentioned in Remark 49.1, the focus never switches to the left-hand side and we are therefore left to prove a collection of sequents of the form $\Gamma \vdash x^+ : (v, r)$ or $\Gamma \vdash f : \sim(N, l)$ for some x^+ or f to be found. In the former case, the part of Γ that stores the unfortunate backup copy of the right-hand side formula that was selected for focus, does not affect whether $(v, r) \in \text{lm}^+(\Gamma)$. In the latter case, only rule **async** can be applied and a sequent of the form $\Gamma; \Delta \vdash c$ is left to be proved (for some c to be found) for every Δ that can decompose (N, l) . For the adequacy with intuitionistic logic to work, it suffices that for every such Δ , the operation $\Gamma; \Delta$ erases from Γ the unfortunate backup copy of the right-hand side formula that was selected for focus. According to Remark 49.2, every such Δ contains **exactly one** item of the form (v, l) or $\sim(P, r)$, i.e. a new right-hand side formula which can overwrite the old one. At least, provided that $(\text{Lab}_+, \text{Lab}_-, \text{Co})$ are defined to do that job.

Having tweaked the decomposition rule for $(\perp^-, l)$ to guarantee Remark 49.2, $(\text{Lab}_+, \text{Lab}_-, \text{Co})$ should also make sure that, for any Γ , the (focussed) sequent $\Gamma \vdash [t^+ : (\perp^-, l)]$ can still be proved for some t^+ to be found, i.e. the sequent $\Gamma \vdash x^+ : (\perp^-, l)$ can be proved for some x^+ to be found, i.e. $(\perp^-, l) \in \text{lm}^+(\Gamma)$ (even when Γ is interpreted as something completely empty). This is easy to do, by having a permanent and special label $x_{(\perp^-, l)}^+ \in \text{Lab}_+$ mapped to $(\perp^-, l)$ in every Γ . This is the same as permanently adding $\perp^-$ on the right-hand side of intuitionistic sequents (as some kind of multi-conclusion), lest that right-hand side ever gets empty: it is harmless to both the intuitionistic provability and the structural theory of proofs (none are added, none are removed).

DEFINITION 68 (Typing contexts)

We assume that we always have $(\perp^-, l) \in \text{lm}^+(\Gamma)$ and that

$$\begin{array}{llll} \text{lm}^+(\Gamma; (l^+, r)) &= \text{lm}^+(\Gamma) \cup \{(l^+, r)\} & \text{lm}^-(\Gamma; a) &= \text{lm}^-(\Gamma) \\ \text{lm}^+(\Gamma; \sim M) &= \text{lm}^+(\Gamma) & \text{lm}^-(\Gamma; \sim(N, l)) &= \text{lm}^-(\Gamma) \cup \{(N, l)\} \\ \text{lm}^\pm(\Gamma; \bullet) &= \text{lm}^\pm(\Gamma) & \text{lm}^\pm(\Gamma; (\Delta_1, \Delta_2)) &= \text{lm}^\pm(\Gamma; \Delta_1; \Delta_2) \\ \text{lm}^+(\Gamma; (v, l)) &= \{(l^+, r) \mid (l^+, r) \in \text{lm}^+(\Gamma)\} \cup \{(v, l), (\perp^-, l)\} \\ \text{lm}^-(\Gamma; \sim(P, r)) &= \{(N, l) \mid (N, l) \in \text{lm}^-(\Gamma)\} \cup \{(P, r)\} \end{array}$$

where again $\pm$ stands for either $+$ or $-$ and v stands for either a negative literal l^- or $\perp^-$. $\ast$

The first three lines are the same as those assumed for $K1$ and $K2$, except it is restricted to those cases where we do not try to add to Γ an atom or a molecule that is interpreted as going to the right-hand side of a sequent. When we want to do that, this atom or molecule should overwrite the previous atom(s) or molecule(s) that was (were) interpreted as being on the right-hand side; this is done in the last two lines, where $\text{lm}^+(\Gamma), \text{lm}^-(\Gamma)$ is completely

erased.

THEOREM 50 (Adequacy between LAF_f^{cf} and LJF)

ϕ satisfies structural adequacy between LAF_f^{cf} and LJF. ※

Proof: The details are similar to those of Theorem 47, relying again on the LJF properties expressed in [LM09, LM11] and following the series of remarks and design decisions that were made above. □

4.4 Examples of labels implementation: De Bruijn's indices and levels

In this section we give some concrete implementations of labels to completely specify the typing context algebras used in the examples of the previous section.

4.4.1 Labels for classical logic

In the instances LAF_{K1} and LAF_{K2} , we have simply made some assumptions on the typing context algebra (in Definition 61). We now give it a full definition satisfying these assumptions and using of De Bruijn's indices.

In fact, we generically build an $(\mathcal{A}, \mathcal{B})$ -context for each pair of sets $\mathcal{A}$ and $\mathcal{B}$, and the typing context algebra will simply be the instance where $\mathcal{A} = \mathbb{A}$ and $\mathcal{B} = \mathbb{M}$.

DEFINITION 69 (Generic context algebras with De Bruijn's indices - classical)

Given two sets $\mathcal{A}$ and $\mathcal{B}$, we define an $(\mathcal{A}, \mathcal{B})$ -context algebra $\text{Co}_{\mathcal{A}, \mathcal{B}}$ as follows:

An $(\mathcal{A}, \mathcal{B})$ -context Γ is a pair (Γ^+, Γ^-) where Γ^+ is a list of elements of $\mathcal{A}$ and Γ^- is a list of elements of $\mathcal{B}$.

Extensions are defined as follows:

$$\begin{aligned} (\Gamma^+, \Gamma^-); a &:= (a :: \Gamma^+, \Gamma^-) & (\Gamma^+, \Gamma^-); \sim b &:= (\Gamma^+, b :: \Gamma^-) \\ (\Gamma^+, \Gamma^-); \bullet &:= (\Gamma^+, \Gamma^-) & (\Gamma^+, \Gamma^-); (\Delta_1, \Delta_2) &:= (\Gamma^+, \Gamma^-); \Delta_1; \Delta_2 \end{aligned}$$

Positive labels and negative labels are two disjoint copies of the set of integers, with elements denoted n^+ and n^- , and we define

$(\Gamma^+, \Gamma^-)[n^+]$ as the $(n+1)^{\text{th}}$ element of Γ^+

$(\Gamma^+, \Gamma^-)[n^-]$ as the $(n+1)^{\text{th}}$ element of Γ^- . ※

These are indeed De Bruijn's indices, since the element accessed by label 0^+ (resp. 0^-) is the head of the list Γ^+ (resp. Γ^-), i.e. the element of the list that has last been added.

Alternatives using De Bruijn's indices are possible:

- the choice we made, when extending a context with (Δ_1, Δ_2) , of first extending the context with Δ_1 and then extending the result with Δ_2 , was completely arbitrary, we could have defined

$$(\Gamma^+, \Gamma^-); (\Delta_1, \Delta_2) := (\Gamma^+, \Gamma^-); \Delta_2; \Delta_1$$

- we could have defined an $(\mathcal{A}, \mathcal{B})$ -context Γ as one single list of atoms and molecules, with $\text{Lab}_+ = \text{Lab}_- = \mathbb{N}$ and $\Gamma[n^+]$ (resp. $\Gamma[n^-]$) being defined only on those integers mapped to atoms (resp. molecules);

- we could have defined an $(\mathcal{A}, \mathcal{B})$ -context Γ as a list of $(\mathcal{A}, \mathcal{B})$ -decompositions, with

$$\Gamma; \Delta := \Delta :: \Gamma$$

and then a positive or negative label would be a pair (n, i) , where the integer n identifies the n^{th} element Δ of the list and i is a string of 0 and 1 representing the path from the root of Δ (seen as a tree) to one of its leaves.

But we can also use De Bruijn’s levels, rather than indices.

One of the drawbacks of the implementation with De Bruijn’s indices, is that the name of a label, declared with a type in a typing context Γ , “changes” when Γ is extended with some typing decomposition Δ . For instance if $\Gamma[0^+]$ is an atom a because a is the head of the list Γ^+ , then in $\Gamma; \Delta$, a may no longer be the head of $(\Gamma; \Delta)^+$ and it will be referred to with an updated label name.

Depending on how the computations of $\Gamma[x^+]$ and $\Gamma[x^-]$ are implemented (imagine we have a HashMap for this), it could be problematic to have to update all the label names at every extension. We could do this update lazily, or we could also go for De Bruijn’s levels: once it has been introduced in a typing context, a label will remain unchanged by the subsequent extensions of the context.

DEFINITION 70 (Context algebras with De Bruijn’s levels - classical)

Positive labels and negative labels are two disjoint copies of the set of integers, with elements denoted n^+ and n^- , and we define

$(\Gamma^+, \Gamma^-)[n^+]$ as the $(|\Gamma^+| - n)^{th}$ element of Γ^+

$(\Gamma^+, \Gamma^-)[n^-]$ as the $(|\Gamma^-| - n)^{th}$ element of Γ^- .

※

In other words, the difference between De Bruijn’s indices and De Bruijn’s levels is that we are counting from the bottom of the list rather than from the head.

All of the above alternatives work for LAF_{K1} and LAF_{K2} , in that the assumptions of Definition 61 are clearly satisfied.

Choosing between them is really a question of implementation, with no theoretical impact.

4.4.2 Labels for intuitionistic logic

In the instance LAF_J , we have made some different assumptions on the typing context algebra (in Definition 68). We adapt our definition of the typing context algebra accordingly.

This time, we directly define it rather than go through the generic definition of an $(\mathcal{A}, \mathcal{B})$ -context algebra for each $\mathcal{A}$ and $\mathcal{B}$, since the assumptions in Definition 68 (unlike those in Definition 61) make a case analysis on the kind of atom (resp. on the kind of molecule) that is added to the typing context. That case analysis would not make sense for arbitrary sets $\mathcal{A}$ and $\mathcal{B}$.

DEFINITION 71 (Context algebras with De Bruijn’s indices - intuitionistic)

The typing context algebra Co is defined as follows:

A typing context Γ is a triple (Γ^+, Γ^-, R) where Γ^+ is a list of atoms, Γ^- is a list of molecules, and R is either an atom of the form (v, l) or a molecule of the form (P, r) .⁴

⁴Intuitively, R represents the right-hand side of the LJF sequent.

Extensions are defined as follows:

$$\begin{array}{llll}
(\Gamma^+, \Gamma^-, R); (l^+, r) & := & ((l^+, r) :: \Gamma^+, \Gamma^-, R) & (\Gamma^+, \Gamma^-, R); \sim(N, l) & := & (\Gamma^+, (N, l) :: \Gamma^-, R) \\
(\Gamma^+, \Gamma^-, R); (v, l) & := & (\Gamma^+, \Gamma^-, (v, l)) & (\Gamma^+, \Gamma^-, R); \sim(P, r) & := & (\Gamma^+, \Gamma^-, (P, r)) \\
(\Gamma^+, \Gamma^-, R); \bullet & := & (\Gamma^+, \Gamma^-) & (\Gamma^+, \Gamma^-, R); (\Delta_1, \Delta_2) & := & (\Gamma^+, \Gamma^-, R); \Delta_1; \Delta_2
\end{array}$$

Again, we use for labels two disjoint copies $\mathbb{N}^+$ and $\mathbb{N}^-$ of the set of integers:

A positive label is either some $n^+ \in \mathbb{N}^+$ or one of two special labels $\star^+$ and $x_{(\perp^-, l)}^+$.

A negative label is either some $n^- \in \mathbb{N}^-$ or the special label $\star^-$.

And we define

$(\Gamma^+, \Gamma^-, R) [n^+]$ as the $(n+1)^{th}$ element of Γ^+

$(\Gamma^+, \Gamma^-, R) [\star^+]$ as R if it is of the form (v, l) (undefined if not)

$(\Gamma^+, \Gamma^-, R) [x_{(\perp^-, l)}^+]$ as $(\perp^-, l)$

$(\Gamma^+, \Gamma^-, R) [n^-]$ as the $(n+1)^{th}$ element of Γ^-

$(\Gamma^+, \Gamma^-, R) [\star^-]$ as R if it is of the form (P, r) (undefined if not). ※

Clearly, the above definition of the typing context algebra satisfies the assumptions in Definition 68.

And again, there are many alternatives for the above definition, including the use of De Bruijn's levels, etc. Choosing between them would again simply be a question of implementation.

In brief, this section (as well as other parts of this dissertation) shows that the theory is able to handle diverse implementations, instead of having the theory commit to a particular choice of formalisation, and then having an implementation depart from it. Here, we can directly see the OCaml modules and module signatures that we can or should implement.

Chapter 5

An abstract focussed sequent calculus - with quantifiers

Contents

5.1	Presentation of the system	108
5.1.1	Quantifying structure	108
5.1.2	Atoms and Molecules	108
5.1.3	Typing decompositions and typing contexts	109
5.1.4	Logical connectives	112
5.1.5	Definition of the system	112
5.2	Extending LAF_{K1} with quantifiers	112

In this chapter we extend the LAF sequent calculus to handle quantifiers.

First, we should notice that the calculus we presented in Chapter 4 can already “handle quantifiers”, in the way the ω -rule does [Hil31, Sch50]. Indeed, we can adapt and extend system LAF_{K1} with an extra rule for the decomposition relation such as

$$\frac{\Delta \Vdash p : \{r/x\} A}{\Delta \Vdash (r, p) : \exists x A}$$

capturing the positive behaviour of the existential quantifier in the synchronous rule.

But this will also determine the asynchronous treatment of the universal quantifier: Ignoring proof-terms for the moment, proving the refutation $\Gamma \vdash \sim \exists x N$ (i.e. intuitively proving $\forall x N^\perp$) requires the use of rule `asyn`, with sub-proofs for each of the sequents

$$\Gamma, \sim \{t/x\} N^\perp \vdash$$

where t ranges over all potential witnesses for x , which is the behaviour of the ω -rule.

In particular if N is of the form $\forall y P$, each of those premisses can then be derived by a proof of the form

$$\frac{\frac{\Gamma \vdash \sim \{t, t'/x, y\} P}{\Gamma \vdash [\exists y \{t/x\} P^\perp]}}{\Gamma, \sim \exists y \{t/x\} P^\perp \vdash}$$

where t' is witness for y whose choice may depend (possibly in a non-uniform way) on the instance t of x .

So, in order to recover a standard rule for $\forall$ -introduction, which uses something like an eigenvariable, we need to enrich LAF, which will now be given by a tuple of parameters

$$(\mathbb{S}, \mathbb{T}, \mathbb{C}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \text{Co}, \text{Pat}, \Vdash)$$

where each parameter is described in Section 5.1.

Section 5.2 then provides an instance illustrating first-order quantification.

5.1 Presentation of the system

5.1.1 Quantifying structure

The first group of parameters $(\mathbb{S}, \mathbb{T}, \mathbb{C}, \Vdash)$ specifies the objects that LAF quantifies over. For logics with quantifiers, the following definition provides a rather general setting: the terms that can be provided as witnesses are multi-sorted, and the sorting may depend on a local sorting context (as we would need for higher-order logic, dependent types, etc).

DEFINITION 72 (Quantifying structure)

LAF is parameterised by a *quantifying structure* $(\mathbb{S}, \mathbb{T}, \mathbb{C}, \Vdash)$, made of

- a fixed set $\mathbb{S}$ of elements called *sorts*, denoted s, s_1 , etc.
- a fixed set $\mathbb{T}$ of elements called *terms*, denoted r, r_1 , etc,
- a fixed set $\mathbb{C}$ of elements called *sorting contexts*, denoted Σ, Σ_1 , etc,
- a *sorting relation*, i.e. a set of elements $(_ \Vdash _ : _) : (\mathbb{C} \times \mathbb{T} \times \mathbb{S})$

✖

5.1.2 Atoms and Molecules

The next group of parameters $(\mathbb{A}, \mathbb{M}, \equiv)$ adapts the notions of atoms and molecules to the presence of quantifiers.

Atoms and molecules now need more structure than in the propositional case, because intuitively, we would like atoms and molecules to refer to terms. Whether it is in the decomposition relation or elsewhere in the LAF inference system, we will have a rule where witnesses for existential variables are picked. This usually involves substituting the witness for the existential variable in the premiss of the rule.

Two reasons suggest to go for a different approach:

First, this requires us to specify how substitutions affect atoms and molecules; which then requires us to specify what variables and terms are for the abstract notions of atoms and molecules; then we would probably need to specify how substitutions affect the decomposition relation. All of which are rather heavy in our abstract setting.

Second, an implementation of proof-search would probably depart from such a rule anyway, as it could be costly to traverse the whole sequent, or even just some of its atoms and molecules, to compute the substitution every time a witness is picked. The substitution would more efficiently be done lazily, keeping the fact that the existential variable “is in fact the witness” in a separate data-structure to be looked up when we finally need the information.

Hence, we will formalise what would be actually closer to implementation, expressing an atom (it will be the same for a molecule) as a pair $(a, \mathbf{r})$ where a is a structure not (explicitly) referring to terms and $\mathbf{r}$ is a list of terms: the former is a *parameterised atoms* and the latter is its list of parameters. The list of parameters $\mathbf{r}$ can be seen as a delayed substitution, in the view that a refers to its parameters by either using something like De Bruijn’s indices, or by having a series of λ -abstractions at its top-level.

For instance, to represent the atoms of first-order logic, we could use a pair

$$(p(4, \#1, \#2), x::5::[])$$

(where p is a ternary predicate symbol and x is an eigenvariable) to represent the atom usually written as $p(4, x, 5)$.

A parameterised atom (such as $p(4, \#1, \#2)$) comes with a notion of *arity*: a list of sorts l describing the sorts of its parameters numbered from 1 to $|l|$ (the arity of $p(4, \#1, \#2)$ would here be a list of length at least 2).

This leads to the following definition.

DEFINITION 73 (Atoms & molecules)

LAF is parameterised by two sets $\mathbb{A}$ and $\mathbb{M}$, whose elements are respectively called (*parameterised atoms*) (denoted $a, a', \dots$) and (*parameterised molecules*) (denoted $M, M', \dots$), each of which is equipped with a function that maps every atom a (resp. molecule M) to a list of sorts denoted $|a|$ (resp. $|M|$) and called its *arity*.

The set $\mathbb{A}_\downarrow$ (resp. $\mathbb{M}_\downarrow$) of *instantiated atoms* (resp. *instantiated molecules*) is the set of pairs of the form $(a, \mathbf{r})$ (resp. $(M, \mathbf{r})$), where a is an atom (resp. M is a molecule) of arity l and $\mathbf{r}$ is a list of terms of length $|l|$.¹

LAF is also parameterised by an equivalence relation $\equiv$ over $\mathbb{A}_\downarrow$, which we call *equality*. $\ast$

The equality relation is what replaces the computation of substitutions: using the previous example, if $(p(4, \#1, \#2), x::5::[])$ “represents” the atom usually written as $p(4, x, 5)$, so do the pairs $(p(4, \#2, \#1), 5::x::[])$, $(p(4, \#1), x::[])$ and $(p(\#3, \#1, \#2), x::5::4::[])$. The equality relation on instantiated atoms is then used to declare all of these pairs be equal.

Interestingly enough, this equality relation will only be used at the leaves of proof-trees when proof-search has to compare two instantiated atoms to close the branch. More surprisingly, there is no need to have a similar equality relation between molecules; they are never compared during proof-search.

5.1.3 Typing decompositions and typing contexts

As we have seen, if the choice of witnesses for existential variables is made in the decomposition relation, the asynchronous phase treats universal variables in the style of the ω -rule. To avoid this, the choice of witnesses cannot be made in the decomposition relation; instead, we “leave a hole” and delay its filling until we inhabit typing decompositions.

Therefore, the notion of typing decomposition itself needs to be enriched with a new construct, denoted $s.\Delta$, that we use to mark a place where an existential variable of sort s was found while decomposing a molecule: For instance we can use the construct in the decomposition relation with a rule (again, forgetting proof-terms) such as

¹We do not relate the sorts specified in l to the sorts of the terms, which only make sense in a specific sorting context.

$$\frac{\Delta \Vdash \left\{ \#1 \middle/ x \right\} A}{s.\Delta \Vdash \exists x^s A}$$

where $\#1$ is a temporary name for the hole (you may think of it as a De Bruijn's index), or with the equivalent rule

$$\frac{\Delta \Vdash A}{s.\Delta \Vdash \exists^s A}$$

if De Bruijn's indices are already used in formulae to represent bound variables.

The choice of witness will then be made when proving/inhabiting $s.\Delta$.

Correspondingly, we extend the notion of typing decompositions as follows:

DEFINITION 74 (Typing decompositions)

The *typing decomposition algebra*, denoted $\mathbb{D}$, is the family of sets $(\mathbb{D}_l)_l$ defined by the following grammar:

$$\Delta^l, \Delta_1^l, \dots ::= a^l \mid \sim M^l \mid \bullet \mid \Delta_1^l, \Delta_2^l \mid s.\Delta^{s::l}$$

where $\Delta^l, \Delta_1^l, \dots$ range over $\mathbb{D}_l$, a^l ranges over parameterised atoms of arity l and M^l ranges over parameterised molecules of arity l (and s still ranges over $\mathbb{S}$).

Elements of $\mathbb{D}_l$ are called *typing decompositions* of arity l .

The set $\mathbb{D}_\downarrow$ of *instantiated typing decompositions* is the set of pairs of the form $(\Delta^l, \mathbf{r})$ where Δ^l is a typing decomposition of arity l and $\mathbf{r}$ is a list of terms of length $|l|$. *

Here, the construct $s.\Delta^{s::l}$ declares a new “hole” of sort s so that the atoms and molecules in $\Delta^{s::l}$ may now depend on this extra parameter.

But notice that typing decompositions (unless instantiated) never mention terms; they will be used to decompose *parameterised* molecules, rather than *instantiated* molecules: this is because, intuitively, the decomposition of a molecule into a typing decomposition only concerns the logical structure of the molecule, not the terms that it contains.²

Now, as in the quantifier-free case, typing decompositions will be used to extend typing contexts, but we do expect the types, in the typing declarations of a typing context, to be *instantiated* atoms and molecules.

This raises two questions when extending a typing context Γ with a typing decomposition Δ^l :

- how do the parameterised atoms and molecules of Δ^l turn into instantiated atoms and molecules in the extended environment?
- how should the new construct $s.\Delta^{s::l}$ impact the extension?

To answer these questions, we anticipate that, as in the quantifier-free case, the typing context Γ gets extended in the *async* rule. In our setting with quantifiers, that rule will be used to refute an *instantiated* molecule. This extension thus takes place in presence of the list of parameters $\mathbf{r}$ of the molecule we are refuting, and the length of this list will match the arity l of Δ^l .

Therefore, the operation that we need to extend environments is of the form $\Gamma; (\Delta^l, \mathbf{r})$, hence the notion of instantiated typing decomposition.

²This requires the distinction between the two to be clear, which will prevent us from modelling higher-order logic.

DEFINITION 75 (Typing contexts)

LAF is parameterised by two sets Lab_+ and Lab_- , of elements called *positive labels* and *negative labels*, respectively.

LAF is then parameterised by an algebra Co of the form

$$\left(\text{Co}, \left(\begin{array}{c} \text{Co} \times \text{Lab}_+ \rightarrow \mathbb{A}_\downarrow \\ (\Gamma, x^+) \mapsto \Gamma[x^+] \end{array} \right), \left(\begin{array}{c} \text{Co} \times \text{Lab}_- \rightarrow \mathbb{M}_\downarrow \\ (\Gamma, x^-) \mapsto \Gamma[x^-] \end{array} \right), \left(\begin{array}{c} \text{Co} \rightarrow \mathbb{C} \\ \Gamma \mapsto \Gamma^e \end{array} \right), \left(\begin{array}{c} \text{Co} \times \mathbb{D}_\downarrow \rightarrow \text{Co} \\ (\Gamma, (\Delta^l, \mathbf{r})) \mapsto \Gamma; (\Delta^l, \mathbf{r}) \end{array} \right) \right)$$

whose elements are called *typing contexts*.

We denote by $\text{dom}^+(\Gamma)$ (resp. $\text{dom}^-(\Gamma)$) the subset of Lab_+ (resp. Lab_-) where $\Gamma[x^+]$ (resp. $\Gamma[x^-]$) is defined. ※

Of course, nothing in the above definition specifies the behaviour of the operation $\Gamma; (\Delta^l, \mathbf{r})$.

This will not be problematic to define the LAF system, nor to define its realisability models; but in order to *build* those, it will be easier if we also know that the typing context algebra satisfies more specific properties. In Section 5.2 we give an example of LAF instance where the behaviour of $\Gamma; (\Delta^l, \mathbf{r})$ is specified.

Also, note the presence of the operation Γ^e that extracts from Γ a sorting context, which will be used in the LAF system to constrain the pick of witnesses.

Notice in the above two definitions (74 and 75) that, in contrast to what we did in the quantifier-free case, we have here directly defined typing decompositions and typing contexts instead of defining them as particular instances of generic decompositions and generic contexts. This is due to the need of taking into account, in those data structures, the parameters $\mathbf{r}$ that are specific to atoms and molecules and non-existent for arbitrary sets $\mathcal{A}$ and $\mathcal{B}$. However, we shall still define generic decompositions and generic contexts, as these will be used for instance to build models of LAF, and also more immediately to define the *structure* of a typing derivation (as we did in the quantifier-free case).

DEFINITION 76 (Generic decomposition algebras and decomposition structures)

Given three set $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{C}$, the $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -*decomposition algebra* $\mathbb{D}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$, whose elements are called $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -*decompositions*, is the free algebra defined by the following grammar:

$$\Delta, \Delta_1, \dots ::= a \mid \sim b \mid \bullet \mid \Delta_1, \Delta_2 \mid c. \Delta$$

where a (resp. b , c) ranges over $\mathcal{A}$ (resp. $\mathcal{B}$, $\mathcal{C}$).

Let $\mathbb{D}_{\text{st}}$ abbreviate $\mathbb{D}_{\text{unit}, \text{unit}, \text{unit}}$, whose elements we call *decomposition structures*.

The (*decomposition*) *structure* of an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -decomposition Δ , denoted $|\Delta|$, is its obvious homomorphic projection in $\mathbb{D}_{\text{st}}$.

The (*decomposition*) *structure* of a typing decomposition Δ^l , denoted $|\Delta^l|$, is defined as follows:

$$\begin{array}{ll} |a^l| & := () \\ |\bullet| & := () \\ |s. \Delta^{s::l}| & := (). |\Delta^{s::l}| \end{array} \quad \begin{array}{ll} |\sim M^l| & := () \\ |\Delta_1^l, \Delta_2^l| & := |\Delta_1^l|, |\Delta_2^l| \end{array}$$

※

Here, we see that the typing decomposition algebra is more subtle than the $(\mathbb{A}, \mathbb{M}, \mathbb{S})$ -decomposition algebra, because arities are taken into account.

Similarly, we here define generic contexts, which will be used in the next chapters.

DEFINITION 77 (Generic contexts)

Given four sets $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$ and $\mathcal{D}$, an $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D})$ -context algebra is an algebra of the form

$$\left(\mathcal{G}, \left(\begin{array}{c} \mathcal{G} \times \text{Lab}_+ \rightarrow \mathcal{A} \\ (\Gamma, x^+) \mapsto \Gamma[x^+] \end{array} \right), \left(\begin{array}{c} \mathcal{G} \times \text{Lab}_- \rightarrow \mathcal{B} \\ (\Gamma, x^-) \mapsto \Gamma[x^-] \end{array} \right), \left(\begin{array}{c} \mathcal{G} \rightarrow \mathcal{D} \\ \Gamma \mapsto \Gamma^e \end{array} \right), \left(\begin{array}{c} \mathcal{G} \times \mathbb{D}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \rightarrow \mathcal{G} \\ (\Gamma, \Delta) \mapsto \Gamma; \Delta \end{array} \right) \right)$$

whose elements are called $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D})$ -contexts.

Again, we denote by $\text{dom}^+(\Gamma)$ (resp. $\text{dom}^-(\Gamma)$) the subset of Lab_+ (resp. Lab_-) where $\Gamma[x^+]$ (resp. $\Gamma[x^-]$) is defined. *

5.1.4 Logical connectives

The concepts of patterns and decomposition relations are unchanged, except they rely on the enriched concepts of atoms, molecules and typing decompositions.

DEFINITION 78 (Patterns & decomposition relation)

LAF is parameterised by a *pattern algebra*, an algebra of the form

$$\left(\text{Pat}, \left(\begin{array}{c} \text{Pat} \rightarrow \mathbb{D}_{\text{st}} \\ p \mapsto |p| \end{array} \right) \right)$$

whose elements are called *patterns*, and by a *decomposition relation* (for every l), i.e. a set of elements

$$(_ \Vdash _ : _) : (\mathbb{D}_l \times \text{Pat} \times \mathbb{M}_l)$$

such that if $\Delta \Vdash p : M$ then the structure of Δ is $|p|$. *

5.1.5 Definition of the system

DEFINITION 79 (Proof-Terms) Proof-terms are defined by the following grammar:

Positive terms	Terms^+	$t^+ ::= pd$
Decomposition terms	Terms^d	$d ::= x^+ \mid f \mid \bullet \mid d_1, d_2 \mid r.d$
Commands	Terms	$c ::= \langle x^- \mid t^+ \rangle \mid \langle f \mid t^+ \rangle$

where p ranges over Pat , x^+ ranges over Lab_+ , x^- ranges over Lab_- , and f ranges over $\text{Pat} \rightarrow \text{Terms}$. *

We can finally present the typing system LAF:

DEFINITION 80 (LAF)

LAF is the inference system of Fig. 27 defining the derivability of three kinds of sequents

$$\begin{array}{ll} (_ \vdash _ : _) & : (\text{Co} \times \text{Terms}^+ \times \mathbb{M}_\downarrow) \\ (_ \vdash _ : _) & : (\text{Co} \times \text{Terms}^d \times \mathbb{D}_\downarrow) \\ (_ \vdash _) & : (\text{Co} \times \text{Terms}) \end{array}$$

We further impose in rule **async** that the domain of function f be exactly those patterns that can decompose M (if $p \in \text{Dom}(f)$ then there exists Δ such that $\Delta \Vdash p : M$).

LAF^{cf} is the inference system LAF without the cut-rule. *

5.2 Extending LAF_{K1} with quantifiers

In this section we extend to multi-sorted first-order logic the example of polarised classical logic (one-sided version LAF_{K1}) in Section 4.3. Such a first-order extension could also be done

$$\frac{\Delta \Vdash p:M \quad \Gamma \vdash d:(\Delta, \mathbf{r})}{\Gamma \vdash [pd:(M, \mathbf{r})]} \text{sync}$$

$$\frac{}{\Gamma \vdash \bullet:(\bullet, \mathbf{r})} \quad \frac{\Gamma \vdash d_1:(\Delta_1, \mathbf{r}) \quad \Gamma \vdash d_2:(\Delta_2, \mathbf{r})}{\Gamma \vdash d_1, d_2:((\Delta_1, \Delta_2), \mathbf{r})} \quad \frac{\Gamma^e \Vdash r':s \quad \Gamma \vdash d:(\Delta, r':\mathbf{r})}{\Gamma \vdash r'.d:s.(\Delta, \mathbf{r})}$$

$$\frac{\Gamma [x^+] \equiv (a, \mathbf{r})}{\Gamma \vdash x^+:(a, \mathbf{r})} \text{init} \quad \frac{\forall p, \forall \Delta, \quad \Delta \Vdash p:M \Rightarrow \Gamma; (\Delta, \mathbf{r}) \vdash f(p)}{\Gamma \vdash f:(\sim M, \mathbf{r})} \text{async}$$

$$\frac{\Gamma \vdash [t^+:\Gamma [x^-]]}{\Gamma \vdash \langle x^- \mid t^+ \rangle} \text{Select} \quad \frac{\Gamma \vdash f:(\sim M, \mathbf{r}) \quad \Gamma \vdash [t^+:(M, \mathbf{r})]}{\Gamma \vdash \langle f \mid t^+ \rangle} \text{cut}$$

Figure 27: LAF

for the two-sided versions of polarised classical logic or intuitionistic logic.

To handle quantifiers, we make a clear separation between *bound variables* and *eigenvariables*: the intuition being that in order to prove $\forall x p(x)$ we prove $p(\mathbf{x})$ for “an arbitrary $\mathbf{x}$ ”, using an eigenvariable $\mathbf{x}$.

Actually, the reasons why we used “labels” instead of “variables” in the quantifier-free system also apply to eigenvariables: both in the perspective of an implementation and for the formalisation of such an abstract system as LAF, it will be convenient to have a deterministic way to name eigenvariables with no notion of α -conversion or equivariance. We will therefore call them *eigenlabels*.

LAF is more flexible regarding bound variables, which could be named and subject to α -conversion. However, already using De Bruijn’s indices to represent binding in the syntax of formulae will be convenient since, as already mentioned in Section 5.1.3, we can simply write

$$\frac{\Delta \Vdash A}{s.\Delta \Vdash \exists^s A}$$

instead of

$$\frac{\Delta \Vdash \left\{ \#^1 / x \right\} A}{s.\Delta \Vdash \exists x^s A}$$

saving us the trouble of defining the substitution operation on formulae.

DEFINITION 81 (Literals, formulae, patterns, decomposition)

Let $\mathbb{S}$ be a set of *sorts* and Ξ be an $\mathbb{S}$ -signature in the sense of multi-sorted first-order logic. Predicate arities are represented as lists of sorts, denoted $l, l', \dots$

Given such an arity l , the set of l -literals is the set of literals over Ξ (well-sorted atomic propositions and their negations) whose free variables are among $\#1, \dots, \#|l|$ with (respective) sorts given by l .

Consider a subset of the set of predicate symbols, whose elements are called *positive predicate symbols*; predicate symbols that are not in that set are called *negative*.

Let the set $\mathbb{A}_l$ of *parameterised atoms of arity l* be the set of l -literals that are either of the form $p(t_1, \dots, t_n)$, with p being a positive predicate symbol, or of the form $\neg p(t_1, \dots, t_n)$, with p being a negative predicate symbol.

Similarly to Definition 60, let the set $\mathbb{M}_l$ of *parameterised molecules of arity l* be the set defined by the first line of the following grammar for (polarised) formulae of classical logic:

$$\begin{aligned} \text{Positive } l\text{-formulae} \quad P^l, \dots &::= a^l \mid \top^+ \mid \perp^+ \mid A^l \wedge^+ B^l \mid A^l \vee^+ B^l \mid \exists^s A^s ::^l \\ \text{Negative } l\text{-formulae} \quad N^l, \dots &::= a^{l^\perp} \mid \top^- \mid \perp^- \mid A^l \wedge^- B^l \mid A^l \vee^- B^l \mid \forall^s A^s ::^l \\ \text{Unspecified } l\text{-formulae} \quad A^l &::= P^l \mid N^l \end{aligned}$$

with a^l ranging over $\mathbb{A}_l$ and $a^{l^\perp}$ ranging over l -literals that are not in $\mathbb{A}_l$.

Similarly to Definition 60, let negation be the involutive function defined as follows:

$$\begin{aligned} (p(t_1, \dots, t_n))^\perp &:= \neg p(t_1, \dots, t_n) \\ (\neg p(t_1, \dots, t_n))^\perp &:= p(t_1, \dots, t_n) \\ \top^{+\perp} &:= \perp^- & \top^{-\perp} &:= \perp^+ \\ \perp^{+\perp} &:= \top^- & \perp^{-\perp} &:= \top^+ \\ (A \wedge^+ B)^\perp &:= A^\perp \vee^- B^\perp & (A \wedge^- B)^\perp &:= A^\perp \vee^+ B^\perp \\ (A \vee^+ B)^\perp &:= A^\perp \wedge^- B^\perp & (A \vee^- B)^\perp &:= A^\perp \wedge^+ B^\perp \end{aligned}$$

and we extend it to sets or multisets of formulae pointwise.

The set Pat of *patterns* extends that of Definition 60 according to the following grammar:

$$p, p_1, p_2, \dots ::= _+ \mid _- \mid \bullet \mid (p_1, p_2) \mid \text{inj}_i(p) \mid \exists p$$

The decomposition relation $(_ \Vdash _ : _) : (\mathbb{D} \times \text{Pat} \times \mathbb{M})$ is the extension of that of Definition 60, as shown in Fig. 28. *

$$\begin{array}{c} \bullet \Vdash \bullet : \top^+ \quad \sim N^\perp \Vdash _ : N \quad a \Vdash _ : a \\[10pt] \frac{\Delta_1 \Vdash p_1 : A_1 \quad \Delta_2 \Vdash p_2 : A_2}{\Delta_1, \Delta_2 \Vdash (p_1, p_2) : A_1 \wedge^+ A_2} \quad \frac{\Delta \Vdash p : A_i}{\Delta \Vdash \text{inj}_i(p) : A_1 \vee^+ A_2} \quad \frac{\Delta \Vdash p : A}{s.\Delta \Vdash \exists p : \exists^s A} \end{array}$$

Figure 28: Decomposition relation for LAF_{K1}

Several concepts are still missing to define an instance of LAF : we need to define the set $\mathbb{T}$ of terms, the set $\mathbb{C}$ of sorting contexts, the sorting relation $\Vdash$, the equality relation $\equiv$ on instantiated atoms, and the typing context algebra Co .

DEFINITION 82 (Terms, sorting and equality)

Let Lab_e be a copy of the set of natural numbers, whose elements are called *eigenlabels* and denoted $n^e, n_1^e, \dots$

The set $\mathbb{T}$ of *terms*, denoted $r, r', \dots$, is defined as the set of first-order terms whose variables are eigenlabels and whose function symbols are those declared in the signature Ξ .

The set $\mathbb{C}$ of *sorting contexts*, denoted $\Sigma, \Sigma', \dots$, is $\mathbf{Lab}_e \rightarrow \mathbb{S}$.

We write $\Sigma \Vdash r:s$ when the term r is of sort s in the sorting context Σ , according to the signature Ξ .

We define the equality relation as follows: $(a^l, \mathbf{r}) \equiv (a^{l'}, \mathbf{r}')$ if the literal $\{ \mathbf{r} / \#1, \dots, \#|l| \} a^l$ ³ is syntactically equal to the literal $\{ \mathbf{r}' / \#1, \dots, \#|l'| \} a^{l'}$.⁴ ※

The last task is to define the typing context algebra $\mathbf{Co}$. We do this by adapting Definitions 69 and 70. We will use De Bruijn's levels for eigenlabels, because as explained in Section 4.4, De Bruijn's levels do not need to be updated once they are introduced (in contrast to De Bruijn's indices).

DEFINITION 83 (Typing context algebra)

We define the support set of $\mathbf{Co}$ as the set of triples $(\Gamma^+, \Gamma^-, \Gamma^e)$ where Γ^+ is a list of elements of $\mathbb{A}_\downarrow$, Γ^- is a list of elements of $\mathbb{M}_\downarrow$, and Γ^e is a list of elements of $\mathbb{T}$.

As in Definition 69, two disjoint copies $\mathbf{Lab}_+$ and $\mathbf{Lab}_-$ of the set of natural numbers are used for positive labels and negative labels, respectively denoted n^+ and n^- , and we define $(\Gamma^+, \Gamma^-, \Gamma^e) [n^+]$ as the $(|\Gamma^+| - n)^{th}$ element of Γ^+ and $(\Gamma^+, \Gamma^-, \Gamma^e) [n^-]$ as the $(|\Gamma^-| - n)^{th}$ element of Γ^- .

We now also define $(\Gamma^+, \Gamma^-, \Gamma^e) [n^e]$ as the $(|\Gamma^e| - n)^{th}$ element of Γ^e , for an eigenlabel n^e .

We turn the resulting structure

$$\left(\mathbf{Co}, \left(\begin{array}{c} \mathbf{Co} \times \mathbf{Lab}_+ \rightarrow \mathbb{A}_\downarrow \\ (\Gamma, n^+) \mapsto \Gamma [n^+] \end{array} \right), \left(\begin{array}{c} \mathbf{Co} \times \mathbf{Lab}_- \rightarrow \mathbb{M}_\downarrow \\ (\Gamma, n^-) \mapsto \Gamma [n^-] \end{array} \right), \left(\begin{array}{c} \mathbf{Co} \rightarrow (\mathbf{Lab}_e \rightarrow \mathbb{T}) \\ \Gamma \mapsto (n^e \mapsto \Gamma [n^e]) \end{array} \right) \right)$$

into a typing context algebra, by adding a notion of typing context extension

$$\left(\begin{array}{c} \mathbf{Co} \times \mathbb{D}_\downarrow \rightarrow \mathbf{Co} \\ (\Gamma, (\Delta^l, \mathbf{r})) \mapsto \Gamma; (\Delta^l, \mathbf{r}) \end{array} \right)$$

which we define as follows:

$$\begin{aligned} (\Gamma^+, \Gamma^-, \Gamma^e); (a^l, \mathbf{r}) &:= ((a^l, \mathbf{r}) :: \Gamma^+, \Gamma^-, \Gamma^e) \\ (\Gamma^+, \Gamma^-, \Gamma^e); (\sim M^l, \mathbf{r}) &:= (\Gamma^+, (M^l, \mathbf{r}) :: \Gamma^-, \Gamma^e) \\ (\Gamma^+, \Gamma^-, \Gamma^e); \bullet &:= (\Gamma^+, \Gamma^-, \Gamma^e) \\ (\Gamma^+, \Gamma^-, \Gamma^e); ((\Delta_1^l, \Delta_2^l), \mathbf{r}) &:= (\Gamma^+, \Gamma^-, \Gamma^e); (\Delta_1^l, \mathbf{r}); (\Delta_2^l, \mathbf{r}) \\ (\Gamma^+, \Gamma^-, \Gamma^e); (s.\Delta^{s::l}, \mathbf{r}) &:= (\Gamma^+, \Gamma^-, s :: \Gamma^e); (\Delta^{s::l}, |\Gamma^e|^e :: \mathbf{r}) \end{aligned}$$

※

The operation of typing context extension adapts to the presence of quantifiers the operation defined in Definition 69 for the quantifier-free case.

The only difference is the third component Γ^e of the typing context, which records the declared eigenlabels together with their sorts. This sorting context is extended whenever the typing context is extended with an instantiated typing decomposition of the form $(s.\Delta^{s::l}, \mathbf{r})$, which creates a new eigenlabel of sort s , which becomes the new head of the sorting context. As we use De Bruijn's levels rather than De Bruijn's indices, the new eigenlabel is therefore $|\Gamma^e|^e$ (and it gets added to the current list of terms). This can be seen as picking the “first available name” for the eigenlabel to be created, a process that is often used in implementations of such systems.

³i.e. the substitution of $\mathbf{r}$ for $\#1, \dots, \#|l|$ in a^l

⁴i.e. the substitution of $\mathbf{r}'$ for $\#1, \dots, \#|l'|$ in $a^{l'}$

With De Bruijn’s indices rather than De Bruijn’s levels, the new eigenlabel would be 0^e , but the price to pay for this is that the previously declared eigenlabels have “changed names”, i.e. would be referred to as $(n+1)^e$ instead of n^e . Every structure referring to those previously declared eigenlabels (namely, the instantiated atoms and molecules in Γ^+ and Γ^- , as well as $\mathbf{r}$ itself) would then need to be updated with the name change.

We could easily define variants of the above system to quantify over other objects than first-order terms, as most of the definitions are rather modular: For example we could quantify over simply-typed λ -terms to design a LAF instance similar to the type theory $\lambda\Pi$ (see e.g. [Bar91]), except our proof-terms do not have the same syntax and typing properties as the λ -terms we quantify over.

For this we take the same definitions as for multi-sorted first-order logic, except in Definition 81 we take $\mathbb{S}$ be the set of *simple types*, and in Definition 82 we take terms to be λ -terms, we take $\Sigma \Vdash r:s$ to be the typing relation of the simply-typed λ -calculus, and we define atom equality with β - (or $\beta\eta$ -) conversion instead of syntactic equality: $(a^l, \mathbf{r}) \equiv (a^{l'}, \mathbf{r}')$ if $(\lambda\#1 \dots \#|l|.a^l) \mathbf{r} \longleftrightarrow_{\beta}^* (\lambda\#1 \dots \#|l'|.b^{l'}) \mathbf{r}'$ (or similarly with $\beta\eta$).

All of the other definitions are the same.

Chapter 6

Realisability models of abstract focussing

Contents

6.1	Model structures and the interpretation of proof-terms	118
6.2	Realisability algebras, interpretation of types & Adequacy	119
6.3	A more concrete class of LAF instances	121
6.3.1	LAF instances with eigenlabels	122
6.3.2	LAF _{K1} is a LAF instance with eigenlabels	124
6.3.3	LAF instances with eigenlabels are LAF instances	125
6.4	A more concrete class of realisability algebras	127
6.5	Example: boolean models to prove Consistency	129

In this chapter we investigate the semantics of LAF. More precisely, we investigate models of proofs / typing derivations with the *Adequacy Lemma* as the main objective: In very generic terms, if t is of type A then in the model we want the interpretation of t to be in the interpretation of A (if that is a set, or we want the interpretation of t to satisfy the interpretation of A , if that is a predicate).

Of course there are many models satisfying the above, starting with the uninformative ones where everything is collapsed.¹ So we investigate here a class of models, as large as possible, and prove the Adequacy Lemma generically for that class; then we will show interesting models in that class for which the Adequacy Lemma (that we now have for free) is informative (e.g. concludes the consistency of LAF, despite the presence of cuts and without proving cut-elimination).

This class of models is that of *abstract realisability algebras*; the specifications that we require of such an algebra do depend on the instance of LAF that we want to model - they will be different if we are to model for instance LAF_{K1}, LAF_{K2}, or LAF_J; but we can give those specifications parametrically and prove the Adequacy Lemma generically.

Hence in this chapter we start by assuming we are given an instance of LAF

¹Interpret every type by the same singleton set and every inhabitant of that type by the inhabitant of the singleton set, and the Adequacy Lemma trivially holds but does not provide any useful information.

$$(\mathbb{S}, \mathbb{T}, \mathbb{C}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \text{Co}, \text{Pat}, \Vdash)$$

Section 6.1 gives the specifications needed to interpret terms, Section 6.2 gives the specifications needed to interpret types and proves the Adequacy Lemma. Finally, Section 6.5 exhibits a simple model from which we derive the consistency of LAF.

6.1 Model structures and the interpretation of proof-terms

In this section we interpret the proof-terms of LAF in a realisability algebra, and for this we introduce the notion of *model structure*.

DEFINITION 84 (Model structure)

A *model structure* is an algebra of the form

$$\left(\mathcal{T}, \mathcal{C}, \mathcal{L}, \mathcal{P}, \mathcal{N}, \perp, \tilde{\text{Co}}, \left(\begin{array}{c} \text{Pat} \rightarrow (\mathbb{D}_{\mathcal{L}, \mathcal{N}, \mathcal{T}} \rightarrow \mathcal{P}) \\ p \mapsto \tilde{p} \end{array} \right), \left(\begin{array}{c} \mathbb{T} \times \mathcal{C} \rightarrow \mathcal{T} \\ (r, \sigma) \mapsto \llbracket r \rrbracket_\sigma \end{array} \right), \left(\begin{array}{c} (\text{Pat} \rightarrow \text{Terms}) \times \tilde{\text{Co}} \rightarrow \mathcal{N} \\ (f, \rho) \mapsto \llbracket f \rrbracket_\rho \end{array} \right) \right)$$

where

- $\mathcal{T}, \mathcal{C}, \mathcal{L}, \mathcal{P}, \mathcal{N}$ are five arbitrary sets of elements called *term denotations*, *valuations*, *label denotations*, *positive denotations*, *negative denotations*, respectively;
- $\perp$ is a relation between negative and positive denotations ($\perp \subseteq \mathcal{N} \times \mathcal{P}$), called the *orthogonality relation*;
- $\tilde{\text{Co}}$ is a $(\mathcal{L}, \mathcal{N}, \mathcal{T}, \mathcal{C})$ -context algebra, whose elements, denoted $\rho, \rho', \dots$, are called *semantic contexts*.

We extend the notation $\llbracket r \rrbracket_\sigma$ to apply to a list of terms $\mathbf{r}$: $\llbracket \mathbf{r} \rrbracket_\sigma$, using the standard map function on lists.

The $(\mathcal{L}, \mathcal{N}, \mathcal{T})$ -decomposition algebra $\mathbb{D}_{\mathcal{L}, \mathcal{N}, \mathcal{T}}$ is abbreviated $\tilde{\mathbb{D}}$; its elements, denoted $\mathfrak{d}, \mathfrak{d}', \dots$, are called *semantic decompositions*. ※

Given a model structure, we can define the interpretation of proof-terms. The model structure already gives an interpretation for the partial functions f from patterns to commands. We extend it to all proof-terms as follows

DEFINITION 85 (Interpretation of proof-terms)

Positive terms (in Terms^+) are interpreted as positive denotations (in $\mathcal{P}$), decomposition terms (in Terms^d) are interpreted as semantic decompositions (in $\tilde{\mathbb{D}}$), and commands (in Terms) are interpreted as pairs in $\mathcal{N} \times \mathcal{P}$ (that may or may not be orthogonal), according to the following definition:

$$\begin{array}{lll} \llbracket pd \rrbracket_\rho & := \tilde{p}(\llbracket d \rrbracket_\rho) & \llbracket x^+ \rrbracket_\rho & := \rho[x^+] & \llbracket \langle x^- \mid t^+ \rangle \rrbracket_\rho & := (\rho[x^-], \llbracket t^+ \rrbracket_\rho) \\ \llbracket f \rrbracket_\rho & & \llbracket f \rrbracket_\rho & := \llbracket f \rrbracket_\rho^2 & \llbracket \langle f \mid t^+ \rangle \rrbracket_\rho & := (\llbracket f \rrbracket_\rho, \llbracket t^+ \rrbracket_\rho) \\ \llbracket \bullet \rrbracket_\rho & & \llbracket \bullet \rrbracket_\rho & := \bullet & & \\ \llbracket d_1, d_2 \rrbracket_\rho & & \llbracket d_1, d_2 \rrbracket_\rho & := \llbracket d_1 \rrbracket_\rho, \llbracket d_2 \rrbracket_\rho & & \\ \llbracket r.d \rrbracket_\rho & & \llbracket r.d \rrbracket_\rho & := \llbracket r \rrbracket_{\rho^e} \cdot \llbracket d \rrbracket_\rho & & \end{array}$$

※

²as given by the model structure

6.2 Realisability algebras, interpretation of types & Adequacy

Again, let us keep in mind the Adequacy Lemma: if t is of type A then the interpretation of t to be in the interpretation of A . We have already defined the interpretation of proof-terms in a model structure. We now proceed to define the interpretation of types.

In system LAF, there are three concepts of “type inhabitation” for atoms and molecules:

- “proving” an atom by finding a suitable positive label in the typing context (the inhabitant is in Lab_+);
- “proving” a molecule by choosing a way to decompose it into a typing decomposition (the inhabitant is in Terms^+);
- “refuting” a molecule by case analysing all the possible ways of decomposing it into a typing decomposition (the inhabitant is in $\text{Pat} \rightarrow \text{Terms}$).

As positive labels are interpreted in $\mathcal{L}$, positive proof-terms are interpreted in $\mathcal{P}$ and functions in $\text{Pat} \rightarrow \text{Terms}$ are interpreted in $\mathcal{N}$, we will correspondingly

- have an interpretation of every atom as a particular subset of $\mathcal{L}$;
- have a *positive* interpretation of every molecule as a particular subset of $\mathcal{P}$;
- have a *negative* interpretation of every molecule as a particular subset of $\mathcal{N}$.

To make sure that we capture, in our notion of abstract realisability algebra, a wide class of models, the first of the three above interpretations will be left as a parameter; we barely impose any specification on this parameter. The other two, however, will be *defined* notions.

Also, we have in LAF a notion of *sorting* for terms, whose counter-part in an abstract realisability algebra is also left as a parameter to be fixed *ad libitum*. This leads to the following definition:

DEFINITION 86 (Realisability algebra)

A *realisability algebra* is

- a model structure
- together with three functions

$$\left(\begin{array}{c} \mathbb{S} \rightarrow \mathbb{P}(\mathcal{L}) \\ s \mapsto \llbracket s \rrbracket \end{array} \right), \quad \left(\begin{array}{c} \mathbb{C} \rightarrow \mathbb{P}(\mathcal{P}) \\ \Sigma \mapsto \llbracket \Sigma \rrbracket \end{array} \right), \quad \left(\begin{array}{c} \mathbb{A}_l \rightarrow (\mathcal{T}^{|l|} \rightarrow \mathbb{P}(\mathcal{L})) \\ a^l \mapsto \llbracket a^l \rrbracket \end{array} \right)$$

satisfying

- if $\Sigma \Vdash r:s$ and $\sigma \in \llbracket \Sigma \rrbracket$ then $(\llbracket r \rrbracket)_\sigma$ is defined and $\llbracket r \rrbracket_\sigma \in \llbracket s \rrbracket$;
- if $(a, \mathbf{r}) \equiv (a', \mathbf{r}')$ then for all $\sigma : \mathcal{C}$ we have $\llbracket a \rrbracket(\llbracket \mathbf{r} \rrbracket_\sigma) = \llbracket a' \rrbracket(\llbracket \mathbf{r}' \rrbracket_\sigma)$.³

✱

Now notice in LAF that the derivability of the typing judgements for atoms and molecules is defined inductively together with the derivability of a typing judgement for typing decompositions; inhabitants of those are decomposition terms.

Therefore, we will also define an interpretation for typing decompositions, as subsets of $\mathbb{D}_{\mathcal{L}, \mathcal{N}, \mathcal{T}}$. For this we need to specify how to “lift relations to typing decompositions”:

³if both sides are defined

DEFINITION 87 (Lifting relations) Given

- two relations $\mathcal{R}_1 \subseteq \mathbb{A}_l \times \mathcal{T}^{|l|} \times \mathcal{L}$ and $\mathcal{R}_2 \subseteq \mathbb{M}_l \times \mathcal{T}^{|l|} \times \mathcal{N}$ (for every arity l)
- a relation $\mathcal{R}_3 \subseteq \mathbb{S} \times \mathcal{T}$,
- an arity l and a list of term denotations $\mathbf{rl}$ of length $|l|$,
- a typing decomposition Δ^l of arity l and a semantic decomposition $\mathfrak{d}$

we say that Δ^l *rl-relates to* $\mathfrak{d}$ according to $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$ if the relation $(\Delta^l, \mathbf{rl}) \mathcal{R} \mathfrak{d}$ can be derived by the following rules:

$$\frac{(a^l, \mathbf{rl}) \mathcal{R}_1 \mathfrak{l}}{(a^l, \mathbf{rl}) \mathcal{R} \mathfrak{l}} \quad \frac{(M^l, \mathbf{rl}) \mathcal{R}_2 \mathfrak{n}}{(\sim M^l, \mathbf{rl}) \mathcal{R} \mathfrak{n}} \quad \frac{}{(\bullet, \mathbf{rl}) \mathcal{R} \bullet}$$

$$\frac{(\Delta_1^l, \mathbf{rl}) \mathcal{R} \mathfrak{d}_1 \quad (\Delta_2^l, \mathbf{rl}) \mathcal{R} \mathfrak{d}_2}{((\Delta_1^l, \Delta_2^l), \mathbf{rl}) \mathcal{R} (\mathfrak{d}_1, \mathfrak{d}_2)} \quad \frac{s \mathcal{R}_3 \mathfrak{r} \quad (\Delta^{s::l}, \mathbf{r}::\mathbf{rl}) \mathcal{R} \mathfrak{d}}{(s.\Delta^{s::l}, \mathbf{rl}) \mathcal{R} \mathfrak{r}.\Delta'}$$

※

Obviously in that case Δ^l and $\mathfrak{d}$ have the same decomposition structure.

The interpretation of types will be defined by simultaneous induction on molecules and typing decompositions. This induction needs to follow a well-founded relation:

DEFINITION 88 (Well-founded LAF instance)

We write $M' \leq M^l$ if there are Δ^l and p such that $\Delta^l \Vdash p:M^l$ and M' is a leaf of Δ^l .

The LAF instance is *well-founded* if $\leq$ is well-founded.

※

It could be the case that the LAF instance is not well-founded, e.g. if molecules contain fixpoints.

Notice 1 In the rest of this chapter, we will assume LAF instances to be well-founded.

Under this assumption, the following interpretations of types are well-defined:

DEFINITION 89 (Interpretation of types and typing contexts) A realisability algebra already provides the interpretation of a parameterised atom of arity l in $(\mathcal{T}^{|l|} \rightarrow \mathbb{P}(\mathcal{L}))$.

We now define

the positive interpretation of a parameterised molecule of arity l in $(\mathcal{T}^{|l|} \rightarrow \mathbb{P}(\mathcal{P}))$;

the negative interpretation of a parameterised molecule of arity l in $(\mathcal{T}^{|l|} \rightarrow \mathbb{P}(\mathcal{N}))$;

the interpretation of a typing decomposition of arity l in $(\mathcal{T}^{|l|} \rightarrow \mathbb{P}(\mathbb{D}_{\mathcal{L}, \mathcal{N}, \mathcal{T}}))$:

$$\begin{aligned} \llbracket M^l \rrbracket^+(\mathbf{rl}) &:= \{\tilde{p}(\mathfrak{d}) \in \mathcal{P} \mid \mathfrak{d} \in \llbracket \Delta^l \rrbracket(\mathbf{rl}), \text{ and } \Delta^l \Vdash p:M^l\} \\ \llbracket M^l \rrbracket^-(\mathbf{rl}) &:= \{\mathfrak{n} \in \mathcal{N} \mid \forall \mathfrak{p} \in \llbracket M^l \rrbracket^+(\mathbf{rl}), \mathfrak{n} \perp \mathfrak{p}\} \\ \llbracket \Delta^l \rrbracket(\mathbf{rl}) &:= \{\mathfrak{d} \in \mathbb{D} \mid \Delta^l \text{ rl-relates to } \mathfrak{d} \text{ according to } \{(a^l, \mathbf{rl}, \mathfrak{l}) \mid \mathfrak{l} \in \llbracket a^l \rrbracket(\mathbf{rl})\} \\ &\quad \{(M^l, \mathbf{rl}, \mathfrak{n}) \mid \mathfrak{n} \in \llbracket M^l \rrbracket^-(\mathbf{rl})\} \\ &\quad \text{and } \{(s, \mathfrak{r}) \mid \mathfrak{r} \in \llbracket s \rrbracket\}\} \end{aligned}$$

We now define the semantics of instantiated atoms, molecules and typing decompositions:

$$\begin{aligned} \llbracket (a^l, \mathbf{r}) \rrbracket_\sigma &:= \llbracket a^l \rrbracket(\llbracket \mathbf{r} \rrbracket_\sigma) & \llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^+ &:= \llbracket M^l \rrbracket^+(\llbracket \mathbf{r} \rrbracket_\sigma) \\ \llbracket (\Delta^l, \mathbf{r}) \rrbracket_\sigma &:= \llbracket \Delta^l \rrbracket(\llbracket \mathbf{r} \rrbracket_\sigma) & \llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^- &:= \llbracket M^l \rrbracket^-(\llbracket \mathbf{r} \rrbracket_\sigma) \end{aligned}$$

We finally define the interpretation of a typing context:⁴

$$\begin{aligned} \llbracket \Gamma \rrbracket := \{ \rho \in \tilde{\mathbf{Co}} \mid & \rho^e \in \llbracket \Gamma^e \rrbracket \\ & \forall x^+ \in \text{dom}^+(\rho), \rho[x^+] \in \llbracket \Gamma[x^+] \rrbracket_{\rho^e} \\ & \forall x^- \in \text{dom}^-(\rho), \rho[x^-] \in \llbracket \Gamma[x^-] \rrbracket_{\rho^e}^- \} \end{aligned}$$

※

Now that we have defined the interpretation of terms and the interpretation of types, we prove the Adequacy Lemma.

LEMMA 51 (Adequacy for LAF) We assume the following hypotheses:

Well-foundedness:

The LAF instance is well-founded.

Typing correlation:

If $\rho \in \llbracket \Gamma \rrbracket$ and $\mathfrak{d} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket_{\rho^e}$ then $(\rho; \mathfrak{d}) \in \llbracket \Gamma; (\Delta^l, \mathbf{r}) \rrbracket$.

Stability:

If $\mathfrak{d} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket_{\sigma}$ for some $\Delta^l, \sigma, \mathbf{r}$ and $\llbracket f(p) \rrbracket_{\rho; \mathfrak{d}} \in \perp$, then $\llbracket f \rrbracket_{\rho} \perp \tilde{p}(\mathfrak{d})$.

We conclude that, for all $\rho \in \llbracket \Gamma \rrbracket$,

1. if $\Gamma \vdash [t^+ : (M^l, \mathbf{r})]$ then $\llbracket t^+ \rrbracket_{\rho} \in \llbracket (M^l, \mathbf{r}) \rrbracket^+$;
2. if $\Gamma \vdash d : (\Delta^l, \mathbf{r})$ then $\llbracket d \rrbracket_{\rho} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket$;
3. if $\Gamma \vdash t$ then $\llbracket t \rrbracket_{\rho} \in \perp$.

※

Proof: See the proof in Coq [GL14].

□

6.3 A more concrete class of LAF instances

Looking at the Adequacy Lemma, the stability condition is traditional: it is the generalisation, to that level of abstraction, of the usual condition on the orthogonality relation in orthogonality models (those realisability models that are defined in terms of orthogonality, usually to model classical proofs [Gir87, DK00, Kri01, MM09, LM08]): orthogonality is “closed under anti-reduction”. Here, we have not yet defined a notion of reduction on LAF proof-terms, but intuitively, we would expect, in order to reduce cuts, to rewrite $\langle f \mid pd \rangle$ to $f(p)$ “substituted by d ”.

On the other hand, the typing correlation property is new, and is due to the level of abstraction we operate at: there is no reason why our data structure for typing contexts would relate to our data structure for semantic contexts, and the extension operation, in both of them, has so far been completely unspecified. Hence, we clearly need such an assumption to relate the two.

However, one may wonder when and why the typing correlation property should be satisfied. Looking at the example of $\mathbf{LAF}_{K1}$, one may anticipate how typing correlation could hold for this instance of LAF: at least in the quantifier-free case (Sections 4.3.1 and 4.4.1), we

⁴In this definition we implicitly require that $\text{dom}^+(\rho) = \text{dom}^+(\Gamma)$, $\text{dom}^-(\rho) = \text{dom}^-(\Gamma)$ and for all $x^+ \in \text{dom}^+(\rho)$ (resp. $x^- \in \text{dom}^-(\rho)$) $\llbracket \Gamma[x^+] \rrbracket_{\rho^e}$ (resp. $\llbracket \Gamma[x^-] \rrbracket_{\rho^e}^-$) is defined.

have a generic definition of $(\mathcal{A}, \mathcal{B})$ -contexts with a parametric operation of extension, which we can use for both typing contexts and semantic contexts.

In this section we generalise this example to a class of LAF systems (and later we identify a corresponding subclass of realisability algebras), where Adequacy holds under a hypothesis that simplifies (and entails) typing correlation, and that is satisfied in particular when the extension of contexts is defined parametrically.

6.3.1 LAF instances with eigenlabels

In this class, the extension of a typing context $\Gamma; (\Delta^l, \mathbf{r})$ is expressed in terms of the extension operation $\Gamma'; \Delta'$ of an $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S}, \mathbb{C})$ -context algebra.

This is done along the same lines as in the example of LAF_{K1} in Section 5.2, i.e. with a notion of eigenlabel and the understanding of sorting contexts (in $\mathbb{C}$) as functions mapping eigenlabels to sorts. Hence, we update and refine previous definitions (and introduce new ones) with this understanding of sorting contexts.

DEFINITION 90 (Three-parameter contexts)

Assume we have three disjoint sets Lab_+ , Lab_- and Lab_e , the union of which $(\text{Lab}_+ \cup \text{Lab}_- \cup \text{Lab}_e)$ we denote Lab .

Given three sets $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$, we abbreviate the terminology $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \text{Lab}_e \rightarrow \mathcal{C})$ -context into $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context.

We also abbreviate $\Gamma^e(x)$ as $\Gamma[x]$, for an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context Γ and an element $x \in \text{Lab}_e$, writing $\text{dom}^e(\Gamma)$ for $\text{Dom}(\Gamma^e)$. Finally, we abbreviate $\text{dom}^+(\Gamma) \cup \text{dom}^-(\Gamma) \cup \text{dom}^e(\Gamma)$ as $\text{dom}(\Gamma)$, and we say that Γ is *empty* if $\text{dom}(\Gamma) = \emptyset$. ※

We now introduce the lifting of relations to generic decompositions and contexts:

DEFINITION 91 (Lifting relations)

Assume we are given three relations $\mathcal{R}_1 \subseteq \mathcal{A} \times \mathcal{A}'$, $\mathcal{R}_2 \subseteq \mathcal{B} \times \mathcal{B}'$ and $\mathcal{R}_3 \subseteq \mathcal{C} \times \mathcal{C}'$.

We say that an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -decomposition Δ relates to an $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ -decomposition Δ' according to $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$ if the relation $\Delta \mathcal{R} \Delta'$ can be derived by the following rules:

$$\frac{a \mathcal{R}_1 a'}{a \mathcal{R} a'} \quad \frac{b \mathcal{R}_2 b'}{b \mathcal{R} b'} \quad \bullet \mathcal{R} \bullet \quad \frac{\Delta_1 \mathcal{R} \Delta'_1 \quad \Delta_2 \mathcal{R} \Delta'_2}{\Delta_1, \Delta_2 \mathcal{R} \Delta'_1, \Delta'_2} \quad \frac{c \mathcal{R}_3 c' \quad \Delta \mathcal{R} \Delta'}{c. \Delta \mathcal{R} c'. \Delta'}$$

We say that an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context Γ relates to an $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ -context Γ' according to $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$ if⁵

$$\begin{aligned} \forall x^+ \in \text{Lab}_+, \quad & \Gamma[x^+] \mathcal{R}_1 \Gamma'[x^+] \\ \forall x^- \in \text{Lab}_-, \quad & \Gamma[x^-] \mathcal{R}_2 \Gamma'[x^-] \\ \forall x \in \text{Lab}_e, \quad & \Gamma[x] \mathcal{R}_3 \Gamma'[x] \end{aligned}$$

Assume we are now given three functions $f_1 : \mathcal{A} \rightarrow \mathcal{A}'$, $f_2 : \mathcal{B} \rightarrow \mathcal{B}'$ and $f_3 : \mathcal{C} \rightarrow \mathcal{C}'$.

we say that Γ' is a map of Γ according to f_1 , f_2 and f_3 if it relates to Γ according to the relations $\{(f_1(a), a) \mid a \in \mathcal{A}\}$, $\{(f_2(b), b) \mid b \in \mathcal{B}\}$ and $\{(f_3(c), c) \mid c \in \mathcal{C}\}$. ※

Using the above two definition, we can now say what a LAF instance with eigenlabels is:

⁵By writing the three conditions we implicitly request $\text{dom}^+(\Gamma) = \text{dom}^+(\Gamma')$, $\text{dom}^-(\Gamma) = \text{dom}^-(\Gamma')$ and $\text{dom}^e(\Gamma) = \text{dom}^e(\Gamma')$.

DEFINITION 92 (LAF instance with eigenlabels)

A LAF *instance with eigenlabels* is given by the following tuple:

$$(\mathbb{S}, \text{Lab}_e, \mathbb{T}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \text{Co}, \text{Pat}, \Vdash, \pi_\Delta^\mathcal{V}, \text{st}_\Delta^\mathcal{V})$$

where

- $\mathbb{S}$ is as in Definition 72 (a set of elements called *sorts*, denoted s, s_1 , etc);
- Lab_e is a set of elements called *eigenlabels*, denoted x, x'_1 , etc;
- $\mathbb{T}$ is a set of elements called *terms*, denoted r, r_1 , etc,
 - that extends the set Lab_e of eigenlabels,
 - and with a systematic way of lifting a function $\text{Lab}_e \rightarrow \text{Lab}_e$ to $\mathbb{T} \rightarrow \mathbb{T}$;

We can then apply a function $\pi : \text{Lab}_e \rightarrow \text{Lab}_e$ to lists of terms (using the standard map function on lists);

- $\Vdash$ is a *sorting relation*, i.e. a set of elements $(_ \Vdash _ : _) : ((\text{Lab}_e \rightarrow \mathbb{S}) \times \mathbb{T} \times \mathbb{S})$, with
 - $\Sigma \Vdash x : s$ if and only if $s = \Sigma(x)$
 - for all $\pi : \text{Lab}_e \rightarrow \text{Lab}_e$, if $\Sigma \circ \pi \Vdash r : s$ then $\Sigma \Vdash \pi(r) : s$;
- $\mathbb{A}, \mathbb{M}, \text{Lab}_+$ and Lab_- are as in Definitions 73 and 75;

We can apply a function $\pi : \text{Lab}_e \rightarrow \text{Lab}_e$ to instantiated atoms and molecules using

$$\pi(a^l, \mathbf{r}) := (a^l, \pi(\mathbf{r})) \text{ and } \pi(M^l, \mathbf{r}) := (M^l, \pi(\mathbf{r}));$$

We then impose that equality on instantiated atoms be stable under any such function $\pi : \text{Lab}_e \rightarrow \text{Lab}_e$: If $(a, \mathbf{r}) \equiv (a', \mathbf{r}')$ then $\pi(a, \mathbf{r}) \equiv \pi(a', \mathbf{r}')$.

- Co is an $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S})$ -context algebra, called the *typing context algebra*, equipped with a *map* operation that associates, to a context Γ and two functions $f_1 : \mathbb{A}_\downarrow \rightarrow \mathbb{A}_\downarrow$, $f_2 : \mathbb{M}_\downarrow \rightarrow \mathbb{M}_\downarrow$, a context $(f_1, f_2) \circ \Gamma$ that is a map of Γ according to f_1, f_2 and the identity on sorts;
- Pat and $\Vdash$ are as in Definition 78.
- We have two functions, respectively called the *renaming policy* and the *fresh naming policy*, of the form

$$\left(\begin{array}{c} \mathbb{P}(\text{Lab}) \times \mathbb{D}_{\text{st} \rightarrow (\text{Lab}_e \rightarrow \text{Lab}_e)} \\ (\mathcal{V}, \Delta) \mapsto \pi_\Delta^\mathcal{V} \end{array} \right), \left(\begin{array}{c} \mathbb{P}(\text{Lab}) \times \mathbb{D}_{\text{st} \rightarrow \mathbb{D}_{\text{unit}, \text{unit}, \text{Lab}_e}} \\ (\mathcal{V}, \Delta) \mapsto \text{st}_\Delta^\mathcal{V} \end{array} \right)$$

We abbreviate $\text{st}_{|\Delta|}^{\text{dom}(\Gamma)}$ as st_Δ^Γ and $\pi_{|\Delta|}^{\text{dom}(\Gamma)}$ as π_Δ^Γ .

※

Most of the above definition is rather straightforward when thinking of sorting contexts as assigning sorts to eigenlabels. What is probably more cryptic is the naming policies $\pi_\Delta^\mathcal{V}$ and $\text{st}_\Delta^\mathcal{V}$ (as well as the map operation of typing contexts): they compensate for the fact that the extension operation of the typing context algebra Co is more basic than in Definition 78. In short, st_Δ^Γ and π_Δ^Γ describe which labels are used in an extended typing context $\Gamma; \Delta$ (especially regarding the labels used in Γ). Section 6.3.3 explains this in details, but we first start with the example of LAF_{K1} .

6.3.2 $\mathbf{LAF}_{K1}$ is a $\mathbf{LAF}$ instance with eigenlabels

In this section we illustrate the above concept by giving an alternative definition for system $\mathbf{LAF}_{K1}$ (from Section 5.2), this time as a $\mathbf{LAF}$ instance with eigenvariables.

Among the parameters

$$(\mathbb{S}, \mathbf{Lab}_e, \mathbb{T}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \mathbf{Lab}_+, \mathbf{Lab}_-, \mathbf{Co}, \mathbf{Pat}, \Vdash)$$

of a $\mathbf{LAF}$ instance with eigenvariables, the context algebra $\mathbf{Co}$ is perhaps the least obvious to identify for $\mathbf{LAF}_{K1}$. We do this now, by going via the definition of a generic family of context algebras as we had done in the quantifier-free version of $\mathbf{LAF}_{K1}$ (Sections 4.3.1 and 4.4.1).

DEFINITION 93 (A generic family of context algebras)

Given three sets $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$, we define $\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$ as the set of elements of the form $(\Gamma^+, \Gamma^-, \Gamma^e)$ where Γ^+ (resp. Γ^- , Γ^e) is a list of elements of $\mathcal{A}$ (resp. $\mathcal{B}$, $\mathcal{C}$).

As in Definition 83, three disjoint copies $\mathbf{Lab}_+$, $\mathbf{Lab}_-$ and $\mathbf{Lab}_e$ of the set of natural numbers are used for positive labels, negative labels and eigenlabels, respectively denoted n^+ , n^- and n^e , and we define

$(\Gamma^+, \Gamma^-, \Gamma^e)[n^+]$ as the $(n+1)^{th}$ element of Γ^+
 $(\Gamma^+, \Gamma^-, \Gamma^e)[n^-]$ as the $(n+1)^{th}$ element of Γ^-
 $(\Gamma^+, \Gamma^-, \Gamma^e)[n^e]$ as the $(n+1)^{th}$ element of Γ^e .

We turn the resulting structure

$$\left(\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}, \left(\begin{array}{c} \mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \times \mathbf{Lab}_+ \rightarrow \mathcal{A} \\ (\Gamma, n^+) \mapsto \Gamma[n^+] \end{array} \right), \left(\begin{array}{c} \mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \times \mathbf{Lab}_- \rightarrow \mathcal{B} \\ (\Gamma, n^-) \mapsto \Gamma[n^-] \end{array} \right), \left(\begin{array}{c} \mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \times \mathbf{Lab}_e \rightarrow \mathcal{C} \\ \Gamma \mapsto (n^e \mapsto \Gamma[n^e]) \end{array} \right) \right)$$

into an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context algebra, by defining the notion of extension as follows:

$$\left(\begin{array}{c} \mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \times \mathbb{D}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \rightarrow \mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}} \\ (\Gamma, \Delta) \mapsto \Gamma; \Delta \end{array} \right)$$

$$\begin{aligned} (\Gamma^+, \Gamma^-, \Gamma^e); a &:= (a::\Gamma^+, \Gamma^-, \Gamma^e) & (\Gamma^+, \Gamma^-, \Gamma^e); \sim b &:= (\Gamma^+, b::\Gamma^-, \Gamma^e) \\ (\Gamma^+, \Gamma^-, \Gamma^e); \bullet &:= (\Gamma^+, \Gamma^-, \Gamma^e) & (\Gamma^+, \Gamma^-, \Gamma^e); (\Delta_1, \Delta_2) &:= (\Gamma^+, \Gamma^-, \Gamma^e); \Delta_1; \Delta_2 \\ (\Gamma^+, \Gamma^-, \Gamma^e); c.\Delta &:= (\Gamma^+, \Gamma^-, c::\Gamma^e); \Delta \end{aligned}$$

※

We can now give the parameters that present $\mathbf{LAF}_{K1}$ as a $\mathbf{LAF}$ instance with eigenlabels:

DEFINITION 94 ($\mathbf{LAF}_{K1}$ as a $\mathbf{LAF}$ instance with eigenlabels)

$\mathbb{S}$ is the set of sorts as in Definition 81.

$\mathbf{Lab}_e$ is a copy of the set of natural numbers and $\mathbb{T}$ is the set of terms, as in Definition 82. Notice that $\mathbb{T}$ does extend $\mathbf{Lab}_e$, and substitution of eigenlabels for terms gives a systematic way to lift a function $\mathbf{Lab}_e \rightarrow \mathbf{Lab}_e$ to a function $\mathbb{T} \rightarrow \mathbb{T}$.

$\Vdash$ is the sorting relation as in Definition 82. Notice that $\Sigma \Vdash x:s$ if and only if $s = \Sigma(x)$, and that for all $\pi : \mathbf{Lab}_e \rightarrow \mathbf{Lab}_e$, if $\Sigma \circ \pi \Vdash r:s$ then $\Sigma \Vdash \pi(r):s$.

$\mathbb{A}$, $\mathbb{M}$ and $\equiv$ are as in Definition 81.

$\mathbf{Lab}_+$ and $\mathbf{Lab}_-$ are as in Definition 83, and the context algebra $\mathbf{Co}$ is the instance $\mathcal{G}_{\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{T}}$ of the generic family of contexts from Definition 93.

Given two functions $f_1 : \mathbb{A}_\downarrow \rightarrow \mathbb{A}_\downarrow$, $f_2 : \mathbb{M}_\downarrow \rightarrow \mathbb{M}_\downarrow$ and an $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{T})$ -context $(\Gamma^+, \Gamma^-, \Gamma^e)$, the result of the map operation $(f_1, f_2) \circ (\Gamma^+, \Gamma^-, \Gamma^e)$ is defined as $(f_1(\Gamma^+), f_2(\Gamma^-), \Gamma^e)$, where $f_1(\Gamma^+)$ and $f_2(\Gamma^-)$ are defined with the standard map operation on lists.

We define $\pi_{\Delta}^{\mathcal{V}}$ as the identity (note that we have $\Gamma[x] = (\Gamma; \Delta)[x]$).

We define $\text{st}_{\Delta}^{\mathcal{V}}$ as the element $\text{st}(\Delta, \text{sup}(\mathcal{V}), (n, \Pi) \mapsto \Pi)$ of $\mathbb{D}_{\text{unit}, \text{unit}, \text{Lab}_e}$, where $\text{st}(\Delta, n, f)$ is defined in continuation-passing style⁶ as follows:

$$\begin{aligned} \text{st}(\cdot, n, f) &:= f(n, \cdot) \\ \text{st}(\sim(\cdot), n, f) &:= f(n, \sim(\cdot)) \\ \text{st}(\bullet, n, f) &:= f(n, \bullet) \\ \text{st}((\cdot). \Delta), n, f) &:= \text{st}(\Delta, n+1, (n', \Pi) \mapsto f(n', (n+1)^e. \Pi)) \\ \text{st}((\Delta_1, \Delta_2), n, f) &:= \text{st}(\Delta_1, n, (n_1, \Pi_1) \mapsto \text{st}(\Delta_2, n_1, (n_2, \Pi_2) \mapsto f(n_2, (\Pi_1, \Pi_2)))) \end{aligned}$$

Finally, Pat and $\Vdash$ are as in Definition 81. ※

The only subtle things in the above definition are that:

- we defined $\pi_{\Delta}^{\mathcal{V}}$ as the identity, since the use of De Bruijn's levels avoids the need to update labels with new names every time a context is extended;⁷
- we defined $\text{st}_{\Delta}^{\mathcal{V}}$ as a data-structure that does nothing but remember the fresh eigenlabels that will be used for each construct $s.\Delta'$ within Δ .

From this alternative definition of LAF_{K1} we now have to describe how the original definition of LAF_{K1} can be recovered. More precisely, the context algebra $\mathcal{G}_{\mathbb{A}_{\downarrow}, \mathbb{M}_{\downarrow}, \mathbb{T}}$ of Definitions 93 and 94 yields the typing context algebra of Definition 83.

This is a particular case of a more general construction that turns every LAF instance with eigenlabels into a LAF instance, which we now present.

6.3.3 LAF instances with eigenlabels are LAF instances

We now show how a LAF instance with eigenlabels forms a LAF instance.

As expected, *sorting contexts* are now partial functions from eigenlabels to sorts (i.e. $\mathbb{C} = \text{Lab}_e \rightarrow \mathbb{S}$).

What remains to do is to turn the $(\mathbb{A}_{\downarrow}, \mathbb{M}_{\downarrow}, \mathbb{S})$ -context algebra Co into a proper typing context algebra in the sense of Definition 75. What is missing is the notion of extension:

$$\left(\begin{array}{c} \text{Co} \times \mathbb{D}_{\downarrow} \rightarrow \text{Co} \\ (\Gamma, (\Delta^l, \mathbf{r})) \mapsto \Gamma; (\Delta^l, \mathbf{r}) \end{array} \right)$$

We define such an extension from the notion of extension that is available in the $(\mathbb{A}_{\downarrow}, \mathbb{M}_{\downarrow}, \mathbb{S})$ -context algebra Co

$$\left(\begin{array}{c} \text{Co} \times \mathbb{D}_{\mathbb{A}_{\downarrow}, \mathbb{M}_{\downarrow}, \mathbb{S}} \rightarrow \text{Co} \\ (\Gamma', \Delta') \mapsto \Gamma'; \Delta' \end{array} \right)$$

and from the naming policies $(\mathcal{V}, \Delta) \mapsto \pi_{\Delta}^{\mathcal{V}}$ and $(\mathcal{V}, \Delta) \mapsto \text{st}_{\Delta}^{\mathcal{V}}$.

More precisely, $\Gamma; (\Delta^l, \mathbf{r})$ is defined as $\Gamma'; \Delta'$, where Γ' is an $(\mathbb{A}_{\downarrow}, \mathbb{M}_{\downarrow}, \mathbb{S})$ -context and Δ' is an $(\mathbb{A}_{\downarrow}, \mathbb{M}_{\downarrow}, \mathbb{S})$ -decomposition, obtained from Γ and $(\Delta^l, \mathbf{r})$ by using the two new functions. These functions allow us to describe the intricacies of the operation $\Gamma; (\Delta^l, \mathbf{r})$ that is completely unspecified in the abstract LAF system:

⁶i.e. for a continuation $f : (\mathbb{N} \times \mathbb{D}_{\text{unit}, \text{unit}, \text{Lab}_e}) \rightarrow \mathbb{D}_{\text{unit}, \text{unit}, \text{Lab}_e}$

⁷With De Bruijn's indices we would there have the opportunity to specify how the eigenlabels in a context Γ should be updated when Γ is extended into $\Gamma; \Delta$; namely, the indices should be raised by the number of new eigenlabels that Δ introduces.

- for instance, imagining that the eigenlabel x is mapped to s by Γ^e , we will need to know what happens to this mapping when Γ is extended into $\Gamma; (\Delta^l, \mathbf{r})$
- also, when Δ^l contains a typing decomposition of the form $s.\Delta^{s::l}$, we expect a new eigenlabel to be mapped to s and we will need to know which one it is.

The two naming policies provide this information:

- in the first example, the renaming policy $\pi_{\Delta^l}^\Gamma(x)$ provides the eigenlabel corresponding to x and mapped to s in the extended environment $\Gamma; (\Delta^l, \mathbf{r})$,⁸
- in the second example, the fresh naming policy $\text{st}_{\Delta^l}^\Gamma$ provides the names of the newly introduced eigenlabels, placed in a $(\text{unit}, \text{unit}, \text{Lab}_e)$ -decomposition with the same structure as Δ^l ; hence, it will contain a decomposition of the form $x.\Pi$ to indicate that x is the eigenlabel we are looking for (mapped to s in the extended environment).

Building the $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S})$ -decomposition Δ' from $(\Delta^l, \mathbf{r})$ thus relies on the following instantiation mechanism:

DEFINITION 95 (Instantiation of typing decompositions)

The *instantiation* $\downarrow_{\mathbf{r}}^\Pi \Delta^l$ of a typing decomposition Δ^l is defined for a list of terms $\mathbf{r}$ of length $|l|$ and a $(\text{unit}, \text{unit}, \text{Lab}_e)$ -decomposition Π that has the same structure as Δ^l , as follows:

$$\begin{array}{lll} \downarrow_{\mathbf{r}}^{()} a^l & := (a^l, \mathbf{r}) & \downarrow_{\mathbf{r}}^{()} (\sim M^l) & := \sim(M^l, \mathbf{r}) \\ \downarrow_{\mathbf{r}}^\bullet \bullet & := \bullet & \downarrow_{\mathbf{r}}^{\Pi_1, \Pi_2} (\Delta_1^l, \Delta_2^l) & := (\downarrow_{\mathbf{r}}^{\Pi_1} \Delta_1^l), (\downarrow_{\mathbf{r}}^{\Pi_2} \Delta_2^l) \\ \downarrow_{\mathbf{r}}^{x.\Pi} (s.\Delta^{s::l}) & := s.(\downarrow_{x::\mathbf{r}}^\Pi \Delta^{s::l}) \end{array}$$

※

DEFINITION 96 (Typing contexts in the sense of LAF instances)

The $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S})$ -context algebra Co of a LAF instance with eigenlabels, is turned into a typing context in the sense of LAF instances by defining the following extension operation:

Given a typing context Γ , a typing decomposition Δ^l of arity l and a list of terms $\mathbf{r}$ of length $|l|$, we define

$$\Gamma; (\Delta^l, \mathbf{r}) := ((\pi_{\Delta^l}^\Gamma, \pi_{\Delta^l}^\Gamma) \circ \Gamma); (\downarrow_{\pi_{\Delta^l}^\Gamma(\mathbf{r})}^{\text{st}_{\Delta^l}^\Gamma} \Delta^l)$$

※

The way we perform this extension can be explained as follows:

- first, the extension will rename the eigenlabels that were declared in Γ ; these eigenlabels are mentioned in the parameters of the instantiated atoms and molecules in Γ , so we use the renaming policy $\pi_{\Delta^l}^\Gamma$ to update with the new names these instantiated atoms and molecules; the result is the context

$$\Gamma' := (\pi_{\Delta^l}^\Gamma, \pi_{\Delta^l}^\Gamma) \circ \Gamma$$

- second, we turn Δ^l into a $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S})$ -decomposition as follows: the instantiated atoms and molecules at the leaves of this decomposition to produce will have their parameters based on $\mathbf{r}$; but the terms in $\mathbf{r}$ may mention the eigenlabels declared in Γ , which are now renamed, so we update $\mathbf{r}$ into $\pi_{\Delta^l}^\Gamma(\mathbf{r})$; then a parameterised atom a (resp. molecule M) at a leaf of Δ^l has an arity of the form $s_1 :: \dots s_n :: l$, and turns into the instantiated atom

⁸In other words, the eigenlabel x has been renamed $\pi_{\Delta^l}^\Gamma(x)$ in the extended environment; depending on how labels are implemented, it might be the case that x keeps its name and $\pi_{\Delta^l}^\Gamma$ is simply the identity.

$(a, x_1 :: \dots x_n :: \pi_{\Delta^l}^\Gamma(\mathbf{r}))$ (resp. molecule $(M, x_1 :: \dots x_n :: \pi_{\Delta^l}^\Gamma(\mathbf{r}))$), where $x_1, \dots, x_n$ are new eigenlabels whose names we get from the fresh naming policy $\text{st}_{\Delta^l}^\Gamma$; this results in the $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S})$ -decomposition

$$\Delta' := \downarrow_{\pi_{\Delta^l}^\Gamma(\mathbf{r})}^{\text{st}_{\Delta^l}^\Gamma} \Delta^l$$

- third, we extend Γ' with Δ' .

THEOREM 52 (The case of $\mathbf{LAF}_{K1}$) The LAF instance $\mathbf{LAF}_{K1}$, defined according to the above methodology from its definition as a LAF instance with eigenlabels (Definition 94), coincides with the direct definition of $\mathbf{LAF}_{K1}$ as a LAF instance (Sections 4.3.1 and 4.4.1). $\ast$

Proof: Clearly we have

$$\Gamma; (\Delta^l, \mathbf{r}) = \Gamma; (\downarrow_{\mathbf{r}}^{\text{st}_{\Delta^l}^\Gamma} \Delta^l)$$

with the left-hand side being defined in Definition 83 and the right-hand side being defined in Definitions 93 and 94. $\square$

As we have seen, $\pi_\Delta^\mathcal{V}(x)$ and $\text{st}_\Delta^\mathcal{V}(x)$ form a naming policy for the eigenlabels used after a (typing) context extension. More generally, the fact that an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context algebra respects this naming policy can be expressed as follows:

DEFINITION 97 (Respecting naming policies) An $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context algebra $\mathcal{G}$ respects the naming policies $(\mathcal{V}, \Delta) \mapsto \pi_\Delta^\mathcal{V}$ and $(\mathcal{V}, \Delta) \mapsto \text{st}_\Delta^\mathcal{V}$ if for all ρ and v we have

1. $\rho[x] = (\rho; v) \left[\pi_{|v|}^{\text{dom}(\rho)}(x) \right]$ for all eigenlabel $x \in \text{dom}^e(\rho)$;
2. and $\text{st}_{|v|}^{\text{dom}(\rho)}$ relates to v according to $(\text{unit} \times \mathbb{A}_\downarrow), (\text{unit} \times \mathbb{M}_\downarrow)$,
and $\{(x, (\rho; v)[x]) \mid x \in \text{dom}^e(\rho; v)\}$. $\ast$

6.4 A more concrete class of realisability algebras

Now the whole point of introducing the subclass of LAF instances that we call “with eigenlabels”, is to have a tighter Adequacy Lemma that relies on a weaker (and more systematically derivable) correlation property than typing correlation.

For this we identify a class of realisability algebras that naturally form models for LAF instances with eigenlabels.

In brief, a realisability algebra with eigenlabels is a realisability algebra where valuations are functions mapping eigenlabels to term denotations.

Assume we have a LAF instance with eigenlabels

$$(\mathbb{S}, \text{Lab}_e, \mathbb{T}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \text{Co}, \text{Pat}, \Vdash, \pi_\Delta^\mathcal{V}, \text{st}_\Delta^\mathcal{V})$$

DEFINITION 98 (Realisability algebras with eigenlabels)

A model structure with eigenlabels is a model structure where $\mathcal{C} = \text{Lab}_e \rightarrow \mathcal{T}$, satisfying

- for all $x \in \text{Lab}_e$, $\sigma : \text{Lab}_e \rightarrow \mathcal{T}$, we have $\llbracket x \rrbracket_\sigma = \sigma(x)$;
- for all $r \in \mathbb{T}$, $\sigma : \text{Lab}_e \rightarrow \mathcal{T}$ and $\pi : \text{Lab}_e \rightarrow \text{Lab}_e$, we have $\llbracket r \rrbracket_{\sigma \circ \pi} = \llbracket \pi(r) \rrbracket_\sigma$;

and where the semantic context algebra respects the naming policies.

A *realisability algebra with eigenlabels* is a realisability algebra whose model structure is a model structure with eigenlabels and where, for all $\Sigma : \mathbf{Lab}_e \rightarrow \mathbb{S}$ and $\sigma : \mathbf{Lab}_e \rightarrow \mathcal{T}$,
 $\sigma \in \llbracket \Sigma \rrbracket$ if and only if for all $x \in \mathbf{Lab}_e$ we have $\sigma(x) \in \llbracket \Sigma(x) \rrbracket$ (and $\text{Dom}(\sigma) = \text{Dom}(\Sigma)$).

※

DEFINITION 99 (Generic correlation)

Given three relations $\mathcal{R}_1 \subseteq \mathcal{A} \times \mathcal{A}'$, $\mathcal{R}_2 \subseteq \mathcal{B} \times \mathcal{B}'$ and $\mathcal{R}_3 \subseteq \mathcal{C} \times \mathcal{C}'$,
 we say that an $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context algebra $\mathcal{G}$ and an $(\mathcal{A}', \mathcal{B}', \mathcal{C}')$ -context algebra $\mathcal{G}'$ satisfy the *correlation property* for $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$ if the following holds:

For all $\Gamma \in \mathcal{G}$, $\Gamma' \in \mathcal{G}'$, $\Delta \in \mathbb{D}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$ and $\Delta' \in \mathbb{D}_{\mathcal{A}', \mathcal{B}', \mathcal{C}'}$
 if Γ relates to Γ' according to $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$
 and Δ relates to Δ' according to $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$
 then $\Gamma; \Delta$ relates to $\Gamma'; \Delta'$ according to $\mathcal{R}_1$, $\mathcal{R}_2$ and $\mathcal{R}_3$;

※

DEFINITION 100 (Correlation with eigenlabels) Given a realisability algebra with eigenlabels (for our LAF instance with eigenlabels), we define three relations⁹

$$\begin{aligned} \mathcal{R}_1^\sigma &:= \{(\mathbf{l}, (a, \mathbf{r})) \mid \mathbf{l} \in \llbracket (a, \mathbf{r}) \rrbracket_\sigma\} && \subseteq \mathcal{L} \times \mathbb{A}_\downarrow \\ \mathcal{R}_2^\sigma &:= \{(\mathbf{n}, (M, \mathbf{r})) \mid \mathbf{n} \in \llbracket (M, \mathbf{r}) \rrbracket_\sigma\} && \subseteq \mathcal{N} \times \mathbb{M}_\downarrow \\ \mathcal{R}_3 &:= \{(\mathbf{r}, s) \mid \mathbf{r} \in \llbracket s \rrbracket\} && \subseteq \mathcal{T} \times \mathbb{S} \end{aligned}$$

for any given $\sigma : \mathbf{Lab}_e \rightarrow \mathcal{T}$.

We say that $\tilde{\mathbf{Co}}$ and $\mathbf{Co}$ satisfy the *correlation with eigenlabels* property if for all $\sigma : \mathbf{Lab}_e \rightarrow \mathcal{T}$, they satisfy the correlation property for $\mathcal{R}_1^\sigma$, $\mathcal{R}_2^\sigma$ and $\mathcal{R}_3$.

※

REMARK 53 $\rho \in \llbracket \Gamma \rrbracket$ if and only if ρ relates to Γ according to $\mathcal{R}_1^{\rho^e}$, $\mathcal{R}_2^{\rho^e}$ and $\mathcal{R}_3$.

※

LEMMA 54 (Correlation with eigenlabels implies typing correlation)

If $\tilde{\mathbf{Co}}$ and $\mathbf{Co}$ satisfy the correlation with eigenlabels property, then they satisfy the typing correlation property: if $\rho \in \llbracket \Gamma \rrbracket$ and $\mathfrak{d} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket_{\rho^e}$ then $(\rho; \mathfrak{d}) \in \llbracket \Gamma; (\Delta^l, \mathbf{r}) \rrbracket$.

※

Proof: See the proof in Coq [GL14]. The main lines are as follows:

From $\rho \in \llbracket \Gamma \rrbracket$ we get that ρ relates to Γ according to $\mathcal{R}_1^{\rho^e}$, $\mathcal{R}_2^{\rho^e}$ and $\mathcal{R}_3$.

Then ρ relates to $(\pi_{\Delta^l}^\Gamma, \pi_{\Delta^l}^\Gamma) \circ \Gamma$ according to $\mathcal{R}_1^{(\rho; \mathfrak{d})^e}$, $\mathcal{R}_2^{(\rho; \mathfrak{d})^e}$ and $\mathcal{R}_3$.

From $\mathfrak{d} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket_{\rho^e}$ we get that $\mathfrak{d}$ relates to $\downarrow_{\pi_{\Delta^l}^\Gamma(\mathbf{r})}^{\text{st}_{\Delta^l}^\Gamma} \Delta^l$ according to $\mathcal{R}_1^{(\rho; \mathfrak{d})^e}$, $\mathcal{R}_2^{(\rho; \mathfrak{d})^e}$ and $\mathcal{R}_3$.

Then correlation with eigenlabels provides that $\rho; \mathfrak{d}$ relates to $\Gamma; (\Delta^l, \mathbf{r})$ according to $\mathcal{R}_1^{(\rho; \mathfrak{d})^e}$, $\mathcal{R}_2^{(\rho; \mathfrak{d})^e}$ and $\mathcal{R}_3$, which means that $(\rho; \mathfrak{d}) \in \llbracket \Gamma; (\Delta^l, \mathbf{r}) \rrbracket$.

At some point in the above proof we use the fact that the semantic context algebra respects the naming policies. $\square$

⁹In this definition we implicitly require $\llbracket (a, \mathbf{r}) \rrbracket_\sigma$ and $\llbracket (M, \mathbf{r}) \rrbracket_\sigma$ to be defined.

LEMMA 55 (Adequacy for LAF with eigenlabels)

We assume the following hypotheses:

Well-foundedness:

The LAF instance with eigenlabels is well-founded.

Correlation with eigenlabels:

$\tilde{\text{Co}}$ and Co satisfy the correlation with eigenlabels property.

Stability:

If $\mathfrak{d} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket_\sigma$ for some $\Delta^l, \sigma, \mathbf{r}$ and $\llbracket f(p) \rrbracket_{\rho; \mathfrak{d}} \in \perp$, then $\llbracket f \rrbracket_\rho \perp \tilde{p}(\mathfrak{d})$.

We conclude that, for all $\rho \in \llbracket \Gamma \rrbracket$,

1. if $\Gamma \vdash [t^+ : (M^l, \mathbf{r})]$ then $\llbracket t^+ \rrbracket_\rho \in \llbracket (M^l, \mathbf{r}) \rrbracket^+$;
2. if $\Gamma \vdash d : (\Delta^l, \mathbf{r})$ then $\llbracket d \rrbracket_\rho \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket$;
3. if $\Gamma \vdash t$ then $\llbracket t \rrbracket_\rho \in \perp$

※

Proof: Corollary of Lemmata 51 and 54. □

This Adequacy Lemma looks similar to Lemma 51, but the correlation assumption is much “weaker”: all the job is done in the extra structure with eigenlabels that we have required from terms, sorting contexts, typing contexts and valuations (and the finer-grained specifications we have imposed on them).

Indeed, correlation with eigenlabels often holds as a particular case of the more general correlation property for all relations $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$, typically when Co and $\tilde{\text{Co}}$ are respectively defined as the two instances $\mathcal{G}_{\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{T}}$ and $\mathcal{G}_{\mathcal{L}, \mathcal{N}, \mathcal{T}}$ of a generic family $(\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}})_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$ of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context algebras whose definition is “sufficiently parametric”. In particular for LAF_{K1} :

REMARK 56 Generic correlation always holds for the family $(\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}})_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$ of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$ -context algebras defined for LAF_{K1} (Definition 93).

In particular, correlation with eigenlabels holds for that system and for any of its realisability algebras where $\tilde{\text{Co}} = \mathcal{G}_{\mathcal{L}, \mathcal{N}, \mathcal{T}}$. If stability also holds for that instance and that realisability algebra, then the conclusions of the Adequacy Lemma hold. ※

The same remark would hold of any LAF instance and any realisability algebra where Co and $\tilde{\text{Co}}$ are defined from a similarly parametric family of context algebras.

6.5 Example: boolean models to prove Consistency

We now exhibit models to prove the consistency of LAF systems.

Assume we have a LAF instance with eigenlabels

$$(\mathbb{S}, \text{Lab}_e, \mathbb{T}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \text{Co}, \text{Pat}, \Vdash)$$

DEFINITION 101 (Boolean realisability algebras)

A *boolean realisability algebra* is a realisability algebra where $\perp = \emptyset$. ※

The terminology comes from the remark that in a boolean realisability algebra, $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$ can only take one of two values: $\emptyset$ or $\mathcal{N}$, depending on whether $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^+$ is empty or not.

REMARK 57 A boolean realisability algebra satisfies Stability. ※

THEOREM 58 (Consistency of LAF instances with eigenvariables)

Assume the LAF instance with eigenlabel is well-founded and assume that we have a boolean realisability algebra with eigenlabels where

- there is an empty semantic context $\rho_\emptyset$;
- $\tilde{\text{Co}}$ and Co satisfy the correlation with eigenlabels property.

Then there is no empty typing context $\Gamma_\emptyset$ and command t such that $\Gamma_\emptyset \vdash t$. ※

Proof: The previous remark provides Stability. If there was such a $\Gamma_\emptyset$ and t , then we would have $\rho_\emptyset \in \llbracket \Gamma_\emptyset \rrbracket$, and the Adequacy Lemma (Lemma 55) would conclude $\llbracket t \rrbracket_{\rho_\emptyset} \in \emptyset$. □

We provide such a realisability model that works with all parametric LAF instances with eigenlabels:

DEFINITION 102 (Trivial model for parametric LAF instances with eigenlabels)

Assume that Co is the instance $\mathcal{G}_{\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{T}}$ of a family of context algebras $(\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}})_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$.

The *trivial boolean model* for it is:

$$\begin{aligned}
 \mathcal{T} &:= \mathcal{L} := \mathcal{P} := \mathcal{N} := \text{unit} \\
 \perp &:= \emptyset \\
 \tilde{\text{Co}} &:= \mathcal{G}_{\text{unit}, \text{unit}, \text{unit}} \\
 \text{and therefore} \\
 \forall \rho \in \tilde{\text{Co}}, \forall x^+ \in \text{dom}^+(\rho), \quad \rho[x^+] &:= () \\
 \forall \rho \in \tilde{\text{Co}}, \forall x^- \in \text{dom}^-(\rho), \quad \rho[x^-] &:= () \\
 \forall \rho \in \tilde{\text{Co}}, \forall x \in \text{dom}^e(\rho), \quad \rho[x] &:= () \\
 \forall \mathfrak{d} \in \mathbb{D}, \quad \tilde{p}(\mathfrak{d}) &:= () \\
 \forall r \in \mathbb{T}, \forall \sigma \in \mathcal{C}, \quad \llbracket r \rrbracket_\sigma &:= () \\
 \forall f : \text{Pat} \rightarrow \text{Terms}, \forall \rho \in \tilde{\text{Co}}, \quad \llbracket f \rrbracket_\rho &:= () \\
 \\
 \forall s \in \mathbb{S}, \quad \llbracket s \rrbracket &:= \text{unit} \\
 \forall a^l \in \mathbb{A}_l, \forall \mathfrak{t}^l \in \mathcal{T}^l, \quad \llbracket a^l \rrbracket(\mathfrak{t}^l) &:= \text{unit}
 \end{aligned}$$

※

It is straightforward to check that the above definition does satisfy the specification of a realisability algebra with eigenlabels.

Note that, not only can $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$ only take one of the two values $\emptyset$ or unit , but $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^+$ can also only take one of the two values $\emptyset$ or unit .

We can now use such a structure to derive consistency for a large class of systems:

COROLLARY 59 (Consistency for parametric LAF instances with eigenlabels)

Assume that the LAF instance with eigenlabels is well-founded and that

- Co is the instance $\mathcal{G}_{\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{T}}$ of a family of context algebras $(\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}})_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$,
- Any two context algebras of the family $(\mathcal{G}_{\mathcal{A}, \mathcal{B}, \mathcal{C}})_{\mathcal{A}, \mathcal{B}, \mathcal{C}}$ satisfy the correlation property for all $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$.
- There is an empty $(\text{unit}, \text{unit}, \text{unit})$ -context in $\mathcal{G}_{\text{unit}, \text{unit}, \text{unit}}$,

Then there is no empty typing context $\Gamma_\emptyset$ and command t such that $\Gamma_\emptyset \vdash t$.

In particular, this is the case for LAF_{K1} . ※

The system $\mathbf{LAF}_J$ does not fall in the above category since the operation of context extension is not parametric enough: when computing $\Gamma; (a^l, \mathbf{r})$ (resp. $\Gamma; (\sim M^l, \mathbf{r})$), we have to make a case analysis on whether a^l is of the form $(l^+, \mathbf{r})$ or $(v, \mathbf{l})$ (resp. whether M^l is of the form $(N, \mathbf{l})$ or $(P, \mathbf{r})$).

But we can easily adapt the above trivial model into a not-as-trivial-but-almost model:

DEFINITION 103 (Trivial model for $\mathbf{LAF}_J$) The *trivial boolean model* for $\mathbf{LAF}_J$ is:

$$\begin{aligned} \mathcal{T} &:= \mathcal{P} := \mathbf{unit} \\ \mathcal{L} &:= \mathcal{N} := \{\mathbf{l}, \mathbf{r}\} \\ \perp &:= \emptyset \\ \tilde{\mathbf{Co}} &\text{ has semantics contexts of the form } (m^+, m^-, m^e, R), \\ &\text{where } m^+, m^-, m^e \in \mathbb{N} \text{ and } R \in \{0, 1\} \\ &\text{and an extension operation defined as follows} \\ (m^+, m^-, m^e, R); \mathbf{r} &:= (m^+ + 1, m^-, R) & (m^+, m^-, m^e, R); \mathbf{l} &:= (m^+, m^-, m^e, 0) \\ (m^+, m^-, m^e, R); \sim \mathbf{l} &:= (m^+, m^- + 1, R) & (m^+, m^-, m^e, R); \sim \mathbf{r} &:= (m^+, m^-, m^e, 1) \\ (m^+, m^-, m^e, R); \bullet &:= (m^+, m^-, m^e, R) \\ (m^+, m^-, m^e, R); () \cdot \Delta &:= (m^+, m^-, m^e + 1, R) \\ (m^+, m^-, m^e, R); (\Delta_1, \Delta_2) &:= (m^+, m^-, m^e, R); \Delta_1; \Delta_2 \end{aligned}$$

and we define

$$\begin{aligned} (m^+, m^-, m^e, R) [n^+] &:= \mathbf{r} \text{ if } n^+ < m^+ & (m^+, m^-, m^e, R) [n^-] &:= \mathbf{l} \text{ if } n^- < m^- \\ (m^+, m^-, m^e, R) [n^+] &\text{undefined otherwise} & (m^+, m^-, m^e, R) [n^-] &\text{undefined otherwise} \\ (m^+, m^-, m^e, 0) [\star^+] &:= \mathbf{l} & (m^+, m^-, m^e, 0) [\star^-] &\text{undefined} \\ (m^+, m^-, m^e, 1) [\star^+] &\text{undefined} & (m^+, m^-, m^e, 1) [\star^-] &:= \mathbf{r} \\ (m^+, m^-, m^e, R) [x_{(\perp, \mathbf{l})}^+] &:= \mathbf{l} \end{aligned}$$

$$\begin{aligned} (m^+, m^-, m^e, R) [n^e] &:= () \text{ if } n^e < m^e \\ (m^+, m^-, m^e, R) [n^e] &\text{undefined otherwise} \end{aligned}$$

$$\begin{aligned} \forall \mathfrak{d} \in \tilde{\mathbb{D}}, \quad \tilde{p}(\mathfrak{d}) &:= () \\ \forall r \in \mathbb{T}, \forall \sigma \in \mathcal{C}, \quad \llbracket r \rrbracket_\sigma &:= () \\ \forall f : \mathbf{Pat} \multimap \mathbf{Terms}, \forall \rho \in \tilde{\mathbf{Co}}, \quad \llbracket f \rrbracket_\rho &:= \mathbf{l} \text{ if every } p \in \mathbf{Dom}(f) \text{ is of the form} \\ &\quad \text{---}_r^+ \mid \text{---}_r^- \mid \bullet_r \mid (p_1, p_2) \mid \text{inj}_i(p) \\ &:= \mathbf{r} \text{ if not} \end{aligned}$$

$$\begin{aligned} \forall s \in \mathbb{S}, \quad \llbracket s \rrbracket &:= \mathbf{unit} \\ \forall (l^+, \mathbf{r}) \in \mathbb{A}_l, \forall \mathbf{tl} \in \mathcal{T}^l, \quad \llbracket (l^+, \mathbf{r}) \rrbracket(\mathbf{tl}) &:= \{\mathbf{r}\} \\ \forall (v, \mathbf{l}) \in \mathbb{A}_l, \forall \mathbf{tl} \in \mathcal{T}^l, \quad \llbracket (v, \mathbf{l}) \rrbracket(\mathbf{tl}) &:= \{\mathbf{l}\} \end{aligned}$$

※

It is straightforward to check that the above definition does satisfy the specification of a realisability algebra with eigenlabels. Moreover, $\mathbf{Co}$ and $\tilde{\mathbf{Co}}$ satisfy the correlation property with eigenlabels.

We can now use such a structure to derive consistency for $\mathbf{LAF}_J$:

THEOREM 60 (Consistency of $\mathbf{LAF}_J$)

There is no command t such that $(\Gamma^+, [], \Gamma^e, (v, \mathbf{l}, \mathbf{r})) \vdash t$ in $\mathbf{LAF}_J$.

※

Proof: Take the trivial boolean model for $\mathbf{LAF}_J$; we have Stability. Take $\rho := (|\Gamma^+|, 0, |\Gamma^e|, \mathbf{l})$;

clearly $\rho \in \llbracket (\Gamma^+, [], \Gamma^e, (v, l, \mathbf{r})) \rrbracket$, and the Adequacy Lemma (Lemma 55) would conclude $\llbracket t \rrbracket_\rho \in \emptyset$. $\square$

Chapter 7

Transforming proofs in the abstract focussed sequent calculus

Contents

7.1 Head reduction	134
7.2 Head normalisation	136
7.3 Re-using proofs	137
7.4 Cut-elimination	140
7.5 Conclusion and further work: Strong normalisation	144

In this chapter we investigate how LAF proofs can be transformed.

First and foremost, we have in mind the key process of structural proof theory: *cut-elimination*. The process is all the more interesting as it relates, through the Curry-Howard correspondence [How80], to the paradigm of computation in functional programming; and in the case of LAF systems, cut-elimination strongly relates to the very concept of pattern-matching, following [Zei09].

Let us also remember that originally, admissibility of cuts was a property used by Gentzen [Gen35] to relate the sequent calculus with cuts, which can easily be proved complete, to the cut-free sequent calculus, which is easily proved consistent. Even though we already have consistency results for LAF systems (with cuts) obtained by semantical methods (see Section 6.5), we are still interested in cut admissibility to get completeness of cut-free LAF systems. Indeed, we identified LAF systems with the perspective of using them as the basis of proof-search implementations, and knowing this property will help organising the exploration of the search-space.

The prospect of implementing proof-search also motivates the study of another kind of proof-transformation: As we shall seek to memoise the proof-search process (tabling all the proofs and sub-proofs we complete to re-use them as often as possible), we will often seek to *adapt* a previously obtained proof to a new sequent to be proved (provided of course this new sequent contains all the necessary ingredients for the proof to be replayed).

In Section 7.1, we identify a notion of *abstract machine* to reduce the proof-terms of LAF, implementing in effect a notion of *head reduction*. In Section 7.2 we prove that this reduction terminates on typed terms, for which the realisability models of Chapter 6 will play

a key role. In Section 7.3 we investigate the re-usability of proofs, by identifying with the concept of *free label* the atoms and molecules of a proved sequent that are necessary for the proof to be replayed on another sequent to prove. In Section 7.4, we will investigate how the transformations explored in the previous sections can be used to prove cut-elimination in LAF. In Section 7.5 we discuss the possibility of more general notions of reduction and the issue of Strong Normalisation.

The proof transformations explored in this chapter will prove particularly useful when using LAF in automated reasoning (see the third part of this dissertation).

7.1 Head reduction

It is natural to want to reduce $\langle f \mid pd \rangle$ to $f(p)$ “substituted by d ”. Indeed, this would be the evaluation rule of pattern-matching: we can think of p as a pattern and d as a way to fill its holes, while f is a pattern-matching function; the rule then selects the branch of f corresponding to p and depending on the pattern’s holes, and computation continues with the code in that branch where the holes have been substituted according to d .

Such a notion of substitution, however, is not yet defined. And so far d is a decomposition term: we can easily imagine using it to extend a context, but it is not a context itself.

Now following the view that “there is no such thing as a free variable” (what is thought of as free in fact bound somewhere else), we can accept that reducing $\langle f \mid pd \rangle$ is in fact done in a context ρ that assigns “values” to the “free labels” of f and d . This view is actually quite natural when thinking of evaluating programs by an abstract machine: evaluation is performed within an “environment” that maps variables to values such as *closures*.

In the case of LAF, this view helps understanding how the term decomposition d can be involved in reductions, as it can now be used to extend the local context ρ in which the command is evaluated:

$$\langle \langle f \mid pd \rangle \mid \rho \rangle \longrightarrow \langle f(p) \mid \rho; d' \rangle$$

where d' is “ d in the context ρ ”. This we could think as simply the pairing (d, ρ) , were it not for the fact that the extension $\rho; d'$ needs d' to be a decomposition, not a pair. Hence, d' will rather be the distribution of ρ down to each leaf of d .

This is formalised as follows:

DEFINITION 104 (Abstract machine for LAF)

Assume we have four sets $\mathbb{V}_+$, $\mathbb{V}_-$, $\mathbb{T}$, $\mathbb{S}$, and a $(\mathbb{V}_+, \mathbb{V}_-, \mathbb{T}, \mathbb{S})$ -context algebra with support set $\mathcal{G}$ such that the set $\mathbb{C} := (\text{Pat} \multimap \text{Terms}) \times \mathcal{G}$ is a subset of $\mathbb{V}_-$.

Elements of $\mathbb{C}$ are called *closures* and denoted $\langle f \mid \rho \rangle$ (where $f : \text{Pat} \multimap \text{Terms}$ and $\rho \in \mathcal{G}$), while elements of $\mathcal{G}$ are called *evaluation contexts*.

An *evaluation decomposition* is a $(\mathbb{V}_+, \mathbb{V}_-, \mathbb{T})$ -decomposition.

An *evaluation triple* is a triple denoted $\langle v \mid pd \rangle$ (overloading the notation for commands) where $v \in \mathbb{V}_-$, $p \in \text{Pat}$ and d is an evaluation decomposition.

A *contextualised command* is a pair denoted $\langle t \mid \rho \rangle$ where t is a command and ρ is an evaluation context.

We assume we have an *instantiation function*

$$\left(\begin{array}{l} \mathbb{T} \times \mathbb{S} \rightarrow \mathbb{T} \\ (r, \sigma) \mapsto \llbracket r \rrbracket_\sigma \end{array} \right)$$

We define the *distribution of an evaluation context ρ over a term decomposition d* , denoted $\llbracket d \rrbracket_\rho$, as the following evaluation decomposition:

$$\begin{aligned} \llbracket x^+ \rrbracket_\rho &:= \rho[x^+] \\ \llbracket f \rrbracket_\rho &:= \llbracket f \mid \rho \rrbracket \\ \llbracket \bullet \rrbracket_\rho &:= \bullet \\ \llbracket d_1, d_2 \rrbracket_\rho &:= \llbracket d_1 \rrbracket_\rho, \llbracket d_2 \rrbracket_\rho \\ \llbracket r.d \rrbracket_\rho &:= \llbracket r \rrbracket_{\rho^e} . \llbracket d \rrbracket_\rho \end{aligned}$$

The reduction relation is defined in two steps: the reduction of a contextualised command to an evaluation triple, and the reduction of an evaluation triple to a contextualised command:

$$\begin{aligned} (\text{head}_1) \quad \llbracket \langle x^- \mid pd \rangle \mid \rho \rrbracket &\longrightarrow \langle \rho[x^-] \mid p \llbracket d \rrbracket_\rho \rangle \\ (\text{head}_2) \quad \llbracket \langle f \mid pd \rangle \mid \rho \rrbracket &\longrightarrow \llbracket \langle f \mid \rho \rangle \mid p \llbracket d \rrbracket_\rho \rrbracket \\ (\text{head}_3) \quad \llbracket \langle f \mid \rho \rangle \mid pd \rrbracket &\longrightarrow \llbracket f(p) \mid \rho; d \rrbracket \end{aligned}$$

We will write $\longrightarrow_{\text{head}_{123}}^*$ for $\longrightarrow_{\text{head}_1, \text{head}_2, \text{head}_3}^*$, which will always be an alternation of $\longrightarrow_{\text{head}_1, \text{head}_2}$ and $\longrightarrow_{\text{head}_3}$.

There are no contextualised commands in normal form and evaluation triples in normal form are those of the form $\langle x^- \mid pd \rangle$.

If a contextualised command or an evaluation triple reduces by $\longrightarrow_{\text{head}_{123}}^*$ to such a normal form, we say that it *head-normalises*. ※

EXAMPLE 11 (Syntactic abstract machine)

Standard examples of abstract machine are *syntactic abstract machines*, where $\mathbb{V}_+ := \mathbf{Lab}_+$ and $\mathbb{T} := \mathbb{T}$, and $\mathbb{V}_- = \mathbf{Lab}_- \cup \mathbb{C}$. In other words, computation can substitute positive labels for positive labels, and substitute either negative labels or closures for negative labels.

Note however that this makes $\mathbb{V}_-$ and $\mathbb{C}$ mutually dependent,¹ so their exact definition can hardly be defined at this abstract level.

But for instance with $\mathbf{LAF}_{K1}$, we can adapt Definition 93 to define $\mathbb{C}$, $\mathbb{V}_-$ and the set $\mathcal{G}$ of evaluation contexts by simultaneous induction:

- $\mathbb{C} := (\mathbf{Pat} \rightarrow \mathbf{Terms}) \times \mathcal{G}$
- $\mathbb{V}_- := \mathbf{Lab}_- \cup \mathbb{C}$
- $\mathcal{G}$ is the set of elements of the form $(\Gamma^+, \Gamma^-, \Gamma^e)$ where Γ^+ (resp. Γ^- , Γ^e) is a list of elements of $\mathbf{Lab}_+$ (resp. $\mathbb{V}_-$, $\mathbb{T}$).

Once the set $\mathcal{G}$ of evaluation contexts is defined, the full evaluation context algebra is simply $\mathcal{G}_{\mathbf{Lab}_+, \mathbb{V}_-, \mathbb{T}}$ (using the notation of Definition 93).

¹Remember that $\mathbb{C}$ is $(\mathbf{Pat} \rightarrow \mathbf{Terms}) \times \mathcal{G}$, where $\mathcal{G}$ is (the support set of) a $(\mathbb{V}_+, \mathbb{V}_-, \mathbb{T}, \mathbb{S})$ -context algebra.

Similarly, the set $\mathbb{S}$ and the function

$$\left(\begin{array}{l} \mathbb{T} \times \mathbb{S} \rightarrow \mathbb{T} \\ (r, \sigma) \mapsto \langle\!\langle r \rangle\!\rangle_\sigma \end{array} \right)$$

can hardly be defined at the abstract level. But for first-order logic it is natural to define $\mathbb{S}$ as the set $\mathbf{Lab}_e \rightarrow \mathbb{T}$ of substitutions, and $\langle\!\langle r \rangle\!\rangle_\sigma$ is simply the application of substitution σ to the first-order term r . ※

7.2 Head normalisation

In this section we show that the abstract machine from Definition 104 terminates, when starting from typed proof-terms.

Mimicking the use of orthogonality models to prove strong normalisation result as in Chapter 2, we prove normalisation of the abstract machine by the use of a realisability model, in the sense of Chapter 6.

DEFINITION 105 (A realisability model for head-normalisation)

Assume we have an abstract machine defined by four sets $\mathbb{V}_+$, $\mathbb{V}_-$, $\mathbb{T}$, $\mathbb{S}$, an evaluation context algebra $\mathcal{G}$, and an instantiation function $(r, \sigma) \mapsto \langle\!\langle r \rangle\!\rangle_\sigma$.

The *head-normalisation model* for this abstract machine is

$$\begin{array}{ll} \mathcal{C} & := \mathbb{S} \\ \mathcal{T} & := \mathbb{T} \\ \mathcal{L} & := \mathbb{V}_+ \\ \mathcal{P} & := \mathbf{Pat} \times \mathbb{D}_{\mathbb{V}_+, \mathbb{V}_-, \mathbb{T}} \\ \mathcal{N} & := \mathbb{V}_- \\ v \perp pd & \text{ if the evaluation triple } \langle v \mid pd \rangle \text{ head-normalises}^2 \\ \tilde{\mathbf{Co}} & := \mathcal{G} \\ \forall \mathfrak{d} \in \tilde{\mathbb{D}}, & \tilde{p}(\mathfrak{d}) & := p\mathfrak{d} \\ \forall r \in \mathbb{T}, \forall \sigma \in \mathcal{C}, & \llbracket r \rrbracket_\sigma & := \langle\!\langle r \rangle\!\rangle_\sigma \\ \forall f : \mathbf{Pat} \rightarrow \mathbf{Terms}, \forall \rho \in \tilde{\mathbf{Co}}, & \llbracket f \rrbracket_\rho & := \langle\!\langle f \mid \rho \rangle\!\rangle \\ \forall s \in \mathbb{S}, & \llbracket s \rrbracket & := \mathbb{T} \\ \forall \Sigma \in \mathbb{C}, & \llbracket \Sigma \rrbracket & := \mathbb{S} \\ \forall a^l \in \mathbb{A}_l, \forall \mathfrak{t} \in \mathcal{T}^l, & \llbracket a^l \rrbracket(\mathfrak{t}) & := \mathbb{V}_+ \end{array}$$

※

REMARK 61 Notice that $\langle\!\langle t \mid \rho \rangle\!\rangle \rightarrow_{\text{head}_1, \text{head}_2} \llbracket t \rrbracket_\rho$. ※

THEOREM 62 (Head-normalisation of an abstract machine)

We assume the following hypotheses:

Well-foundedness:

The LAF instance is well-founded.

Typing correlation:

If $\rho \in \llbracket \Gamma \rrbracket$ and $\mathfrak{d} \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket$ then $(\rho; \mathfrak{d}) \in \llbracket \Gamma; (\Delta^l, \mathbf{r}) \rrbracket$.

We conclude that, for all $\rho \in \llbracket \Gamma \rrbracket$, if $\Gamma \vdash t$ then $\langle\!\langle t \mid \rho \rangle\!\rangle$ head-normalises. ※

²for the reduction relation defined in Definition 104

Proof: Stability is obvious for the head-normalisation model:

Assume $\llbracket f(p) \rrbracket_{\rho;d} \in \perp$. Following the previous remark, this entails that $\langle f(p) \mid \rho; d \rangle$ head-normalises. Hence, $\langle \langle f \mid \rho \rangle \mid pd \rangle$ head-normalises, which is literally what $\in \llbracket f \rrbracket_{\rho} \perp \tilde{p}(d)$ means.

So we can apply the Adequacy Lemma (Lemma 51), and obtain that $\llbracket t \rrbracket_{\rho}$ head-normalises, from which get that $\langle t \mid \rho \rangle$ head-normalises. $\square$

We now see how this applies to a syntactic abstract machine. Assume we have a well-founded LAF instance, and a syntactic abstract machine for it that features identity evaluation contexts, i.e. a family of contexts id satisfying $\text{id}[x^+] = x^+$ and $\text{id}[x^-] = x^-$.

COROLLARY 63 (Head normalisation) Assume that the evaluation context algebra $\mathcal{G}$ and Co satisfy the typing correlation.

If $\Gamma \vdash t$ then $\langle t \mid \text{id} \rangle$ head-normalises.³ ※

Proof: The valuation id^e is in $\llbracket \Gamma^e \rrbracket = \mathbb{S}$.

Every positive label x^+ is in $\llbracket (a^l, \mathbf{r}) \rrbracket_{\sigma} = \mathbb{V}_+ = \text{Lab}_+$ (for every σ, a^l and $\mathbf{r}$).

Every negative label x^- is in $\llbracket (M^l, \mathbf{r}) \rrbracket_{\sigma}^-$ (for every σ, M^l and $\mathbf{r}$), since x^- is in $\mathcal{N} = \mathbb{V}_- = \text{Lab}_- \cup \mathbb{C}$ and $x^- \perp pd$ for all $(p, d) \in \llbracket (M^l, \mathbf{r}) \rrbracket_{\sigma}^+$.⁴

Hence, the identity evaluation context id is in $\llbracket \Gamma \rrbracket$. $\square$

In particular, LAF_{K1} , LAF_{K2} , LAF_J are all head normalising.

7.3 Re-using proofs

Now, in order to have strong normalisation, and even just cut-elimination itself, our notion of abstract machine above is too weak, as it only (and deterministically) performs “head reduction”.

A state of a syntactic machine such as $\langle v \mid pd \rangle$ could almost be read back as a real command, if only we could *compute* closures such as $\langle f \mid \rho \rangle$, which we never do: just as in the weak reduction in λ -calculus, we never propagate the evaluation context ρ (which can be seen as a substitution) into f (in other words propagate it under the abstraction represented by the meta-level function f).

We could compute a closure $\langle f \mid \rho \rangle$ as a function $\langle f \rangle_{\rho} : \text{Pat} \rightarrow \text{Terms}$ such that

$$\langle f \rangle_{\rho}(p) = \langle f(p) \rangle_{\rho'; \text{idd}}$$

with the recursively defined propagation of an evaluation context ρ into a command c denoted $\langle c \rangle_{\rho}$, and where

- idd is an “identity decomposition term”, to create identity bindings for the labels in $f(p)$ introduced by the application of f to p ;
- ρ' is the update of ρ , providing the same bindings as ρ but taking care that the labels might have changed after the context extension with idd .

³for the evaluation context id with $\text{dom}^+(\text{id}) = \text{dom}^+(\Gamma)$ and $\text{dom}^-(\text{id}) = \text{dom}^-(\Gamma)$

⁴Indeed, $\langle x^- \mid pd \rangle$ is head-normalising since it cannot be reduced by the abstract machine.

But so far a LAF instance does not tell us how to infer idd and ρ' from ρ .

This is exactly the same situation as with the eigenlabels for which a LAF instance with eigenlabels provided two functions st_Δ^Γ and π_Δ^Γ to do exactly that.

We therefore enrich the concept of a LAF instance with eigenlabels as follows:

DEFINITION 106 (LAF instance with explicit label updates)

A LAF instance with explicit label updates is given by the following tuple:

$$(\mathbb{S}, \text{Lab}_e, \mathbb{T}, \Vdash, \mathbb{A}, \mathbb{M}, \equiv, \text{Lab}_+, \text{Lab}_-, \text{Co}, \mathbb{R}, \text{Pat}, \Vdash, \pi_\Delta^\vee, \text{st}_\Delta^\vee)$$

whose components are exactly as in the definition of a LAF instance with eigenlabels, except that

- The map operation of the typing context algebra Co satisfies the following property:
For all $f_1 : \mathbb{A}_\downarrow \rightarrow \mathbb{A}_\downarrow$ and $f_2 : \mathbb{M}_\downarrow \rightarrow \mathbb{M}_\downarrow$, all $(\mathbb{A}_\downarrow, \mathbb{M}_\downarrow, \mathbb{S})$ -decompositions Δ and Δ' , and all typing contexts Γ ,
If Δ relates to Δ' according to $\{(a, f_1(a)) \mid a \in \mathbb{A}_\downarrow\} \cup \{(m, f_2(m)) \mid m \in \mathbb{M}_\downarrow\}$ and the identity relation on sorts,
then $(f_1, f_2) \circ (\Gamma; \Delta)$ relates to $((f_1, f_2) \circ \Gamma; \Delta')$ according to the identity relations.
- There is a $(\text{Lab}_+, \text{Lab}_-, \text{Lab}_e)$ -context algebra $\mathbb{R}$ called the *renaming context algebra*, and equipped with a *renaming composition* that combines two renaming contexts π and π' into $\pi \circ \pi'$ so that $\pi \circ \pi'[x] = \pi[\pi'[x]]$ (resp. $\pi \circ \pi'[x^+] = \pi[\pi'[x^+]]$ and $\pi \circ \pi'[x^-] = \pi[\pi'[x^-]]$);
- we require the naming policies $\pi_\Delta^\vee$ and $\text{st}_\Delta^\vee$ to give information not only on eigenlabels, but also on positive and negative labels:

$$\left(\begin{array}{ccc} \mathbb{P}(\text{Lab}) \times \mathbb{D}_{\text{st}} & \rightarrow & \mathbb{R} \\ (\mathcal{V}, \Delta) & \mapsto & \pi_\Delta^\vee \end{array} \right) \quad \left(\begin{array}{ccc} \mathbb{P}(\text{Lab}) \times \mathbb{D}_{\text{st}} & \rightarrow & \mathbb{D}_{\text{Lab}_+, \text{Lab}_-, \text{Lab}_e} \\ (\mathcal{V}, \Delta) & \mapsto & \text{st}_\Delta^\vee \end{array} \right)$$

Clearly, we can extract from those naming policies the policies in the sense of Definition 92 (with types $(\mathbb{P}(\text{Lab}) \times \mathbb{D}_{\text{st}} \rightarrow (\text{Lab}_e \rightarrow \text{Lab}_e))$ and $(\mathbb{P}(\text{Lab}) \times \mathbb{D}_{\text{st}} \rightarrow \mathbb{D}_{\text{unit}, \text{unit}, \text{Lab}_e})$).

Finally, we require that $\mathbb{R}$ respect those naming policies.

✱

REMARK 64 It is straightforward to define $\mathbb{R}$, $\pi_{|\Delta|}^{\text{dom}(\Gamma)}$, and $\text{st}_{|\Delta|}^{\text{dom}(\Gamma)}$ in LAF_{K1} and LAF_{K2} to make them LAF instances with explicit label updates.

✱

With this information, we can now properly define the free labels of a proof-term, something which we surprisingly did not need so far, but that will indicate which parts of a typing environment are actually used in a proof.

DEFINITION 107 (Free labels) The free labels of a positive term (resp. decomposition term, command) that is typed in a typing context Γ , are defined by the rules of Fig. 29. ✱

Knowing what free variables are, we are now able, given a proof of a sequent, to replay the proof for any other sequent whose typing context contains the atoms and molecules that type the free variables of the original proof.

For this we define the renaming of a term:

DEFINITION 108 (Renaming) The renaming, denoted $\pi \cdot t^+$ (resp. $\pi \cdot d$, $\pi \cdot t$), by a renaming context π , of a positive term (resp. decomposition term, command) that is typed in a typing context Γ , is defined by the rules of Fig. 30.

✱

$\text{FL}(pd)$	$:= \text{FL}(d)$
$\text{FL}(x^+)$	$:= \{x^+\}$
$\text{FL}(f)$	$:= \bigcup_{p \in \text{Dom}(f)} \pi^{-1}(\text{FL}(f(p)))$
$\text{FL}(\bullet)$	$:= \emptyset$
$\text{FL}(d_1, d_2)$	$:= \text{FL}(d_1) \cup \text{FL}(d_2)$
$\text{FL}(r.d)$	$:= \text{FL}(r) \cup \text{FL}(d)$
$\text{FL}(\langle x^- \mid t^+ \rangle)$	$:= \{x^-\} \cup \text{FL}(t^+)$
$\text{FL}(\langle f \mid t^+ \rangle)$	$:= \text{FL}(f) \cup \text{FL}(t^+)$

where π is the function in $(\text{Lab}_+ \rightarrow \text{Lab}_+) \cup (\text{Lab}_- \rightarrow \text{Lab}_-) \cup (\text{Lab}_e \rightarrow \text{Lab}_e)$ mapping every $x^+ \in \text{Lab}_+$ to $\pi_{|p|}^{\text{dom}(\Gamma)}[x^+]$ (resp. $x^- \in \text{Lab}_-$ to $\pi_{|p|}^{\text{dom}(\Gamma)}[x^-]$, and $x \in \text{Lab}_e$ to $\pi_{|p|}^{\text{dom}(\Gamma)}[x]$).

Figure 29: Free labels

$\pi \cdot pd$	$:= p(\pi \cdot d)$
$\pi \cdot x^+$	$:= \pi[x^+]$
$\pi \cdot f$	$:= p \mapsto \left(\left(\pi_{ p }^{\text{dom}(\Gamma)} \circ \pi \right); \text{st}_{ p }^{\text{dom}(\Gamma)} \right) \cdot f(p)$
$\pi \cdot \bullet$	$:= \bullet$
$\pi \cdot (d_1, d_2)$	$:= (\pi \cdot d_1), (\pi \cdot d_2)$
$\pi \cdot (r.d)$	$:= \pi^e(r).(\pi \cdot d)$
$\pi \cdot \langle x^- \mid t^+ \rangle$	$:= \langle \pi[x^-] \mid (\pi \cdot t^+) \rangle$
$\pi \cdot \langle f \mid t^+ \rangle$	$:= \langle (\pi \cdot f) \mid (\pi \cdot t^+) \rangle$

Figure 30: Renaming

In the renaming of a function f , $\pi_{|p|}^{\text{dom}(\Gamma)} \circ \pi$ updates the co-domain of π as we went “through a binding”, and composing with $\text{st}_{|p|}^{\text{dom}(\Gamma)}$ adds the “identity bindings” for the labels introduced by the application of f to p .

Now as mentioned before, when we have a proof for a particular sequent, we want to identify when it can be replayed for another sequent. For this we define what it means for a typing context Γ' to at least contain the instantiated atoms and molecules of a typing context Γ : this is done by identifying a renaming π that will map the labels in Γ' to some labels in Γ that have the same type.

DEFINITION 109 (Context embedding)

We say that Γ embeds into Γ' along a renaming context π , written $\Gamma \sqsubseteq_{\pi} \Gamma'$, if for all x (resp. x^+, x^-) in $\text{dom}^e(\pi)$ (resp. $\text{dom}^+(\pi), \text{dom}^-(\pi)$) we have $\Gamma[x] = \Gamma'[\pi[x]]$ (resp. $\Gamma[x^+] = \Gamma'[\pi[x^+]]$, $\Gamma[x^-] = \Gamma'[\pi[x^-]]$). ※

Notice that the domain of π might be smaller than that of Γ , so that π does not necessarily map *every* label declared in Γ . This is a feature (rather than a bug) that will allow us to ignore those instantiated atoms and molecules in Γ that are not used in the proof that we

want to replay (the renaming π may only be defined on those labels that are free in the proof-term).

THEOREM 65 (Replaying a proof) Assume the following property:

Renaming correlation:

For all Γ, Γ', π , if $\Gamma \sqsubseteq_{\pi} \Gamma'$ then $\Gamma; \Delta \sqsubseteq_{\pi'} \Gamma'; \Delta$, where $\pi' = (\pi|_{\text{dom}(\Gamma)} \circ \pi); \text{st}_{|p|}^{\text{dom}(\Gamma)}$.

We conclude that, for all $\pi \in \mathbb{R}$ such that $((\rho^e, \rho^e) \circ \Gamma) \sqsubseteq_{\pi} \Gamma'$,

1. if $\text{FL}(t^+) \subseteq \text{Dom}(\pi)$ and $\Gamma \vdash [t^+ : (M^l, \mathbf{r})]$ then $\Gamma' \vdash [(\pi \cdot t^+) : (M^l, \mathbf{r})]$
2. if $\text{FL}(d) \subseteq \text{Dom}(\pi)$ and $\Gamma \vdash d : (\Delta^l, \mathbf{r})$ then $\Gamma' \vdash (\pi \cdot d) : (\Delta^l, \mathbf{r})$
3. if $\text{FL}(t) \subseteq \text{Dom}(\pi)$ and $\Gamma \vdash t$ then $\Gamma' \vdash (\pi \cdot t)$

※

Proof: See the Coq proof [GL14]. □

A LAF instance with explicit label updates thus allows us to apply a renaming to a proof to get a proof of a new sequent. This will be used heavily in an implementation of proof-search that memoises proofs in order to paste them as often as possible.

7.4 Cut-elimination

We now show how to use substitution and the substitution lemma to define a normalisation procedure, in a LAF instance with explicit label updates, to produce cut-free terms.

DEFINITION 110 (Normalisation) We take a syntactic abstract machine, whose evaluation context algebra is equipped with a *renaming operation* that associates, to a renaming context π and an evaluation context ρ , an evaluation context $\pi \circ \rho$ such that

- for all $x \in \text{dom}^e(\rho)$, we have $\pi \circ \rho[x] = \pi^e(\rho[x])$
- for all $x^+ \in \text{dom}^+(\rho)$, we have $\pi \circ \rho[x^+] = \pi[\rho[x^+]]$
- for all $x^- \in \text{dom}^-(\rho)$, we have
 - if $\rho[x^-] = y^-$ then $\pi \circ \rho[x^-] = \pi[y^-]$
 - if $\rho[x^-] = \langle f \mid \rho' \rangle$ then $\pi \circ \rho[x^-] = \langle f \mid \pi \circ \rho' \rangle$

We define the *big-step semantics* of the LAF instance with explicit label updates as two relations

- one denoted $d \Downarrow d'$ between an evaluation decomposition d and a cut-free decomposition term d' ,
- one denoted $\langle t \mid \rho \rangle \Downarrow t'$ between a contextualised command $\langle t \mid \rho \rangle$ and a cut-free command t' ,

defined by simultaneous induction by the rules of Fig. 31.

We say that a contextualised command $\langle t \mid \rho \rangle$ (resp. an evaluation decomposition d) *normalises* if there is some t' such that $\langle t \mid \rho \rangle \Downarrow t'$ (resp. some d' such that $d \Downarrow d'$). Notice in that case that t' (resp. d') is cut-free. We also say that an evaluation triple $\langle v \mid pd \rangle$ *normalises* if $\langle v \mid pd \rangle \longrightarrow_{\text{head}_{123}} \langle x^- \mid p'd' \rangle$ and d' normalises. ※

In order to show that this forms a cut-elimination procedure, we need to

$$\begin{array}{c}
\frac{\forall p \in \text{Dom}(f), \quad \llbracket f(p) \mid \left(\left(\pi_{|p|}^{\text{dom}(\rho)} \circ \rho \right) ; \text{st}_{|p|}^{\text{dom}(\rho)} \right) \rrbracket \Downarrow f'(p)}{\llbracket f \mid \rho \rrbracket \Downarrow f'} \quad \frac{}{x^- \Downarrow x^-} \\
\\
\frac{\frac{}{x^+ \Downarrow x^+} \quad \frac{}{\bullet \Downarrow \bullet} \quad \frac{d_1 \Downarrow d'_1 \quad d_2 \Downarrow d'_2}{d_1, d_2 \Downarrow d'_1, d'_2} \quad \frac{d \Downarrow d'}{r.d \Downarrow r.d'}}{\frac{\llbracket t \mid \rho \rrbracket \longrightarrow_{\text{head}_{123}}^* \langle x^- \mid pd \rangle \quad d \Downarrow d'}{\llbracket t \mid \rho \rrbracket \Downarrow \langle x^- \mid pd' \rangle}}
\end{array}$$

Figure 31: Cut-elimination

- give typing rules for contextualised commands, evaluation decomposition, and evaluation triples;
- show that $\longrightarrow_{\text{head}_{123}}^*$ satisfies Subject Reduction with these rules;
- show that every typed contextualised command normalises.

DEFINITION 111 (Typing the elements of a syntactic abstract machine)

The typing rules for the elements of a syntactic abstract machine are given in Fig. 32, where $\vdash_{\text{LAF}}$ denotes the derivability of sequents in LAF (Fig. 27). ※

THEOREM 66 (Subject Reduction)

1. If $\Gamma \vdash \llbracket t \mid \rho \rrbracket$ and $\llbracket t \mid \rho \rrbracket \longrightarrow_{\text{head}_1, \text{head}_2} \langle v \mid pd \rangle$ then $\Gamma \vdash \langle v \mid pd \rangle$.
 2. If $\Gamma \vdash \langle v \mid pd \rangle$ and $\langle v \mid pd \rangle \longrightarrow_{\text{head}_3} \llbracket t \mid \rho \rrbracket$ then $\Gamma \vdash \llbracket t \mid \rho \rrbracket$.
 3. If $\Gamma \vdash d : (\Delta, \mathbf{r})$ and $d \Downarrow d'$ then $\Gamma \vdash_{\text{LAF}} d' : (\Delta, \mathbf{r})$.
 4. If $\Gamma \vdash \llbracket t \mid \rho \rrbracket$ and $\llbracket t \mid \rho \rrbracket \Downarrow t'$ then $\Gamma \vdash_{\text{LAF}} t'$.
- ※

Proof: The first two points are given by a simple rearrangement of the sub-derivation trees. The last two points are proved by induction on the normalisation derivations. □

Finally, we adapt the realisability model for head normalisation (Definition 105) to prove cut-elimination:

DEFINITION 112 (A realisability model for normalisation)

Evaluation decompositions		
$\frac{}{\Gamma \vdash \bullet : (\bullet, \mathbf{r})}$	$\frac{\Gamma \vdash d_1 : (\Delta_1, \mathbf{r}) \quad \Gamma \vdash d_2 : (\Delta_2, \mathbf{r})}{\Gamma \vdash d_1, d_2 : ((\Delta_1, \Delta_2), \mathbf{r})}$	$\frac{\Gamma^e \Vdash r' : s \quad \Gamma \vdash d : (\Delta, r' :: \mathbf{r})}{\Gamma \vdash r'.d : s.(\Delta, \mathbf{r})}$
$\frac{\Gamma [x^+] \equiv (a, \mathbf{r})}{\Gamma \vdash x^+ : (a, \mathbf{r})}$	$\frac{\Gamma \vdash \rho : \Gamma' \quad \Gamma' \vdash_{\text{LAF}} f : (\sim M, \mathbf{r})}{\Gamma \vdash \langle\langle f \mid \rho \rangle\rangle : (\sim M, \mathbf{r})}$	

Evaluation triples	
$\frac{\Delta \Vdash p : M \quad \Gamma \vdash d : (\Delta, \mathbf{r})}{\Gamma \vdash \langle x^- \mid pd \rangle}$	$\Gamma [x^-] = (M, \mathbf{r})$
$\frac{\Gamma \vdash \langle\langle f \mid \rho \rangle\rangle : (\sim M, \mathbf{r}) \quad \Delta \Vdash p : M \quad \Gamma \vdash d : (\Delta, \mathbf{r})}{\Gamma \vdash \langle\langle\langle f \mid \rho \rangle\rangle \mid pd \rangle}$	

Contextualised commands	
$\frac{\Gamma \vdash \rho : \Gamma' \quad \Gamma' \vdash_{\text{LAF}} t}{\Gamma \vdash \langle\langle t \mid \rho \rangle\rangle}$	

Evaluation contexts	
$\left(\begin{array}{l} \forall x \in \text{dom}^e(\Gamma'), \quad \Gamma^e \Vdash \rho [x] : \Gamma' [x] \\ \forall x^+ \in \text{dom}^+(\Gamma'), \quad \Gamma [\rho [x^+]] \equiv \Gamma' [x^+] \\ \text{either } \Gamma [y^-] = \Gamma' [x^-] \quad \text{if } \rho [x^-] = y^- \\ \text{or } \Gamma \vdash \langle\langle f \mid \rho' \rangle\rangle : (\sim M, \mathbf{r}) \quad \text{if } \rho [x^-] = \langle\langle f \mid \rho' \rangle\rangle \text{ and } \Gamma' [x^-] = (M, \mathbf{r}) \end{array} \right)$	
$\Gamma \vdash \rho : \Gamma'$	

Figure 32: Typing a syntactic abstract machine

The *normalisation model* for this syntactic abstract machine is

$\mathcal{C}$	$:= \mathbb{S}$	
$\mathcal{T}$	$:= \mathbb{T}$	
$\mathcal{L}$	$:= \text{Lab}_+$	
$\mathcal{P}$	$:= \text{Pat} \times \mathbb{D}_{\text{Lab}_+, \text{Lab}_- \cup \mathbb{C}, \mathbb{T}}$	
$\mathcal{N}$	$:= \text{Lab}_- \cup \mathbb{C}$	
$v \perp (p, d)$	if the evaluation triple $\langle v \mid pd \rangle$ normalises	
$\tilde{\mathcal{C}}o$	$:= \mathcal{G}$	
$\forall \mathfrak{d} \in \tilde{\mathbb{D}},$	$\tilde{p}(\mathfrak{d})$	$:= (p, \mathfrak{d})$
$\forall r \in \mathbb{T}, \forall \sigma \in \mathcal{C},$	$\llbracket r \rrbracket_\sigma$	$:= \langle\langle r \rangle\rangle_\sigma$
$\forall f : \text{Pat} \rightarrow \text{Terms}, \forall \rho \in \tilde{\mathcal{C}}o,$	$\llbracket f \rrbracket_\rho$	$:= \langle\langle f \mid \rho \rangle\rangle$
$\forall s \in \mathbb{S},$	$\llbracket s \rrbracket$	$:= \mathbb{T}$
$\forall \Sigma \in \mathbb{C},$	$\llbracket \Sigma \rrbracket$	$:= \mathbb{S}$
$\forall a^l \in \mathbb{A}_l, \forall \mathfrak{rl} \in \mathcal{T}^l,$	$\llbracket a^l \rrbracket(\mathfrak{rl})$	$:= \text{Lab}_+$

※

REMARK 67

Notice this is the same definition as Definition 105, except for the orthogonality relation which we have strengthened by requiring normalisation instead of head-normalisation.

Of course we still have $\langle t \mid \rho \rangle \rightarrow_{\text{head}_{12}} \llbracket t \rrbracket_\rho$. ※

THEOREM 68 (Normalisation of a syntactic abstract machine)

We assume the following hypotheses:

Well-foundedness:

The LAF instance is well-founded.

Correlation with eigenlabels:

$\tilde{\text{Co}}$ and Co satisfy the correlation with eigenlabels property.

We conclude that, for all $\rho \in \llbracket \Gamma \rrbracket$, if $\Gamma \vdash t$ then $\langle t \mid \rho \rangle$ normalises. ※

Proof: Stability is obvious for the normalisation model:

Assume $\llbracket f(p) \rrbracket_{\rho; d} \in \perp$. Following the previous remark, this entails that $\langle f(p) \mid \rho; d \rangle$ normalises. Hence, $\langle \langle f \mid \rho \rangle \mid pd \rangle$ normalises, which is literally what $\in \llbracket f \rrbracket_\rho \perp \tilde{p}(d)$ means.

So we can apply the Adequacy Lemma (Lemma 55), and obtain that $\llbracket t \rrbracket_\rho$ normalises, from which get that $\langle t \mid \rho \rangle$ normalises. □

Again, assume we have a LAF instance with explicit label updates, a syntactic machine for it that features identity evaluation contexts, i.e. a family of contexts id satisfying $\text{id}[x^+] = x^+$, $\text{id}[x^-] = x^-$ and $\text{id}[x] = x$.

LEMMA 69 (Normalisation and renaming)

If d (resp. t) normalises then $\pi \cdot d$ (resp. $\pi \cdot t$) normalises. ※

Proof: By induction on the normalisation derivation. □

LEMMA 70 (Evaluation decompositions in the model are normalising)

1. For all typing decomposition Δ^l , for all $\mathbf{r}, \sigma$ and all $d \in \mathbb{D}_{\text{Lab}^+, \text{Lab}^-, \text{Lab}^e}$ with the same structure as Δ^l , $d \in \llbracket (\Delta^l, \mathbf{r}) \rrbracket_\sigma$.
2. For all molecules M^l , for all $\mathbf{r}, \sigma$ and all $\langle f \mid \rho \rangle$ in $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$, $\langle f \mid \rho \rangle$ normalises.
3. For all typing decomposition Δ^l , for all $\mathbf{r}, \sigma$ and all d in $\llbracket (\Delta^l, \mathbf{r}) \rrbracket_\sigma$, d normalises.
4. For all molecules M^l , for all $\mathbf{r}, \sigma$ and all negative labels $x^-, x^- \in \llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$. ※

Proof: By simultaneous induction on Δ^l and M^l , using the well-founded property of the LAF instance.

For point 1: by induction on Δ^l , the base case being point 4.

For point 2: by unfolding the definition of $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$,

$\langle f \mid \pi_{|p|}^{\text{dom}(\rho)} \circ \rho \rangle$ is orthogonal to $(p, \text{st}_{|p|}^{\text{dom}(\rho)})$ (using point 1).

For point 3: by induction on Δ^l , the base case being point 2.

For point 4: for all $(p, d) \in \llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^+$, we have the evaluation decomposition d in some $\llbracket (\Delta^l, \mathbf{r}) \rrbracket_\sigma$, and by point 3 d normalises; hence $\langle x^- \mid pd \rangle$ normalises (i.e. $x^- \perp pd$), so x^- is in $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$. □

COROLLARY 71 (Cut-elimination) Assume that the evaluation context algebra $\mathcal{G}$ and Co satisfy the correlation with eigenlabels property.

If $\Gamma \vdash t$ then $\langle\langle t \mid \text{id} \rangle\rangle$ normalises.⁵

Therefore the LAF instance with explicit label updates admits cuts. ※

Proof: Every eigenlabel x is in $\llbracket s \rrbracket = \mathbb{T} = \mathbb{T}$ (for every s).

Every positive label x^+ is in $\llbracket (a^l, \mathbf{r}) \rrbracket_\sigma = \mathbb{V}_+ = \text{Lab}_+$ (for every σ, a^l and $\mathbf{r}$).

Every negative label x^- is in $\llbracket (M^l, \mathbf{r}) \rrbracket_\sigma^-$ (previous lemma).

Hence, the identity evaluation context id is in $\llbracket \Gamma \rrbracket$, and we can apply the previous theorem.

Combined with Subject Reduction (Theorem 66), we can transform every proof with cuts into a cut-free proof. □

In particular, LAF_{K1} and LAF_{K2} admit cuts.

7.5 Conclusion and further work: Strong normalisation

Now, we have proved cut-elimination but not strong normalisation, as we have used a big-step operational semantics to reduce proof-terms to cut-free forms, but we still have not defined a non-deterministic reduction relation for which strong normalisation might be interesting. For this we would definitely need to compute closures (which we still have avoided so far), by pushing down evaluation contexts with the rules of Fig. 33.

$$\begin{array}{ll}
 \rho \cdot pd & := p(\rho \cdot d) \\
 \\
 \rho \cdot x^+ & := \pi[x^+] \\
 \rho \cdot f & := p \mapsto \left(\left(\pi_{|p|}^{\text{Dom}(\Gamma)} \circ \rho \right) ; \text{st}_{|p|}^{\text{Dom}(\Gamma)} \right) \cdot f(p) \\
 \rho \cdot \bullet & := \bullet \\
 \rho \cdot (d_1, d_2) & := (\rho \cdot d_1), (\rho \cdot d_2) \\
 \rho \cdot (r.d) & := \pi^e(r).(\rho \cdot d) \\
 \\
 \rho \cdot \langle x^- \mid t^+ \rangle & := \langle \pi[x^-] \mid (\rho \cdot t^+) \rangle \\
 \rho \cdot \langle f \mid t^+ \rangle & := \langle (\rho \cdot f) \mid (\rho \cdot t^+) \rangle
 \end{array}$$

Figure 33: Substitution

Using this to define a non-deterministic reduction relation, Subject Reduction for the latter would rely on a typability result for the substitution operation, similar to Theorem 65 for renamings. On the other hand, we would then avoid introducing all the extra typing rules of Fig. 32, as the reduction relation would not rely on the constructs of an abstract machine but would directly operate on terms and commands.

We conjecture that the normalisation model of Definition 112 would work, exactly as it is, to show that typed terms and commands are strongly normalising.⁶ However, we would probably need to prove the equivalent, for LAF, of the substitution lemma in λ -calculus:

⁵for an identity evaluation context id with the same domains as Γ

⁶Well, not exactly “as it is”, since instead of closures we would directly take functions from $\text{Pat} \rightarrow \text{Terms}$.

$$\left\{ \frac{P}{y} \right\} \left\{ \frac{N}{x} \right\} M = \left\{ \left\{ \frac{P}{y} \right\} \frac{N}{x} \right\} \left\{ \frac{P}{y} \right\} M$$

which we would also need if we are to prove confluence of the (non-deterministic) reduction relation. Such a lemma would broach the topic of equality between LAF proof-terms, a question that we carefully managed to avoid so far as it involves considering which equality we take on the meta-level functions $f:\mathbf{Pat} \rightarrow \mathbf{Terms}$ (extensional, intensional?).

We therefore leave all of these questions for future work.

Part III

Theorem proving

Introduction

The Sequent Calculus, even in Gentzen’s original formulation [Gen35], is not only a formalism to represent complete proofs, but it also specifies a natural, non-deterministic proof-search procedure: the gradual completion of incomplete proof-trees, starting from the one-node tree carrying a sequent to be proved, and extending the incomplete branches step-by-step until a complete proof-tree is obtained. This is called *bottom-up proof-search*, or *root-first proof-search*. It is the basic mechanism of e.g. *tableaux methods* (see e.g. [DGHP99, BG01]), and it was also used to describe and extend the logic programming paradigm [MNPS91].

As mentioned in Chapter 3, focussing was originally introduced [AP89, And92] in the framework of linear logic [Gir87], with motivations for logic programming. In other words, focussing helped designing proof-search procedures. As described in Chapter 3 (and in Part II), focussing had a important impact on the theory of complete proofs (and their semantics). We now come back to the view of focussing as an algorithmic methodology for completing incomplete proof-trees.

In [LDM11] we used a focussed sequent calculus to describe type inhabitation / proof-construction in Pure Type Systems [Bar92] (and higher-order unification), which provides (the basis for) the type theory behind several proof assistants such as Coq [Coq] or Twelf [Twe].

In this dissertation, we illustrate how the above methodology can apply to theorem proving in classical logic. This was mostly the object of our PSI project [PSI], with contributions shared with my student Mahfuza Farooque [FLM12b, FLM12a, FGL13, FGLM13, Far13].

Analytic tableaux probably form the proof-search procedures that are closest to the sequent calculus. Being much more procedure-oriented than the sequent calculus (whose theory handles complete proofs), tableaux offer an important difference in their explicit management of existential variables during search (variables that may be instantiated to conclude provability or refutability of the input), for instance via *first-order unification* in the case of pure first-order logic.

Clause tableaux provide variants of tableaux procedures that exploit a clausal formulation of the formulae to be refuted (which they share with resolution-based techniques or even SAT-solving techniques). In Farooque’s Ph.D. [Far13], clause tableaux were shown to be simulated (in a strong sense) by root-first proof-search in the focussed sequent calculus LKF (see e.g. Chapter 3 or [LM09]), when the latter is extended with the ability to change the polarity of atoms *on-the-fly* during proof-search. More interestingly, clause tableaux that satisfy *connection* properties (strong and weak connections) were shown to correspond to the construction of LKF proofs that abide by specific polarisation policies:

Connections require that, when a branch of an incomplete proof-tree (or tableau) is extended by expanding on a clause $l_1 \vee \dots \vee l_n$, thus creating n new sub-branches, then at least

one of them is closed immediately by connecting its corresponding literal l_i with a literal that was obtained earlier on the branch.

Similarly, when root-first proof-search in LKF focusses on (the negation of)⁷ the clause $l_1 \vee \dots \vee l_n$, a policy that forces the polarity of $\vee$ to be negative and forces the polarity of one literal among $l_1 \dots l_n$ to also be negative, may be used so that the synchronous phase of LKF forces the immediate closing of the branch corresponding to that literal, by “connecting it” to a previously obtained literal.

This was formalised in [Far13], which also broached the topic of reasoning *modulo a theory*. Building on the idea of (clause) *tableaux-modulo-theories* suggested by Tinelli [Tin07] in connection with SAT-modulo-theories solving (SMT-solving), we developed in the PSI project [PSI] an extension of LKF with a decision procedure, and showed its application to SMT-solving.

This is what is presented in Chapter 8. The motivation behind it is to propose a focussed sequent calculus framework where different techniques for automated (or interactive!) theorem proving can be simulated: tuning the polarities or the polarisation policies determines (or contributes to determining) the proof-search strategies that capture the said techniques, switching for instance from a tableau procedure to an SMT-procedure (such as DPLL($\mathcal{T}$)) by a simple change of polarity policy.

This aim gave rise to the implementation of the PSYCHE prototype, which is still in the early development phase, which is the object of the system description [GL13] and which is presented in Chapter 9. The system is designed as a platform for implementing the proof-search strategies that capture different theorem proving techniques. Doing so raises the question of trust, and of the correctness of an output produced by any of these implemented strategies. What the platform offers is an architecture that lets various strategies and techniques be experimented, and implemented as *plugins* via an API with PSYCHE’s *kernel*, while guaranteeing the correctness of the output. This is obtained by a somewhat transformed LCF-architecture [GMW79]. A potential application of such a platform is to offer it as a backend prover for the proof obligations produced by verification tools such as Why3 [BFM⁺13, FP13]; the strategy programming and experimenting facilities of PSYCHE could then be used to tune the behaviour of the proof-search to the specific kind of proof obligations that need to be proved, without worrying about correctness.

We then conclude this dissertation in Chapter 10 with the perspectives of PSYCHE’s development as impacted by the material developed in this dissertation, and an opening to the numerous connections with the automated reasoning literature that remain to be investigated.

⁷In sequent calculus we try to prove the negation of the formulae that tableaux methods seek to refute.

Chapter 8

DPLL($\mathcal{T}$) as proof-search in a focussed sequent calculus

Contents

8.1	A version of LKF to work modulo a theory: $LK^p(\mathcal{T})$	152
8.1.1	Background	152
8.1.2	Definitions	153
8.2	Bisimulation with the DPLL($\mathcal{T}$) procedure	156
8.2.1	The elementary DPLL($\mathcal{T}$) procedure	156
8.2.2	Simulation of the elementary DPLL($\mathcal{T}$) procedure in $LK^p(\mathcal{T})$	158
8.2.3	Completing the bisimulation	161
8.2.4	More advanced features	163
8.3	Future work: Relation to abstract focussing	163
8.3.1	On-the-fly polarisation	164
8.3.2	Extending LAF to LAF($\mathcal{T}$)	164

This chapter focusses on automated techniques for solving the *Satisfiability Modulo Theories* (SMT) family of problems, illustrating how these can be available in a system based on goal-directed proof-search. Such problems generalise propositional SAT-problems: instead of considering the satisfiability of conjunctive normal forms (CNF) over propositional variables, SMT problems concern the satisfiability of CNF over atomic propositions from a theory $\mathcal{T}$ such as linear arithmetic or bit vectors. Given a procedure deciding the consistency -with respect to $\mathcal{T}$ - of a conjunction of atoms or negated atoms, SMT-solving organises a cooperation between this procedure and SAT-solving techniques, thus providing a decision procedure for SMT-problems. This smart extension of the successful SAT-solving techniques opened a prolific area of research and led to the implementation of ever-improving tools, namely SMT-solvers, now crucial to a number of applications in software verification. The architecture of SMT-solvers is based on the extension of the Davis, Putnam, Logemann and Loveland (DPLL) procedure [DP60, DLL62] for solving SAT-problems to a procedure called DPLL($\mathcal{T}$) [NOT06] addressing SMT-problems.

This chapter does not try to improve the DPLL($\mathcal{T}$) technique itself, or current SMT-solvers based on it, but makes a step towards the integration of the technique into a sequent calculus

framework. More precisely, we investigate how we can perform each of the *steps* of $\text{DPLL}(\mathcal{T})$ as bottom-up proof-search in sequent calculus. This allows the $\text{DPLL}(\mathcal{T})$ algorithm to be applied *up-to-a-point*, where a switch to another technique can be made (depending on the newly generated goals).¹ This simulation can be seen as a first step toward a better proof-theoretical understanding of how different proof-search strategies (e.g. tableaux, resolution, $\text{DPLL}(\mathcal{T}), \dots$), geared toward different logical fragments, could efficiently cooperate inside a common platform for theorem proving.

The polarities and the focussing properties of the sequent calculus we use allow us to derive a stronger result than the mere simulation of $\text{DPLL}(\mathcal{T})$: the proofs that are the images of $\text{DPLL}(\mathcal{T})$ runs finishing on UNSAT can be characterised by a simple criterion only involving the way polarities are assigned to literals and the way formulae are placed into the focus of sequents. From this criterion we directly get a simple proof-search strategy that is bisimilar to $\text{DPLL}(\mathcal{T})$ runs: that which performs the depth-first completion of incomplete proof-trees (starting with the leftmost open leaf), using any inference steps satisfying the given criterion on polarities and focusing. The bisimulation ensures that bottom-up proof-search in sequent calculus can be as efficient as the $\text{DPLL}(\mathcal{T})$ procedure.

Section 8.1 presents the variant of System LKF (from Section 3.2) that we use to describe $\text{DPLL}(\mathcal{T})$ in terms of proof-search. Section 8.2 describes the details of how $\text{DPLL}(\mathcal{T})$ is captured: we first identify an *elementary* version of $\text{DPLL}(\mathcal{T})$ that is the direct extension of the *Classical DPLL procedure* to a background theory $\mathcal{T}$, as well as being a restriction of the full *Abstract DPLL DPLL Modulo Theories* system,² both of which can be found in [NOT06]; then we prove the bisimulation result and discuss the $\text{DPLL}(\mathcal{T})$ mechanisms that are not in our elementary version. Section 8.3 concludes by connecting the above to the abstract LAF system(s) developed in Part II of this dissertation.

8.1 A version of LKF to work modulo a theory: $\text{LK}^p(\mathcal{T})$

8.1.1 Background

Clearly in root-first proof-search, asynchronous rules can be applied eagerly (i.e. can be chained, without creating backtrack points and losing completeness), since they are invertible. Focussing says that applying synchronous rules (although possibly creating backtrack points) can also be chained without losing completeness. This forced chaining of synchronous rules can be seen for instance in the LKF system of Section 3.2, where sequent may feature a formula in its *focus*.

A sequent with a positive atom in focus must be proved immediately by an axiom on that atom; hence, the polarity of atoms greatly affects the shape of proofs. As illustrated in e.g. [LM09], the following sequent expresses the *Fibonacci* logic program (in some language where addition is primitive) and a goal $\text{fib}(n, p)$ (where n and p are closed terms):

¹In contrast, other approaches to integrating SAT- or SMT-solving to a wider theorem proving framework usually rely on the automated technique to perform a *full* run; this is the case of [Web11, AFG⁺11, BCP11, BBP11] and Lescuyer’s solver [LC09] that runs within the Coq proof assistant thanks to its *reflection* ability.

²that allows more advanced features such as *backjumping* and *clause learning*

$$\begin{aligned} & \text{fib}(0, 0), \\ & \text{fib}(1, 1), \\ & \forall i p_1 p_2 (\text{fib}(i, p_1) \Rightarrow \text{fib}(i + 1, p_2) \Rightarrow \text{fib}(i + 2, p_1 + p_2)) \\ & \vdash \text{fib}(n, p) \end{aligned}$$

The goal will be proved with backward-reasoning if the `fib` atoms are negative (yielding a proof of exponential size in n), and forward-reasoning if they are positive (yielding many proofs, one of which being linear).

In classical logic, polarities of connectives and atoms do not affect the provability of formulae, but still greatly affect the shape of proofs, and hence the basic proof-construction steps. This chapter shows how the $DPLL(\mathcal{T})$ steps correspond to proof-construction steps for an appropriate management of polarities. For this we use a variant of the LKF system [LM09] presented in Section 3.2: the sequent calculus $LK^p(\mathcal{T})$ [FGLM13, Far13].

In order to make logical sense of e.g. the primitive addition in the Fibonacci example above, we only enrich LKF with the ability to call a decision procedure to decide the consistency of conjunctions of literals w.r.t. a theory (i.e. the same as for $DPLL(\mathcal{T})$): for a theory that equates $1 + 1$ and 2 , a call to the procedure proves $p(2), p^\perp(1 + 1) \vdash$ in one step (unlike LKF's syntactic checks).

System LKF also assumes that all atoms come with a pre-determined polarity, whereas $LK^p(\mathcal{T})$ allows *on-the-fly* polarisation of atoms: the root of a proof-tree might have none of its atoms polarised, but atoms may become positive or negative as progress is made in the proof-search.

8.1.2 Definitions

In this section we present the quantifier-free fragment of the focussed sequent calculus $LK^p(\mathcal{T})$ [FGLM13, Far13]. This fragment concerns propositional classical logic *modulo a theory* and will be sufficient for the simulation of $DPLL(\mathcal{T})$.

This sequent calculus (and this logic) involves a notion of *literal* and a notion of *theory*. The reader can safely see behind this terminology the standard notions from proof theory and automated reasoning. However at this point, very little is required from or assumed about those two notions.

DEFINITION 113 (Literals)

Let $\mathcal{L}$ be a set of elements called *literals*, equipped with an involutive function called *negation* from $\mathcal{L}$ to $\mathcal{L}$. In the rest of this chapter, a possibly primed or indexed lowercase l always denotes a literal, and $l^\perp$ its negation. ✱

Another ingredient of $LK^p(\mathcal{T})$ is a *theory* $\mathcal{T}$, given in the form of an *inconsistency predicate*, a notion that we now introduce:

DEFINITION 114 (Inconsistency predicates)

An *inconsistency predicate* is a predicate over sets of literals

- satisfied by the set $\{l, l^\perp\}$ for every literal l ;
- that is upward closed (if a subset of a set satisfies the predicate, so does the set);
- such that if the sets $\mathcal{P}, l$ and $\mathcal{P}, l^\perp$ satisfy it then so does $\mathcal{P}$.

The smallest inconsistency predicate is called the *syntactical inconsistency* predicate³. If a set $\mathcal{P}$ of literals satisfies the syntactical inconsistency predicate, we say that $\mathcal{P}$ is *syntactically inconsistent*, denoted $\mathcal{P} \models$. Otherwise $\mathcal{P}$ is *syntactically consistent*.

The theory $\mathcal{T}$ in the notation $\text{LK}^p(\mathcal{T})$ is described by means of an(other) inconsistency predicate, called the *semantical inconsistency predicate*, which will be a formal parameter of the inference system defining $\text{LK}^p(\mathcal{T})$.

If a set $\mathcal{P}$ of literals satisfies the semantical inconsistency predicate, we say that $\mathcal{P}$ is *semantically inconsistent* or *inconsistent modulo theory*, denoted by $\mathcal{P} \models_{\mathcal{T}}$. Otherwise $\mathcal{P}$ is *semantically consistent* or *consistent modulo theory*. *

DEFINITION 115 (Formulae, negation)

Let $\mathcal{L}$ be a set of literals. The formulae of propositional polarised classical logic are given by the following grammar:

$$\begin{array}{lcl} \text{Formulae } A, B, \dots & ::= & l \quad \text{where } l \text{ ranges over } \mathcal{L} \\ & | & A \wedge^+ B \mid A \vee^+ B \mid \top^+ \mid \perp^+ \\ & | & A \wedge^- B \mid A \vee^- B \mid \top^- \mid \perp^- \end{array}$$

The size of a formula A , denoted $\sharp(A)$, is its size as a tree (number of nodes).

Let $\mathcal{P} \subseteq \mathcal{L}$ be syntactically consistent. Intuitively, it represents the set of literals declared to be positive.

We define $\mathcal{P}$ -positive formulae and $\mathcal{P}$ -negative formulae as the formulae generated by the following grammars:

$$\begin{array}{lcl} \mathcal{P}\text{-positive formulae } P, \dots & ::= & p \mid A \wedge^+ B \mid A \vee^+ B \mid \top^+ \mid \perp^+ \\ \mathcal{P}\text{-negative formulae } N, \dots & ::= & p^\perp \mid A \wedge^- B \mid A \vee^- B \mid \top^- \mid \perp^- \end{array}$$

where p ranges over $\mathcal{P}$.

Let $\mathcal{U}_{\mathcal{P}}$ be the set of all $\mathcal{P}$ -unpolarised literals, i.e. literals that are neither $\mathcal{P}$ -positive nor $\mathcal{P}$ -negative.

Negation is recursively extended into an involutive map on formulae as follows:

$(A \wedge^+ B)^\perp \quad := \quad A^\perp \vee^- B^\perp$	$(A \wedge^- B)^\perp \quad := \quad A^\perp \vee^+ B^\perp$
$(A \vee^+ B)^\perp \quad := \quad A^\perp \wedge^- B^\perp$	$(A \vee^- B)^\perp \quad := \quad A^\perp \wedge^+ B^\perp$
$(\top^+)^\perp \quad := \quad \perp^-$	$(\top^-)^\perp \quad := \quad \perp^+$
$(\perp^+)^\perp \quad := \quad \top^-$	$(\perp^-)^\perp \quad := \quad \top^+$

*

REMARK 72 Note that, given a syntactically consistent set $\mathcal{P}$ of literals, negations of $\mathcal{P}$ -positive formulae are $\mathcal{P}$ -negative and vice versa. *

NOTATION 116 A possibly primed or indexed Γ always denotes a set of formulae. By Γ_{lit} we denote the subset of elements of Γ that are literals, and we write $l \in \Gamma$ if l or $l^\perp$ appears in Γ . By $\text{lit}_{\mathcal{P}}(\Gamma)$ we denote the sub-multiset of Γ consisting of its $\mathcal{P}$ -positive literals (i.e. $\mathcal{P} \cap \Gamma$ as a set). *

³It is the predicate that is true of a set $\mathcal{P}$ of literals iff $\mathcal{P}$ contains both l and $l^\perp$ for some $l \in \mathcal{L}$.

DEFINITION 117 (System $LK^p(\mathcal{T})$)

The system $LK^p(\mathcal{T})$ is the sequent calculus defined by the rules of Fig. 34, which fall into three categories: *synchronous*, *asynchronous*, and *structural* rules, and manipulate two kinds of sequents:

$$\begin{array}{l} \Gamma \vdash^{\mathcal{P}} [A] \quad \text{where the formula } A \text{ is in the } \textit{focus} \text{ of the sequent} \\ \Gamma \vdash^{\mathcal{P}} \Gamma' \end{array}$$

where $\mathcal{P}$ is a syntactically consistent set of literals declared to be positive.

A sequent of the second kind where Γ' is empty is called *developed*. ※

Synchronous rules

$$\begin{array}{lll} (\wedge^+) \frac{\Gamma \vdash^{\mathcal{P}} [A] \quad \Gamma \vdash^{\mathcal{P}} [B]}{\Gamma \vdash^{\mathcal{P}} [A \wedge^+ B]} & (\vee^+) \frac{\Gamma \vdash^{\mathcal{P}} [A_i]}{\Gamma \vdash^{\mathcal{P}} [A_1 \vee^+ A_2]} \\ (\top^+) \frac{}{\Gamma \vdash^{\mathcal{P}} [\top^+]} & (\text{Init}_1) \frac{\text{lit}_{\mathcal{P}}(\Gamma), l^\perp \models_{\mathcal{T}} l \in \mathcal{P}}{\Gamma \vdash^{\mathcal{P}} [l]} & (\text{Release}) \frac{\Gamma \vdash^{\mathcal{P}} N}{\Gamma \vdash^{\mathcal{P}} [N]} \quad N \text{ is not } \mathcal{P}\text{-positive} \end{array}$$

Asynchronous rules

$$\begin{array}{lll} (\wedge^-) \frac{\Gamma \vdash^{\mathcal{P}} A, \Delta \quad \Gamma \vdash^{\mathcal{P}} B, \Delta}{\Gamma \vdash^{\mathcal{P}} A \wedge^- B, \Delta} & (\vee^-) \frac{\Gamma \vdash^{\mathcal{P}} A_1, A_2, \Delta}{\Gamma \vdash^{\mathcal{P}} A_1 \vee^- A_2, \Delta} \\ (\perp^-) \frac{\Gamma \vdash^{\mathcal{P}} \Delta}{\Gamma \vdash^{\mathcal{P}} \Delta, \perp^-} & (\top^-) \frac{}{\Gamma \vdash^{\mathcal{P}} \Delta, \top^-} & (\text{Store}) \frac{\Gamma, A^\perp \vdash^{\mathcal{P}; A^\perp} \Delta}{\Gamma \vdash^{\mathcal{P}} A, \Delta} \quad \begin{array}{l} A \text{ is a literal} \\ \text{or is } \mathcal{P}\text{-positive} \end{array} \end{array}$$

Structural rules

$$(\text{Select}) \frac{\Gamma, P^\perp \vdash^{\mathcal{P}} [P]}{\Gamma, P^\perp \vdash^{\mathcal{P}}} \quad P \text{ is not } \mathcal{P}\text{-negative} \quad (\text{Init}_2) \frac{\text{lit}_{\mathcal{P}}(\Gamma) \models_{\mathcal{T}}}{\Gamma \vdash^{\mathcal{P}}}$$

where $\mathcal{P}; A := \mathcal{P}, A$ if $A \in \mathcal{U}_{\mathcal{P}}$
 $\mathcal{P}; A := \mathcal{P}$ if not

 Figure 34: System $LK^p(\mathcal{T})$

The gradual proof-tree construction defined by the bottom-up application of the inference rules of $LK^p(\mathcal{T})$, is a goal-directed mechanism whose intuition can be given as follows:

Asynchronous rules are invertible: $(\wedge^-)$ and $(\vee^-)$ are applied eagerly when trying to construct the proof-tree of a given sequent; (Store) is applied when hitting a positive formula or a negative literal on the right-hand side of a sequent, storing its negation on the left.

When the right-hand side of a sequent becomes empty (i.e. the sequent is *developed*), a sanity check can be made with (Init_2) to check the semantical consistency of the stored literals (w.r.t. the theory), otherwise a choice must be made to place a positive formula in focus, using rule (Select) , before applying synchronous rules like $(\wedge^+)$ and $(\vee^+)$. Each such rule decomposes the formula in focus, keeping the revealed sub-formulae in the focus of the corresponding premises, until a positive literal or a non-positive formula is obtained: the former case must be closed immediately with (Init_1) calling the decision procedure, and the latter case uses the (Release) rule to drop the focus and start applying asynchronous

rules again. The synchronous and the structural rules are in general not invertible,⁴ so each application of those yields in general a backtrack point in the proof-search.

Notice that an invariant of such a proof-tree construction process is that the left-hand side of a sequent only contains negative formulae and positive literals.

NOTATION 118 When F is a formula of unpolarised propositional logic and Ψ is a set of such formulae, $\Psi \models F$ means that Ψ entails F in propositional classical logic. Given a theory $\mathcal{T}$ (given by a semantical inconsistency predicate), we define the set of all *theory lemmas* as $\Psi_{\mathcal{T}} := \{l_1 \vee \dots \vee l_n \mid l_1^\perp, \dots, l_n^\perp \models_{\mathcal{T}}\}$ and generalise the notation $\models_{\mathcal{T}}$ to write $\Psi \models_{\mathcal{T}} F$ when $\Psi_{\mathcal{T}}, \Psi \models F$. In that case we say that F is a *semantical consequence* of Ψ . For any polarised formula A , let $\underline{A}$ be the unpolarised formula obtained by removing all polarities on connectives. ※

THEOREM 73 (Cut-elimination and Completeness of $\mathbf{LK}^p(\mathcal{T})$)

- The following rules are admissible in $\mathbf{LK}^p(\mathcal{T})$:

$$(\text{Pol}) \frac{\Gamma \vdash^{\mathcal{P}, l}}{\Gamma \vdash^{\mathcal{P}}} \quad l \in \Gamma \text{ and } \text{lit}_{\mathcal{P}}(\Gamma), l^\perp \models_{\mathcal{T}} \quad (\text{cut}) \frac{\Gamma \vdash^{\mathcal{P}} l \quad \Gamma \vdash^{\mathcal{P}} l^\perp}{\Gamma \vdash^{\mathcal{P}}} \quad l \in \Gamma$$

provided the bottom sequent satisfies some property called *safety* [FGL13, Far13].

- If $\models_{\mathcal{T}} F$, then for all A such that $\underline{A} = F$, we can prove $\vdash^\emptyset A$ in $\mathbf{LK}^p(\mathcal{T})$. ※

The meta-theory of $\mathbf{LK}^p(\mathcal{T})$, in particular the proofs of the above, can be found in [FGL13, Far13].

8.2 Bisimulation with the DPLL($\mathcal{T}$) procedure

8.2.1 The elementary DPLL($\mathcal{T}$) procedure

Intuitively, DPLL($\mathcal{T}$) aims at proving the inconsistency of a set of *clauses* with respect to a theory. We therefore retain from the previous section the notion of literal and inconsistencies, and introduce clauses:

DEFINITION 119 (Clause)

A *clause* is a finite set of literals, which can be seen as their disjunction.

In the rest of the chapter, a possibly indexed upper cased C always denotes a clause. The empty clause is denoted by $\perp$. The number of literals in a clause C is denoted $\sharp(C)$. The possibly indexed symbol ϕ always denotes finite sets of clauses $\{C_1, \dots, C_n\}$, which can also be seen as a Conjunctive Normal Form (CNF). We use $\sharp(\phi)$ to denote the sum of the sizes of the clauses in ϕ . Finally $\text{lit}(\phi)$ denotes the set of literals that appear in ϕ or whose negations appear in ϕ .

Viewing clauses as disjunctions of literals and sets of clauses as CNF, we will generalise Notation 118, writing for instance $\phi \models C^\perp$ or $\phi \models C$, as well as $\phi \models_{\mathcal{T}} C^\perp$ or $\phi \models_{\mathcal{T}} C$. ※

⁴(but they may be so, e.g. $(\wedge^+)$)

DEFINITION 120 (Decision literals and sequences)

We consider a (single) copy of the set $\mathcal{L}$ of literals, denoted $\mathcal{L}^d$, whose elements are called *decision literals*, which are just tagged clones of the literals in $\mathcal{L}$. Decision literals are denoted⁵ by l^d .

We use the possibly indexed symbol Δ to denote a finite sequence of possibly tagged literals, with $\emptyset$ denoting the empty sequence. We also use Δ_1, Δ_2 and Δ_1, l, Δ_2 to denote the suggested concatenation of sequences.

For such a sequence Δ , we write $|\Delta|$ for the subset of $\mathcal{L}$ containing all the literals in Δ with their potential tags removed. The sequences that DPLL($\mathcal{T}$) will construct will always be duplicate-free, so the difference between Δ and $|\Delta|$ is just a matter of tags and ordering. When the context is unambiguous, we will sometimes use Δ when we mean $|\Delta|$.

We define $\text{Sat}(\Delta) := \{l \mid \Delta, l^\perp \models_{\mathcal{T}}\}$, the closure of a sequence Δ by semantical entailment. For any set of clauses ϕ , the set of literals occurring in ϕ that are semantically entailed by Δ is denoted by $\text{Sat}_\phi(\Delta) := \text{Sat}(\Delta) \cap \text{lit}(\phi)$. ※

REMARK 74 Semantical consequences are the analogues of the consequences of a partial boolean assignment in the context of a DPLL procedure for propositional logic without theory.

Obviously, if $l \in \Delta$, then $l \in \text{Sat}(\Delta)$. If $\phi_1 \subseteq \phi_2$, then for any Δ , $\text{Sat}_{\phi_1}(\Delta) \subseteq \text{Sat}_{\phi_2}(\Delta)$. ※

We can now describe the elementary DPLL($\mathcal{T}$) procedure as a transition system between states.

DEFINITION 121 (Elementary DPLL($\mathcal{T}$))

A *state* of the DPLL($\mathcal{T}$) procedure is either the state UNSAT, or a pair denoted $\Delta \parallel \phi$, where ϕ is a set of clauses and Δ is a sequence of possibly tagged literals. The *transition rules* of the elementary DPLL($\mathcal{T}$) procedure are given in Fig. 35. ※

Decide	$\Delta \parallel \phi$	$\Rightarrow \Delta, l^d \parallel \phi$	where $l \in \text{lit}(\phi)$	and $l \notin \Delta$ and $l^\perp \notin \Delta$.
Propagate	$\Delta \parallel \phi, C \vee l$	$\Rightarrow \Delta, l \parallel \phi, C \vee l$	where $\Delta \models C^\perp$	and $l \notin \Delta$ and $l^\perp \notin \Delta$.
Propagate $_{\mathcal{T}}$	$\Delta \parallel \phi$	$\Rightarrow \Delta, l \parallel \phi$	where $l \in \text{Sat}_\phi(\Delta)$	and $l \notin \Delta$ and $l^\perp \notin \Delta$.
Fail	$\Delta \parallel \phi, C$	$\Rightarrow \text{UNSAT}$,	where $\Delta \models C^\perp$	and there is no decision literal in Δ .
Fail $_{\mathcal{T}}$	$\Delta \parallel \phi$	$\Rightarrow \text{UNSAT}$,	where $\Delta \models_{\mathcal{T}}$	and there is no decision literal in Δ .
Backtrack	$\Delta_1, l^d, \Delta_2 \parallel \phi, C$	$\Rightarrow \Delta_1, l^\perp \parallel \phi, C$	where $\Delta_1, l, \Delta_2 \models C^\perp$	and there is no decision literal in Δ_2 .
Backtrack $_{\mathcal{T}}$	$\Delta_1, l^d, \Delta_2 \parallel \phi$	$\Rightarrow \Delta_1, l^\perp \parallel \phi$	where $\Delta_1, l, \Delta_2 \models_{\mathcal{T}}$	and there is no decision literal in Δ_2 .

Figure 35: Elementary DPLL($\mathcal{T}$)

This transition system is an extension of the *Classical DPLL procedure*, as presented in [NOT06], to the background theory $\mathcal{T}$.⁶ The first four rules are explicitly taken from the Abstract DPLL Modulo Theories system of [NOT06].⁷ The other rules of that system (namely $\mathcal{T}$ -Backjump, $\mathcal{T}$ -Learn, $\mathcal{T}$ -Forget, etc), are not considered here in their full generality, but specific cases and combinations are covered by the rest of our elementary DPLL($\mathcal{T}$) sys-

⁵This exponent tag is a standard notation, standing for “decision”.

⁶We removed the Pure Literal rule, in general unsound in presence of a theory $\mathcal{T}$.

⁷Unit Propagate and Theory Propagate are renamed as Propagate and Propagate $_{\mathcal{T}}$ for consistency with the other rule names.

tem, so that it is logically complete.⁸ Note that this transition system is not deterministic: for instance the **Decide** rule can be applied from any state and it furthermore does not enforce a strategy for picking the literal to be tagged among the eligible elements of $\text{lit}(\phi)$. At the level of implementation, this (non deterministic) transition system is turned into a deterministic algorithm, whose efficiency crucially relies on the strategies adopted to perform the choices left unspecified by $\text{DPLL}(\mathcal{T})$.

We illustrate those rules, in the theory $\mathcal{T}$ of *Linear Rational Arithmetic*, with the two basic examples of elementary $\text{DPLL}(\mathcal{T})$ runs presented in Fig. 36 (where Δ and ϕ always refer to the current state $\Delta \parallel \phi$).

$\emptyset$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate
$x > 0$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate
$x > 0, (x + y > 0)^\perp$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate$_{\mathcal{T}}$ $((y > 0)^\perp \in \text{Sat}_\phi(\Delta))$
$x > 0, (x + y > 0)^\perp, (y > 0)^\perp$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate$_{\mathcal{T}}$ $((x = -1)^\perp \in \text{Sat}_\phi(\Delta))$
$x > 0, (x + y > 0)^\perp, (y > 0)^\perp, (x = -1)^\perp$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Fail on clause $(y > 0 \vee x = -1)$
UNSAT		

$\emptyset$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate
$x > 0$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate
$x > 0, (x + y > 0)^\perp$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate$_{\mathcal{T}}$ $((y > 0)^\perp \in \text{Sat}_\phi(\Delta))$
$x > 0, (x + y > 0)^\perp, (y > 0)^\perp$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Propagate $(x = -1)$ from $(y > 0 \vee x = -1)$
$x > 0, (x + y > 0)^\perp, (y > 0)^\perp, (x = -1)$	$x > 0, (x + y > 0)^\perp, (y > 0 \vee x = -1)$	Fail$_{\mathcal{T}}$ $x > 0, x = -1$ inconsistent with $\mathcal{T}$
UNSAT		

Figure 36: Examples of elementary $\text{DPLL}(\mathcal{T})$ runs

A reason to introduce rule **Fail $_{\mathcal{T}}$** is to allow the second run to finish with the same output as the first: Indeed, the last **Propagate** step has created a $\mathcal{T}$ -inconsistency from which we could not derive **UNSAT** without a **Fail $_{\mathcal{T}}$** step.⁹

8.2.2 Simulation of the elementary $\text{DPLL}(\mathcal{T})$ procedure in $\text{LK}^p(\mathcal{T})$

The aim of this section is to describe how the elementary $\text{DPLL}(\mathcal{T})$ procedure can be transposed into a proof-search process for sequents of the $\text{LK}^p(\mathcal{T})$ calculus. A complete and successful run of the $\text{DPLL}(\mathcal{T})$ procedure is a sequence of transitions $\emptyset \parallel \phi \Rightarrow^* \text{UNSAT}$, which ensures that the set of clauses ϕ is inconsistent modulo the theory. Hence, we are devising a proof-search process aiming at building an $\text{LK}^p(\mathcal{T})$ proof-tree for sequents of the form $\phi' \vdash$, where ϕ' represents the set of clauses ϕ as a sequent calculus structure, in the following

⁸**Backtrack** is a restricted version of **$\mathcal{T}$ -Backjump** (this holds on the basis that the full system satisfies some basic invariant -Lemma 3.6 of [NOT06]), **Fail $_{\mathcal{T}}$** (resp. **Backtrack $_{\mathcal{T}}$**) is a combination of **$\mathcal{T}$ -Learn**, **Fail** (resp. **Backtrack**), and **$\mathcal{T}$ -Forget** steps.

⁹(or, alternatively, a **$\mathcal{T}$ -Learn** step in [NOT06])

sense:

DEFINITION 122 (Representation of clauses as formulae)

An $\text{LK}^p(\mathcal{T})$ formula C' represents a $\text{DPLL}(\mathcal{T})$ clause $\{l_j\}_{j=1\dots p}$ if $C' = l_1 \vee^- \dots \vee^- l_p \vee^- \perp^-$.
 A set of formulae ϕ' represents a set of clauses ϕ if there is a bijection f from ϕ to ϕ' such that for all clauses C in ϕ , $f(C)$ represents C . *

REMARK 75 If C' represents C , then $\sharp(C') \leq 2\sharp(C)$ (there are fewer symbols $\vee^-$ than there are literals in C). *

Note here that we carefully use the negative disjunction connective to translate $\text{DPLL}(\mathcal{T})$ clauses. This is crucial not only to mimic $\text{DPLL}(\mathcal{T})$ without duplicating formulae but more generally to control the search space.

Now, in order to construct a proof of $\phi' \vdash$ from a run $\emptyset \parallel \phi \Rightarrow^* \text{UNSAT}$, we proceed gradually by considering the intermediate steps of the $\text{DPLL}(\mathcal{T})$ run:

$$\emptyset \parallel \phi \Rightarrow^* \Delta \parallel \phi \Rightarrow^* \text{UNSAT}$$

In the intermediate $\text{DPLL}(\mathcal{T})$ state $\Delta \parallel \phi$, the sequence Δ is a log of both the search space explored so far (in $\emptyset \parallel \phi \Rightarrow^* \Delta \parallel \phi$) and the search space that remains to be explored (in $\Delta \parallel \phi \Rightarrow^* \text{UNSAT}$). In this log, a tagged decision literal l^d indicates a point where the procedure has made an exploratory choice (the case where l is true has been/is being explored, the case where $l^\perp$ is true remains to be explored), while untagged literals in Δ are predictable consequences of the decisions made so far and of the set of clauses ϕ to be falsified.

If we are to express the $\text{DPLL}(\mathcal{T})$ procedure as the gradual construction of a $\text{LK}^p(\mathcal{T})$ proof-tree, we should get from $\emptyset \parallel \phi \Rightarrow^* \Delta \parallel \phi$ a proof-tree that is not yet complete and get from $\Delta \parallel \phi \Rightarrow^* \text{UNSAT}$ some (complete) proof-tree(s) that can be “plugged into the holes” of the incomplete tree. We should read in Δ the “interface” between the incomplete tree that has been constructed and the complete sub-trees to be constructed.

We use the plural here since there can be more than one sub-tree left to construct: $\Delta \parallel \phi \Rightarrow^* \text{UNSAT}$ contains the information to build not only a proof of $|\Delta|, \phi' \vdash$, but also proofs of the sequents corresponding to the other parts of the search space to be explored, characterised by the tagged literals in Δ . For instance, a run from $l_1, l_2^d, l_3, l_4^d \parallel \phi \Rightarrow^* \text{UNSAT}$ contains the information to build a proof of $l_1, l_2, l_3, l_4, \phi' \vdash$ but also the proofs of $l_1, l_2, l_3, l_4^\perp, \phi' \vdash$ and $l_1, l_2^\perp, \phi' \vdash$. Those extra sequents are obtained by collecting from a sequence Δ its “backtrack points” as follows:

DEFINITION 123 (Backtrack points)

The *backtrack points* $\llbracket \Delta \rrbracket$ of a sequence Δ of possibly tagged literals is the list of sets of untagged literals recursively defined by the following rules, where $[]$ and $::$ are the standard list constructors.

$\llbracket () \rrbracket$	$:= []$
$\llbracket \Delta, l \rrbracket$	$:= \llbracket \Delta \rrbracket$
$\llbracket \Delta, l^d \rrbracket$	$:= \Delta, l^\perp :: \llbracket \Delta \rrbracket$

*

REMARK 76 The length of $\llbracket \Delta \rrbracket$ is the number of decision literals in Δ . *

Now, coming back to the $\text{DPLL}(\mathcal{T})$ transition sequence $\emptyset \parallel \phi \Rightarrow^* \Delta \parallel \phi$ and its intuitive counterpart in sequent calculus, we have to formalise the notion of *incomplete* proof-tree together with the notion of “filling its holes”:

DEFINITION 124 (Incomplete proof-tree, extension)

An *incomplete proof-tree* in $\text{LK}^p(\mathcal{T})$ is a tree labelled with sequents,

- whose leaves are tagged as either *open* or *closed*;
- whose open leaves are labelled with developed sequents;
- and such that every node that is not an open leaf, together with its children, forms an instance of the $\text{LK}^p(\mathcal{T})$ rules.

The *size* of an incomplete proof-tree is its number of nodes.

An incomplete proof-tree π' is an *extension* of π , if there is a tree (edge and nodes preserving) homomorphism from π to π' . It is an *n-extension* of π , if moreover the difference of size between π' and π is less than or equal to n . ✱

REMARK 77 An incomplete proof-tree that has no open leaf is (isomorphic to) a well-formed complete $\text{LK}^p(\mathcal{T})$ proof of the sequent labelling its root. In that case, we say the proof-tree is *complete*. ✱

The intuition that an intermediate DPLL($\mathcal{T}$) state describes an “interface” between an incomplete proof-tree and the complete proof-trees that should be plugged into its holes, is formalised as follows:

DEFINITION 125 (Correspondence)

An incomplete proof-tree π *corresponds to* a DPLL($\mathcal{T}$) state $\Delta \parallel \phi$ if:

- the length of $|\Delta| :: \llbracket \Delta \rrbracket$ is the number of open leaves of π ;
- if Δ_i is the i^{th} element of $|\Delta| :: \llbracket \Delta \rrbracket$, then the i^{th} open leaf of π (taken left-to-right) is labelled by a developed sequent of the form $\Delta'_i, \phi'_i \vdash^{\Delta_i}$, where:
 - ϕ'_i represents ϕ (in the sense of Definition 122);
 - $\text{Sat}_\phi(\Delta_i) = \text{Sat}_\phi(\Delta'_i)$.

An incomplete proof-tree π *corresponds to* the state UNSAT if it has no open leaf. ✱

REMARK 78 In the general case, different incomplete proof-trees might correspond to the same DPLL($\mathcal{T}$) state (just like different DPLL($\mathcal{T}$) runs may reach that state from the initial one).

Note that we do not require anything from the conclusion of an incomplete proof-tree corresponding to $\Delta \parallel \phi$: just as our correspondence says nothing about the DPLL($\mathcal{T}$) transitions taking place after $\Delta \parallel \phi$ (nor about the trees to be plugged into the open leaves), it says nothing about the transitions taking place before $\Delta \parallel \phi$ (nor about the incomplete proof-tree, except for its open leaves).

If an incomplete proof-tree π corresponds to a DPLL($\mathcal{T}$) state $\Delta \parallel \phi$ where there are no decision literals in Δ , then there is exactly one open leaf in π , and it is labelled by a sequent of the form $\Delta', \phi' \vdash^{|\Delta|}$, where ϕ' represents ϕ and $\text{Sat}_\phi(\Delta) = \text{Sat}_\phi(\Delta')$.

To the initial state $\emptyset \parallel \phi$ of a run of the DPLL($\mathcal{T}$) procedure corresponds the incomplete proof-tree consisting of one node (both root and open leaf) labelled with the sequent $\phi' \vdash$, where ϕ' represents ϕ . ✱

The simulation theorem below provides a systematic way of interpreting any DPLL($\mathcal{T}$) transition as a completion of incomplete proof-trees that preserves the correspondence given in Definition 125 and controls the growth of the proof trees.

THEOREM 79 (Simulation of DPLL($\mathcal{T}$) in $\text{LK}^p(\mathcal{T})$)

If $\Delta \parallel \phi \Rightarrow \mathcal{S}_2$ is a valid DPLL($\mathcal{T}$) transition, and π_1 is an incomplete proof tree in $\text{LK}^p(\mathcal{T})$ corresponding to $\Delta \parallel \phi$, then there exists a $(2\sharp(\phi) + 3)$ -extension π_2 of π_1 that corresponds to $\mathcal{S}_2$. ※

Proof: See [FGLM13, Far13]. By case analysis on the nature of the transition, completing the leftmost open leaf of π_1 . Basically:

• Fail using clause C	corresponds to	Select on $C^\perp$
• Fail $_{\mathcal{T}}$	corresponds to	Init ₂ rule
• Backtrack using clause C	corresponds to	Select on $C^\perp$
• Backtrack $_{\mathcal{T}}$	corresponds to	Init ₂ rule
• Propagate using clause C	corresponds to	Select on $C^\perp$
• Fail $_{\mathcal{T}}$	corresponds to	Pol rule
• Decide	corresponds to	cut rule

□

COROLLARY 80

If $\emptyset \parallel \phi \Rightarrow^n \text{UNSAT}$ and ϕ' represents ϕ then there is a complete proof in $\text{LK}^p(\mathcal{T})$ of $\phi' \vdash$, of size smaller than $(2\sharp(\phi) + 3)n$. ※

8.2.3 Completing the bisimulation

Now the point of having mentioned quantitative information in Theorem 79, via the notion of n -extension, is to motivate the idea that performing proof-search directly in $\text{LK}^p(\mathcal{T})$ is in essence not less efficient than running DPLL($\mathcal{T}$): we have a linear bound in the length of the DPLL($\mathcal{T}$) run (and the proportionality ratio is itself an affine function of the size of the original problem).

We also need to make sure that this final proof-tree is indeed found as efficiently as running DPLL($\mathcal{T}$), which can be done by identifying, in $\text{LK}^p(\mathcal{T})$, a (complete) search space that is isomorphic to (and hence no wider than) that of DPLL($\mathcal{T}$). We analyse for this a proof-search strategy, in $\text{LK}^p(\mathcal{T})$, that exactly captures the proof-extensions that we have used in the simulation of DPLL($\mathcal{T}$), i.e. the proof of Theorem 79:

DEFINITION 126 (DPLL($\mathcal{T}$)-extensions)

An incomplete proof tree π_2 is a *DPLL($\mathcal{T}$)-extension* of an incomplete proof tree π_1 if

1. it extends π_1 by replacing its leftmost open leaf with an incomplete proof-tree of one of the forms:

$$\frac{\dots}{\Gamma, A^\perp \vdash^{\mathcal{P}} [A]} \begin{matrix} (b) \\ (a) \end{matrix} \quad \frac{\Gamma \vdash^{\mathcal{P}} l \quad \Gamma \vdash^{\mathcal{P}} l^\perp}{\Gamma \vdash^{\mathcal{P}}} l \in \Gamma$$

$$\frac{\Gamma \vdash^{\mathcal{P}, l}}{\Gamma \vdash^{\mathcal{P}}} (c) \quad \frac{}{\Gamma \vdash^{\mathcal{P}}} \Gamma_{\text{lit}} \models_{\mathcal{T}}$$

where

- (a) A is a (positive) conjunction of literals that are all in $\mathcal{P}$ except maybe one that is $\mathcal{P}$ -unpolarised

- (b) the only instances of (Pol) in the above proof are of the form $\frac{\Gamma \vdash^{\mathcal{P}, l^\perp} l}{\Gamma \vdash^{\mathcal{P}} l}$

- (c) $l \in \Gamma$ with $\Gamma_{\text{lit}}, l^\perp \models_{\mathcal{T}}$

2. any incomplete proof-tree satisfying point 1. and extended by π_2 is π_2 itself.

※

Given a DPLL($\mathcal{T}$)-extension, we can now identify a DPLL($\mathcal{T}$) transition that the extension simulates, in the sense of Theorem 79:

THEOREM 81 (Simulation of the strategy back into DPLL($\mathcal{T}$))

If π_2 is a DPLL($\mathcal{T}$)-extension of π_1 , and π_1 corresponds to $\Delta \parallel \phi$, then there is a (unique) DPLL($\mathcal{T}$) transition $\Delta \parallel \phi \Rightarrow \mathcal{S}_2$ such that π_2 corresponds to $\mathcal{S}_2$. ※

Proof: See [FGLM13, Far13]. □

If a complete proof-tree of $\text{LK}^{\mathcal{P}}(\mathcal{T})$, whose conclusion is an SMT-problem,¹⁰ systematically uses the rules in the way described by the above shapes, then it is the image of a DPLL($\mathcal{T}$) run.

While it could be envisaged to simulate DPLL($\mathcal{T}$) in a Gentzen-style sequent calculus (with a variant of Theorem 79), the above definition and theorem reveal the advantage of using a focused sequent calculus for polarised logic: Definition 126 presents¹¹ different ways of *starting* the extension of an open branch (whose leaf sequent is developed), each one of them corresponding to a specific DPLL($\mathcal{T}$) transition; then *focussing* takes care of the following steps of the extension so that, when hitting developed sequents again, the exact simulation of the DPLL($\mathcal{T}$) transition has been performed.

In order for proof-search mechanisms to exactly match DPLL($\mathcal{T}$) transitions, focussing therefore provides the right level of granularity and (together with an appropriate management of polarities) the right level of determinism.

¹⁰i.e. it corresponds to an initial state of DPLL($\mathcal{T}$)

¹¹mostly by specifying the management of polarities

COROLLARY 82 (Bisimulation) The correspondance relation (see Definition 125) between incomplete proof trees and $DPLL(\mathcal{T})$ states is a bisimulation for the transition system defined on incomplete proof-trees of $LK^p(\mathcal{T})$ by the strategy of $DPLL(\mathcal{T})$ -extensions and on states by $DPLL(\mathcal{T})$. ✱

8.2.4 More advanced features

Finally, obtaining this tight result is the reason why we identified the *elementary* $DPLL(\mathcal{T})$ system, a restriction of the Abstract DPLL Modulo Theories system of [NOT06]:

Modern SMT-solvers feature some mechanisms that are not part of our (logically complete) *elementary* $DPLL(\mathcal{T})$ system but increase efficiency, such as *backjumping* and *lemma learning* (cf. rules $\mathcal{T}$ -Backjump, $\mathcal{T}$ -Learn in [NOT06]).

It is possible to simulate those rules in $LK^p(\mathcal{T})$ by using general cuts, by extending with identical steps several open branches of incomplete proof-trees, and possibly by using explicit weakenings (depending on whether we adapt the correspondence between $DPLL(\mathcal{T})$ states and incomplete proof-trees). Again, the details of this can be found in [FGLM13, Far13].

However, with such “parallel extensions” of incomplete proof-trees, it is not clear how to count the *sizes* of proofs and extensions in a meaningful way, so the quantitative aspects of Theorem 79 and Corollary 80 are compromised; neither is it clear which criterion on proof-trees (and on how to extend them) identifies the proof-construction strategy that is the exact image of a $DPLL(\mathcal{T})$ procedure featuring those advanced mechanisms. In other words, it is not clear how to obtain such a tight correspondence.

Nonetheless, understanding backjumping and lemma learning in terms of “parallel extensions” of incomplete proof-trees, gives some concrete leads on how to integrate these features to a root-first proof-search procedure as described in this chapter. One of them is to use *memoisation* for the proof-search function. This is used to close, in one single step, any branch that would otherwise be closed by repeating the same steps as in a subproof that has already been found. In particular, doing this avoids repeating, several times, the proof-construction steps of a “parallel extension” corresponding to a single backjump.

Memoisation is also a way of performing clause-learning: a learnt clause C is a clause for which we know that $\phi \models_{\mathcal{T}} C$, and that is made available for **Fail**, **Backtrack** or **Propagate**. Such a clause corresponds to a key $\phi', C^{\perp} \vdash$ of the memoisation table, with its proof as value. A state where C can be used for **Fail** or **Backtrack** is necessarily a sequent weakening $\phi', C^{\perp} \vdash$ with extra formulae or literals, so the proof recorded in the memoisation table can be plugged there to close the current branch. When C can be used for **Propagate**, it suffices to make a cut on the missing literal: one branch will be closed by plugging-in the proof recorded in the memoisation table, while the other branch will continue the simulation.

In the next chapter, we expand on the implementation of the above results in the form of a specific *plugin* for the PSYCHE system.

8.3 Future work: Relation to abstract focussing

In this section, we give some hints as to how we could develop, in the abstract LAF system, the methodology of simulating $DPLL(\mathcal{T})$ as bottom-up proof-search in a focussed sequent calculus. We already know that we can capture LKF as the LAF instance LAF_{K1} . Since we

$$\begin{array}{c}
\frac{}{\bullet \Vdash \bullet : (\top^+, \mathcal{P})} \quad \frac{}{\sim(N^\perp, \mathcal{P}) \Vdash _ : (N, \mathcal{P})} \quad \frac{N \text{ is } \mathcal{P}\text{-negative}}{a \Vdash _ : (a, \mathcal{P})} \quad \frac{}{a \in \mathcal{P}} \\
\\
\frac{\Delta_1 \Vdash p_1 : (A_1, \mathcal{P}) \quad \Delta_2 \Vdash p_2 : (A_2, \mathcal{P})}{\Delta_1, \Delta_2 \Vdash (p_1, p_2) : (A_1 \wedge^+ A_2, \mathcal{P})} \quad \frac{\Delta \Vdash p : (A_i, \mathcal{P})}{\Delta \Vdash \text{inj}_i(p) : (A_1 \vee^+ A_2, \mathcal{P})}
\end{array}$$

Figure 37: Decomposition relation for LAF_{K1p}

have slightly modified LKF for the purpose of capturing $\text{DPLL}(\mathcal{T})$, it is natural to ask whether that fits the LAF framework as well.

8.3.1 On-the-fly polarisation

The first difference between LKF and $\text{LK}^p(\mathcal{T})$, is that we have on-the-fly polarisation of atoms. This is a feature that can easily be integrated as another LAF instance:

DEFINITION 127 (The LAF instance for on-the-fly polarisation)

The definition of the instance LAF_{K1p} is the same as that of LAF_{K1} , except that

- molecules are now pairs $(A, \mathcal{P})$ made of a formula A and a polarisation set $\mathcal{P}$ such that A is $\mathcal{P}$ -positive;
- atoms are literals.

The decomposition relation is defined in Fig. 37 (which adapts Fig. 24).

Typing contexts are defined similarly to Definition 69, but with the extra information about polarities:

A typing context is given by $(\Gamma^+, \Gamma^-, \mathcal{P})$, where Γ^+ is a list of literals and Γ^- is a list of formulae;

$(\Gamma^+, \Gamma^-, \mathcal{P})[n^+]$ is the $(n+1)^{\text{th}}$ element of Γ^+

$(\Gamma^+, \Gamma^-, \mathcal{P})[n^-]$ is $(A, \mathcal{P})$, where A is the $(n+1)^{\text{th}}$ element of Γ^- .

Context extension updates $\mathcal{P}$ just as in rule **Store** of $\text{LK}^p(\mathcal{T})$, so that (for instance)

$$(\Gamma^+, \Gamma^-, \mathcal{P}); a = ((a :: \Gamma^+), \Gamma^-, (\mathcal{P}; a))$$

✱

That instance being defined, it is however not completely clear how to integrate to LAF_{K1p} the two admissible rules of $\text{LK}^p(\mathcal{T})$ (at least in their current form):

$$\begin{array}{c}
\text{(Pol)} \quad \frac{\Gamma \vdash^{\mathcal{P}, l}}{\Gamma \vdash^{\mathcal{P}}} \quad l \in \Gamma \text{ and } \text{lit}_{\mathcal{P}}(\Gamma), l^\perp \models_{\mathcal{T}} \quad \text{(cut)} \quad \frac{\Gamma \vdash^{\mathcal{P}} l \quad \Gamma \vdash^{\mathcal{P}} l^\perp}{\Gamma \vdash^{\mathcal{P}}} l \in \Gamma
\end{array}$$

and how to use them in a bottom-up proof-search procedure based on LAF_{K1p} . This is future work.

8.3.2 Extending LAF to $\text{LAF}(\mathcal{T})$

The second difference between LKF and $\text{LK}^p(\mathcal{T})$, is of course the theory $\mathcal{T}$ and its decision procedure, used in rules Init_1 and Init_2 .

Notice that LAF can accommodate a weak form of “modulo theory”, at least according to the definition of Chapter 5: The equality on atoms is a parameter that we can use to identify for instance $a(3+4)$ with $a(7)$, in particular in the rule typing positive labels

$$\frac{\Gamma[x^+] \equiv (a, \mathbf{r})}{\Gamma \vdash x^+ : (a, \mathbf{r})} \text{init}$$

However in **LAF**, we cannot close a branch in one step by involving several atoms of Γ , as we do for instance in $\text{LK}^p(\mathcal{T})$ when we call e.g. a simplex algorithm to check the consistency of (the positive literals of) Γ (which in **LAF** would be Γ^+).

For this we would need to extend **LAF** with a decision procedure. We could think of doing it in the following way:

- replace the notion of positive label by a notion of *focussed justification*, and abstract away the part Γ^+ of a typing context Γ , which is no longer a function from positive labels to instantiated atoms but an abstract data structure called a *positive typing context*;
- replace the notion of equality between instantiated atoms by a typing relation of the form $\Gamma^+ \models [s^+ : (a, \mathbf{r})]$, where Γ^+ is a positive typing context, and s^+ is a positive justification.
- add a notion of *justification* and a typing relation of the form $\Gamma^+ \models s$, where Γ^+ is a positive typing context, and s is a justification.

We would then get:

DEFINITION 128 (Proof-Terms) Let $\mathbb{J}$ be a set of elements called *justifications*, and $\mathbb{J}^+$ be a set of elements called *focussed justifications*.

Proof-terms are defined by the following syntax:

Positive terms	Terms ⁺	$t^+ ::= pd$
Decomposition terms	Terms ^d	$d ::= s^+ \mid f \mid \bullet \mid d_1, d_2 \mid r.d$
Commands	Terms	$c ::= \{s\} \mid \langle x^- \mid t^+ \rangle \mid \langle f \mid t^+ \rangle$

where p ranges over patterns, s ranges over justifications, s^+ ranges over focused justifications, x^- ranges over Lab_- , f ranges over functions from patterns to commands. $\ast$

And now we can give the **LAF** system parameterised by a “theory” given by the pair of typing relations $_ \models [_ : _]$ and $_ \models _$, which we may call $\mathcal{T}$, and which plays the same role as the semantical inconsistency predicate in $\text{LK}^p(\mathcal{T})$.

DEFINITION 129 (LAF($\mathcal{T}$)) Let Co^+ be the family of positive typing contexts.

Assume we are given a pair $\mathcal{T}$ of two relations

$(_ \models [_ : _]) : (\text{Co}^+ \times \mathbb{J}^+ \times \mathbb{A}_\downarrow)$ and $(_ \models _) : (\text{Co}^+ \times \mathbb{J})$.

We define in Fig. 38 the derivability of three typing judgements

- $(_ \vdash [_ : _]) : (\text{Co} \times \text{Terms}^+ \times \mathbb{M}_\downarrow)$
- $(_ \vdash _ : _) : (\text{Co} \times \text{Terms}^d \times \mathbb{D}_\downarrow)$
- $(_ \vdash _) : (\text{Co} \times \text{Terms})$

$\ast$

The empty theory could be recovered by having positive labels, having positive typing contexts as maps from positive labels to instantiated atoms, having focussed justifications be exactly positive labels, and by setting setting $\Gamma^+ \models [s^+ : (a, \mathbf{r})]$ if and only if $\Gamma^+(s^+) \equiv (a, \mathbf{r})$ and never having $\Gamma^+ \models s$.

Linear arithmetic could be defined by a relation $\Gamma^+ \models s$ that would check the consistency of the instantiated atoms in Γ^+ , and a relation $\Gamma^+ \models [s^+ : (a, \mathbf{r})]$ that would check the consistency of the instantiated atoms in Γ^+ together with $(a^\perp, \mathbf{r})$.

$$\frac{\Delta \Vdash p:M \quad \Gamma \vdash d:(\Delta, \mathbf{r})}{\Gamma \vdash [pd:(M, \mathbf{r})]} \text{sync}$$

$$\frac{}{\Gamma \vdash \bullet:(\bullet, \mathbf{r})} \quad \frac{\Gamma \vdash d_1:(\Delta_1, \mathbf{r}) \quad \Gamma \vdash d_2:(\Delta_2, \mathbf{r})}{\Gamma \vdash d_1, d_2:((\Delta_1, \Delta_2), \mathbf{r})} \quad \frac{\Gamma^e \Vdash r':s \quad \Gamma \vdash d:(\Delta, r':\mathbf{r})}{\Gamma \vdash r'.d:s.(\Delta, \mathbf{r})}$$

$$\frac{\Gamma^+ \models [s^+:(a, \mathbf{r})]}{\Gamma \vdash s^+:(a, \mathbf{r})} \text{init}_1 \quad \frac{\forall p, \forall \Delta, \quad \Delta \Vdash p:M \Rightarrow \Gamma; (\Delta, \mathbf{r}) \vdash f(p)}{\Gamma \vdash f:(\sim M, \mathbf{r})} \text{async}$$

$$\frac{\Gamma^+ \models s}{\Gamma \vdash \{s\}} \text{init}_2 \quad \frac{\Gamma \vdash [t^+:\Gamma[x^-]]}{\Gamma \vdash \langle x^- \mid t^+ \rangle} \text{Select} \quad \frac{\Gamma \vdash f:(\sim M, \mathbf{r}) \quad \Gamma \vdash [t^+:(M, \mathbf{r})]}{\Gamma \vdash \langle f \mid t^+ \rangle} \text{cut}$$

Figure 38: LAF($\mathcal{T}$)

For congruence closure we could have the same approach, or we could give a special role to $(a, \mathbf{r})$ in defining when $\Gamma^+ \models [s^+:(a, \mathbf{r})]$ holds.

The abstract focussing system could be seen as a *functor* (in the programming language sense) $\mathcal{T} \mapsto \text{LAF}(\mathcal{T})$ that takes a pair of typing relations (focussed, unfocussed) and returns a new pair of typing relations (focussed, unfocussed). In that view, the functor could be composed with others, and iterated. We conjecture that second-order logic or higher-order logic could be captured by the fixpoint of this functor, together with one that can convert an atom into a molecule, etc.

In every theory, the justifications could be dummy objects, if we do not have proof objects to produce when running the decision procedure. Or they could be as informative as one would like; in particular, it would be useful if s (resp. s^+) could at least indicate which part of Γ^+ is actually used to derive $\Gamma^+ \models [s^+:(a, \mathbf{r})]$ (resp. $\Gamma^+ \models s$). This could be done via a notion of free labels, so that we can apply the same methodology as that of Section 7.3 to re-use proofs in different contexts.

The formal study of such a LAF system, together with the adaptation of its realisability models, is left for future work.

Chapter 9

The PSYCHE system

Contents

9.1	Motivation	168
9.2	Overview and general architecture	170
9.3	PSYCHE's Kernel	171
9.4	Plugins	172
9.4.1	Specifications and implemented instances	172
9.4.2	Memoisation and lemma learning	173
9.5	Decision procedures	175
	Conclusion: Testing and perspectives	176

In this chapter, we describe PSYCHE [Psy], a system programmed in OCaml that implements, among other things, the ideas developed in the previous chapter(s). In particular, it uses polarities and focussing as a way to organise proof-search.

PSYCHE is a highly modular proof-search engine designed as a platform for either interactive or automated theorem proving, and the acronym stands for the *Proof-Search factorY for Collaborative Heuristics*. By platform, we mean that its architecture is organised around a *kernel* that interacts with *plugins* to be programmed via a specific API. The goal of this architecture is to allow the implementation of various theorem proving techniques while guaranteeing correctness of the output: whether an input formula is provable or not provable. As a platform, it can also be used to implement the collaboration of various techniques which, once programmed as plugins, share the same notion of *proof-search state*.

The aim is therefore to provide a high level of confidence about the output of the theorem proving process, no matter how programmers have implemented their plugins, which is done by adopting and somewhat transforming the LCF architecture [GMW79].

Finally, PSYCHE features the ability to call *decision procedures* such as those used in Sat-Modulo-Theories provers. We therefore illustrate PSYCHE by using it for SMT-solving.

In brief:

- The kernel is based on a proof-search engine *à la* Prolog, offering an API to perform incremental and goal-directed constructions of proof-trees in a focussed Sequent Calculus, which can be seen as a *tableaux method* [DGHP99].
- PSYCHE can produce proof objects.

- Plugins can be programmed to drive the kernel, using its API, through the search space towards an answer *provable* or *not provable*; correctness of the answer only relies on the kernel via the use of a private type for answers (similar to LCF’s *theorem* type).
- Plugins can be interactive.
- PSYCHE offers a *memoisation* feature to help programming efficient plugins.
- The kernel is parameterised by a procedure deciding the consistency of collections of literals with respect to a background theory, just as in SAT-modulo-theories (SMT) solvers.

The current version 2.0 of PSYCHE is distributed

- with a kernel designed for first-order logic modulo a theory $\mathcal{T}$;
- with a plugin whose behaviour on quantifier-free problems is $\text{DPLL}(\mathcal{T})$, using *watched literals* to propagate literals or close branches, and PSYCHE’s memoisation feature to learn lemmas;
- with decision procedures for: pure propositional logic (for SAT-solving), pure first-order logic, quantifier-free Linear Rational Arithmetic (LRA), and Congruence Closure;
- with a DIMACS parser and an SMTLib2 parser¹;
- as a program of about 6700 lines of OCaml 4.00 (the kernel itself is only 800 lines), using hash-consing and Patricia tries for efficiency reasons.

PSYCHE does not claim to be a better SAT- or SMT-solver or first-order theorem prover than any existing one: for instance the heuristics for applying $\text{DPLL}(\mathcal{T})$ rules in the aforementioned plugin are still basic, and so is the decision procedure for LRA (it is not incremental). What we offer here is a platform and its modularity: anyone with better (or different) heuristics and decision procedures can simply write them as OCaml modules of our predefined module types, and PSYCHE will seamlessly run with them, keeping the same LCF-style guarantees.

In Section 9.1, we give more motivation for the development of PSYCHE. In Section 9.2, we describe the general architecture of the system, in particular we explain how the guarantee of correctness is enforced, using the kernel API and a private type for answers. In Section 9.3, we briefly review how the kernel works, connecting to the theory described in the previous chapters. Section 9.4 then describes what the specifications required of a plugin, and the way our distributed plugin simulates $\text{DPLL}(\mathcal{T})$ according to the results from Chapter 8. Section 9.5 describes the specifications of decision procedures and parsers, while Section 9.6 concludes with some tests and perspectives.

9.1 Motivation

PSYCHE’s architecture is designed for the ambition of allowing various theorem proving techniques (generic or problem-specific) to collaborate on a common platform, whilst giving high confidence in the answers produced.

Interfacing the numerous techniques and tools available for theorem proving is legitimately receiving a lot of attention: Automated Theorem Provers, SAT-solvers, SMT-solvers, Proof assistants, etc. While trust is already an issue even for a single tool running on its own, it becomes even more of an issue when different tools interact. *Proof-checking* is one way of addressing this, i.e. being very permissive in the algorithms used for theorem proving,

¹The latter is taken without modification from the Alt-Ergo SMT prover.

as long as they output some proof objects that can be checked. Another way is the LCF-style [GMW79], where only a small kernel of primitives needs to be trusted, and anything smarter (e.g. the interaction between sophisticated techniques) boils down to calls to the kernel’s primitives.

In the context of proof-checking (such as in Coq [Coq]), a natural way to interact with different (already implemented) techniques, is the black box approach, where an external tool is called and its output is converted back into a proof that can be checked by the system [AFG⁺11, BCP11]. It is somewhat more surprising that, despite the highly programmable possibilities of the LCF architecture (from which the ML languages come), the most successful integration of automated reasoning techniques in an LCF-based proof assistant such as Isabelle [NPW02, Isa] seems to also use variants of the black box approach (as very impressively demonstrated by Sledgehammer) [Web11, PB12, BBP11].

PSYCHE aims at producing answers that are correct by construction, not having to rely on proof-checking; it therefore adopts the LCF philosophy (although it can produce proof objects), also because having a simple trusted kernel is a convenient starting point for different techniques to collaborate. But the goal here is to open the black boxes and program their algorithms directly with calls to the kernel’s API, as plugins for PSYCHE.

Such a deeper level of integration opens up the perspective of interleaving the use of different techniques: An external tool requires an input problem that it can entirely treat; but implementing the *steps* of its algorithm as small progressions in the search-space covered by the main system, allows more possibilities, such as running the technique *up-to-a-point*, where a switch to another technique may be appropriate (e.g. depending on newly generated goals).

The challenge is for the kernel to offer an appropriate API of proof-search or proof-construction primitives, to allow the efficient implementation of theorem proving techniques as plugins. Most LCF-style systems offer primitives corresponding to the inference rules of Natural deduction, or a Hilbert-style system. This is a very fine-grained level, that leaves most (if not all) of the work to the plugin; requiring it to use the kernel’s primitives is less of an aid and more of a constraint: it does ensure that, in case the output is provable, a proof has been constructed (at least theoretically), but it is a computational overhead for the plugin’s work.

PSYCHE makes the choice of a bigger grain, and leaves to the kernel some real proof-search computation, but where no decision needs to be made. For this we use the focussed sequent calculus $LK^p(\mathcal{T})$ [FGL13, FGLM13], whose quantifier-free version has been presented in Chapter 8. Not only can polarities and focussing be used to describe effective proof-search strategies in Sequent Calculus (narrowing the search-space offered by Gentzen’s original rules), but in our case, they also specify a sensible division of labour between PSYCHE’s kernel and PSYCHE’s plugins, re-designing the standard LCF-style API.

This new design makes PSYCHE guarantee the correctness of both types of answers: *provable* or *not provable*, while the traditional LCF style only guarantees the correctness of answers of the form *provable*.

9.2 Overview and general architecture

The kernel contains the mechanisms for exploring the proof-search space, taking into account branching and backtracking. It has no *a priori* regarding the order in which branches are explored, and this lack of intelligence makes its code rather short. If it reaches a proof, then that proof is correct by construction, and if the entire search space is explored and no proof is found, then the kernel correctly outputs that no proof exists.

The plugins then drive the kernel by specifying in which order the branches of the search space should be explored and to which depth, something that is expected to depend on the kind of problem that is being treated. The quality of the plugin is how fast it drives the kernel towards a answer *provable* or *not provable*.

This already departs from the traditional LCF-style in that some actual proof-search computation is performed in the kernel, not just atomic steps of proof-construction:

In traditional LCF, each inference rule of the logic

$$\frac{\text{prem}_1 \quad \dots \quad \text{prem}_n}{\text{conc}} \text{ name}$$

gives rise to a primitive of the kernel’s API, whose type declares n arguments:

name: `thm -> ... -> thm -> thm`

In PSYCHE’s kernel, such an inference rule is “wrapped” in the kernel’s unique API primitive:

machine: `statement -> output`

such that `search(conc)` will trigger the recursive calls `search(prem_1), ..., search(prem_n)`, as bottom-up proof-search should do.

PSYCHE’s general architecture is illustrated by its main top-level call (slightly reworded for clarity):

`Plugin.solve(Kernel.machine(Parser.parse input))`

PSYCHE has a collection of parsers (currently one for DIMACS and one for SMTLib2) and calls the appropriate one on PSYCHE’s input. The resulting abstract syntax tree is fed to the kernel’s `machine` function that will initiate the search.² This produces a value of type `output` that is given to the plugin to work with, and the plugin must solve the problem by outputting an answer *provable* or *not provable*.

This could give the impression that the plugin performs computation after the kernel has finished his, but this is not quite true, as illustrated by the nature of type `output`:

type output = `Jackpot of answer | InsertCoin of coin -> output`

which describes the kernel as a *slot machine*: when it is run, it outputs

- either a definitive answer *provable* or *not provable*
- or an intermediate output that represents unfinished computation: in order for computation to continue, the plugin needs to “insert another coin in the slot machine”; depending on the kind of coin inserted, proof-search will resume in a certain way.

To summarise, the kernel performs proof-search as long as there is no decision to be made (on which backtrack may later be needed), and when it hits such a point, it stops and asks for another coin to indicate how to proceed next. The plugin drives the kernel in the exploration

²In fact, the kernel module is created with the choice of a background theory that is either guessed from the input or specified by the user on the command line.

of the proof-search space by inserting carefully chosen coins, hoping that one day the machine will stop with the “jackpot”: a value of the form `Jackpot(...)`.

Now while this architecture somewhat departs from LCF, it does share with it the distrust of anything outside the kernel: when concerned with the soundness of the answer (whichever it be), the plugin is here considered as an adversary, so PSYCHE defines the type `answer` as a private type that only the kernel can inhabit (just like the `thm` type of LCF). PSYCHE's type

`answer = private Provable of statement*proof | NotProvable of statement`

can be read by the plugin and the top-level if need be, but cannot be inhabited by them. That way, a plugin cannot cheat about PSYCHE's answer: the worst it can do is of course to crash PSYCHE's runs or diverge. In PSYCHE as in traditional LCF, inhabitation of the abstract type (in case of PSYCHE, with a value of the form `Provable(...)`) explicitly or implicitly constructs a proof of the statement. But contrary to LCF, PSYCHE also gives guarantees when the output is *not provable*: it can only occur when the kernel has entirely explored the search-space unsuccessfully.

Such a use of typing prompted for an ML-language to implement PSYCHE, and we chose OCaml (4.0).

9.3 PSYCHE's Kernel

As described above, the kernel's API has the slot machine as its only primitive, controlled by the *coins* that are inserted in it. In order for efficient plugins to be conveniently programmed, the kernel's primitive needs to accept a rather expressive range of coins that can specify a smart exploration of the search-space. This depends on the inference system that is used in the kernel for the incremental and bottom-up construction of proof-trees, and on identifying the inference rules that the kernel will perform automatically from those that will pause computation and prompt the plugin for new directions.

This is where focussed sequent calculi for polarised logic(s) come in. Focussing is what we use for the division of labour between PSYCHE's kernel and PSYCHE's plugins:

The kernel applies the asynchronous steps automatically without any instruction from the plugin, and then stops and asks for another coin describing the next synchronous phase, where smart choices may have to be made (starting with the choice of the positive formula to work on).

An important consequence of this division of labour is that **every kernel call terminates**, because the length of each phase is bounded by the size of the formula(e) being decomposed. Therefore, infinite proof-search has to go through an infinite interaction between the kernel and the plugin (unless the plugin itself loops before inserting the next coin).

The choice of polarities on connectives and literals affects the kernel-plugin interaction. For instance the polarity of $\vee$ will determine whether it will be decomposed automatically by the kernel (second rule, asynchronous) or with a smart choice by the plugin (first rule, synchronous):

$$\frac{\Gamma \vdash A_i}{\Gamma \vdash A_1 \vee^+ A_2} \quad \frac{\Gamma \vdash A_1, A_2, \Gamma'}{\Gamma \vdash A_1 \vee^- A_2, \Gamma'}$$

The polarity of literal being also crucial, PSYCHE offers the plugin the possibility to polarise literals *on-the-fly*, during the search (which is very useful for the plugins we implemented).

In PSYCHE 2.0, the kernel implements the sequent calculus $LK^p(\mathcal{T})$ [FGLM13, Far13], whose quantifier-free version was presented in Chapter 8. But PSYCHE does implement the full system with quantifiers, with specific mechanisms dealing with eigenvariables (introduced when proving a universal formula) and meta-variables (introduced when proving an existential formula).

The different coins that the plugin can insert thus correspond to the smart application of the non-asynchronous inference rules of $LK^p(\mathcal{T})$: a formula to select, a side to choose when decomposing $\vee^+$, a literal to polarise in a certain way, a cut to be made ($LK^p(\mathcal{T})$ admits cuts), or a consistency check of the current sequent with the given background theory (a global parameter of the kernel).

Finally, the plugin can also instruct the kernel to move in the search-space: when it gets tired of investigating the current branch, it can abandon it temporarily and explore the next success/failure branch to the left/right.

The code of the kernel is rather small (around 800 lines) and purely functional. Continuation-Passing-Style (CPS) is used to minimise the use of the stack and provide a natural way to represent the progression of the kernel within the search space: the API function

```
machine: statement->output
```

actually wraps a real (tail-)recursive function

```
search: statement->(output->'a)->'a
```

with the identity continuation. Continuations are heavily used for branching and backtracking (e.g. when `search` applies a rule with several premises, it makes a recursive call on one of the branches and stacks up the others in the continuation that is passed; when the plugin chooses to explore one branch, the kernel records in a similar way the other branches that are not being explored yet -forcing in the end the entire exploration of the search-space), and naturally provide the values implementing a slot machine waiting for its coin.

9.4 Plugins

9.4.1 Specifications and implemented instances

A plugin is any OCaml module implementing the following identified module type (bearing in mind that `answer` is for the plugin a private type that it cannot inhabit by itself):

```
module type PluginType = sig
  ...
  solve: output->answer
end
```

However, it is likely that the sophisticated strategies/heuristics that the plugin is meant to implement rely on some clever choice of data-structures for formulae, sets of formulae, sets of literals. So the plugin and the kernel have to agree on those three data-structures that are communicated both ways during the interaction. In PSYCHE 2.0, the kernel provides the data-structure to represent formulae, but the plugin can embark in the data-structure the information that it needs to treat them efficiently. The data-structures implementing sets of

formulae and sets of literals, on the other hand, are parameters of the kernel, and the plugin provides them.³

We first tested PSYCHE’s architecture with a basic plugin **Naive**, which

- implements sets (of formulae, literals) with OCaml’s lists;
- inserts the first available coin in the slot machine, whenever asked.

This works fine for small tautologies, printable on a screen.

More recently, Jean-Marc Notin provided a module for interactive theorem proving, via a command-line interface: it still implements sets using OCaml’s lists, but every time a coin needs to be inserted in the machine, the interface prompts the user for the coin to insert.

A more ambitious aim for automated reasoning was to capture in PSYCHE some propositional SAT and SAT-Modulo-Theories solving techniques, making $\text{DPLL}(\mathcal{T})$ technology available in a generic tableau-like / Prolog-like / goal-directed proof-search framework like PSYCHE.

For this we implemented the simulation of $\text{DPLL}(\mathcal{T})$, expressed rather canonically as a transition system [NOT06], as a simple bottom-up proof-construction mechanism in $\text{LK}^p(\mathcal{T})$, as described in Chapter 8. More practically, every rule of $\text{DPLL}(\mathcal{T})$ can be seen as the insertion of a particular coin in PSYCHE’s slot machine.

We implemented this as two different plugins for PSYCHE: `DPLL_Pat` and `DPLL_WL`. These remain toy plugins, because, although it is now clear, from Chapter 8, how to perform each rule of $\text{DPLL}(\mathcal{T})$ in PSYCHE, we still have to decide *which* rule to apply and *when*. So the two plugins

- embark, in the kernel’s representation of formulae, a flat representation of them as sets of literals when the formulae happen to be clauses;
- implement sets (of formulae/clauses, literals) using Patricia tries;
- implement a basic strategy to apply $\text{DPLL}(\mathcal{T})$ rules; in the case of propositional logic: apply **Fail** or **Backtrack** if possible, if not try **Unit Propagate**, if not do **Decide** on some random literal.

The two plugins differ in the way they look up for the applicability of **Fail** / **Backtrack** / **Unit Propagate**: `DPLL_Pat` looks it up using the Patricia tries implementing sets of clauses, while `DPLL_WL` looks it up using the technique of *watched literals* [MMZ⁺01] (keeping a small watching table in the plugin). This technique was originally implemented in PSYCHE by student Matthieu Vegreville, and the plugin seems on average 1.5 faster than that using Patricia tries.

9.4.2 Memoisation and lemma learning

Such plugins would not be efficient at all if no backjumping and clause learning was done while performing $\text{DPLL}(\mathcal{T})$. In [FGLM13] we also show how to do this using general cuts, and either accept to extend several open branches of an open proof-tree with identical steps or depart from the bottom-up proof construction paradigm that we have used so far. We

³This is admittedly a security problem, since a bug in the plugin’s data-structure for sets could affect the way a sequent is transformed by the kernel when it applies an inference rule bottom-up. The next version of PSYCHE will adopt a double representation of sets (one for the kernel, one for the plugin) to completely avoid the kernel relying on plugin code.

opted for a generic mechanism to avoid re-doing, for some open branch, the same steps as those used in a previously completed branch: *memoisation*. In Chapter 8 we explained how the use of memoisation emulates the use of a learnt clause for **Fail**, **Unit Propagate**, etc.

Indeed, nothing prevents a plugin from recording the sub-trees completed by the kernel, and proposing them later for another branch where the same proof-tree is relevant. PSYCHE 2.0 therefore offers a memoisation module, to be used by plugins to record values of (the abstract) type **answer**. And the kernel's slot machine accepts from the plugin, as a special coin carrying such a value, "here is an already found answer that also applies to the current goal". The kernel accepts the value as closing the current branch (one way or another) **without any proof-checking** (since the abstract type ensures the value came as an earlier output of the kernel); it only checks that the value applies to the current goal.

The memoisation table is filled-in by clause-learning: our plugin adds an entry whenever it builds a complete proof of some sequent $\Delta \vdash$ and no previous entry $\Delta' \vdash$ exists with $\Delta' \subseteq \Delta$, or whenever it concludes that some sequent $\Delta \vdash$ is not provable and no previous entry $\Delta' \vdash$ exists with $\Delta \subseteq \Delta'$.

Now for a memoised answer *Provable* to be reusable as often as possible, a pre-processing step is applied to a proof-tree before it enters the table: it is pruned from every formula that is not used in the proof. This is easy to do for the complete proofs of $LK^p(\mathcal{T})$ (eager weakening are applied a posteriori by inspection of the inductive structure). PSYCHE's kernel actually performs the pruning on-the-fly whenever an inference is added to complete proofs, so that, whenever it outputs **Jackpot(sequent,proof)**, the sequent is already pruned.

Since proof-completion can be seen as finding a *conflict* (a situation where the current partial model contradicts the set of clauses), pruning by eager weakening is a *conflict analysis* process naturally provided by structural proof theory:

Conflict analysis is a process used in SAT- and SMT-solving aims at identifying, in a situation of **Fail**, **Backtrack**, etc, which literals of the current model are sufficient to contradict, when taken together, the set of clauses; the disjunction of their negations forms a new clause that can be learnt and re-used later. Techniques to compute this can be based on *graph analysis*; the kind of conflict analysis performed by the pruning mechanism of PSYCHE turns out to be a particular form the graph analysis mechanism.

Of course, just as in SMT-solving, the efficiency of conflict analysis relies on the efficiency of the decision procedure in providing a small inconsistent subset whenever it decides that a set of literals is inconsistent.

Another feature sometimes used in SMT-solving, in conjunction with clause learning, is the use of *restarts*: at some point of the $DPLL(\mathcal{T})$ run, computation resumes with the empty model:

$$\Delta \parallel \phi \Rightarrow \emptyset \parallel \phi$$

This is only useful if the current set of clauses ϕ is different from the original one, i.e. some clauses have been learnt: in that case, restarting from the empty model but with all the learnt clauses, might be faster than closing all the branches that have been opened (corresponding to the decision literals in Δ).

In PSYCHE, restarts can be done the same way: since the plugin is in charge of computation, it can record the first output that the kernel produced, and later come back to it: the side effect that makes it different from the first run is that the memoisation table has been filled with valuable information; this information may allow the search to find a proof

more quickly than by closing all the open branches of the current incomplete proof-tree. This has been implemented in PSYCHE by students Zelda Mariet and Clément Pit-Claudel, with convincing experimental results.

9.5 Decision procedures

Decision procedures and parsers integrate PSYCHE's code the same way as plugins: we offer a module type for decision procedures and one for parsers. Someone with a decision procedure or a parser can implement a module of the corresponding type and run PSYCHE with it.

In the case of decision procedures for quantifier-free problems (i.e. with ground literals), the output of a decision procedure for the background theory $\mathcal{T}$ should be able to decide whether a conjunction of literals is consistent with $\mathcal{T}$ or not.

```
module type GroundDecProc = sig
  ...
  type literals
  ...
  consistency: literals set -> (literals set) option
end
```

The decision procedure provides the type of literals, so as to run efficiently, while the kernel accepts any type for literals since it will not inspect its values.

The output of the `consistency` function is not a boolean: it should be `None` if the input is a set of literals consistent with the theory $\mathcal{T}$, and `Some(s)` otherwise, with `s` being a subset of the input that is already inconsistent. Indeed, conflict analysis requires such subsets to be produced when an inconsistency is found, and the smaller the subsets, the more efficient clause learning and memoisation will be.

In the case of problems with quantifiers, the decision procedure should answer whether there is a way to instantiate meta-variables so as to make the conjunction of literals inconsistent with the theory: only in this case will the current branch be immediately closed, propagating the instantiation of meta-variables to the remaining open branches. In case such an instantiation fails another branch, we should backtrack to the current branch and propose another way of closing it. Therefore, the decision should not only be able to decide whether there is an instantiation of meta-variables that makes the literals inconsistent, but it should be able to *enumerate all* possible instantiations that make the literals inconsistent. Instead of a boolean answer, we thus expect a *stream* of solutions, each of which is a set of literals together with a working instantiation:

```
module type DecProc = sig
  ...
  type literals
  type constraints
  ...
  val consistency : literals set -> (literals set, constraints) stream
end
```

The idea of using streams of solutions is natural, and proposed in the form of *instance streams* in [Gie00] in a proof-search methodology that deliberately avoids backtracking. Although proof-search in PSYCHE does backtrack, we should investigate the connection between PSYCHE's methodology and that of [Gie00].

As evoked by the module type above, instantiations in PSYCHE are actually called *constraints*: in pure first-order logic, a constraint would simply be a (most general) unifier σ that makes two literals l_1 and l_2 of the input set such that $l_1\sigma = l_2\sigma$. It would be easy to enumerate all such constraints, by enumerating all pairs l_1 and l_2 of the input set that can be unified in the above sense: they are in finite numbers.

But for other theories we could imagine different kinds of constraints on meta-variables: for instance in Linear Rational Arithmetic we could imagine a constraint imposing that meta-variable ?X3 be in the interval $[0; \frac{3}{2}]$. Therefore, the decision procedure provides the notion of constraint that is appropriate for the background theory $\mathcal{T}$, while the type for constraints is abstract for the kernel, which will only propagate constraints from branch to branch, but not inspect them.

The exact specifications that should be met by the constraint structure so that proof-search using them is sound and complete with respect to the formalisation of the theory $\mathcal{T}$ without meta-variables, is the object of a paper being currently written with student Damien Rouhling.

Finally, PSYCHE is modular in its parsers: a *parser* is any module of a pre-defined module type, and should in particular implement a function

```
parse: string->((statement option)*(boolean option))
```

that turns a string input into a statement to be proved (or `None` if no statement was parsed in the input), and possibly an expected result *Provable/NotProvable* that the input string may indicate.

Conclusion: Testing and perspectives

PSYCHE 2.0 is run from the command-line, taking as input one or more file(s) or directory(ies), or, if none indicated, the standard input:

```
psyche [OPTION]... [FILE/DIR]...
```

Available options are:

- theory selects theory (among empty, lra, cc, first-order; default is empty)
- gplugin selects generic plugin (among naive, dpll_pat, dpll_wl; default is dpll_wl)
- latex allows latex output of proof-trees
- alphasort treats input files in alphabetical order (default is from smaller to bigger)
- examples treats theory examples instead of standard input
- nocuts disallows cuts
- fair ensures fairness between formulae for focus
- noweakenings disables conflict analysis
- nomemo disables memoisation
- restarts selects a restart strategy
- help displays this list of options

As illustrated by the options, PSYCHE can produce proofs (of $LK^p(\mathcal{T})$) and print them as inference trees in $\text{\LaTeX}$ format (but proofs can quickly get too big for $\text{\LaTeX}$).

We ran PSYCHE on instances of SAT in DIMACS format, and QF_LRA instances in SMTLib2 format and the results are available on Psyche’s website [Psy]. Psyche works well on small instances and its performance starts declining between 20Kb and 100Kb of input problem size (of course this is no appropriate measure of difficulty). This is of course very far from current SAT benchmarks, perhaps a bit less from SMTLib2 ones (our instances were download from the up-to-date library). But as we said, the current plugins and decision procedures are illustrative toys. PSYCHE is a platform where people knowing good and efficient techniques should be able to program them.

In the short-to-medium terms, we plan to

- improve the decision procedure for LRA (making it incremental, and returning smaller sets);
- implement other theories and combine them (congruence closure, Linear Integer Arithmetic, bit vectors, etc);
- improve DPLL($\mathcal{T}$) plugins to better handle non-clausal formulae;
- implement other theorem proving techniques as plugins: *analytic tableaux* are the closest to our sequent calculus, but theoretical developments have already shown that *clausal tableaux* (including *connection tableaux*) can also be done [Far13], as well as *resolution* proofs.

Finally, we can imagine using a proof assistant to prove PSYCHE’s correctness, since the kernel seems small enough (800 lines) and the plugins need not be certified.

We will develop our long-term plans for PSYCHE in the conclusion of this dissertation.

Chapter 10

Conclusion and further work

Contents

10.1 Summary of the topics covered by this dissertation	179
10.2 Further work	180

10.1 Summary of the topics covered by this dissertation

In the first part of this dissertation, we reviewed the computational interpretation of proofs in terms of Call-by-Name and Call-by-Value evaluation of programs [CH00, Sel01]; we approached realisability semantics by a systematic construction of orthogonality models [Par97, DK00, Kri01, Miq11], used for instance to prove strong normalisation results or classical witness extraction. From this the concepts of polarities and focussing naturally emerged [MM09], and a computational interpretation of focussed proofs was given in terms of pattern-matching [Zei08a, Zei08b].

In the second part of this dissertation, we developed this approach into an abstract focussed sequent calculus LAF with proof-terms, of which several focussed calculi of the literature are instances, such as LKF and LJF [LM09]. We used this framework to formally relate classical realisability with the computational interpretation of focussing as pure pattern-matching, again via the construction of orthogonality models.

In the third part of the thesis we explored a specific approach to theorem proving benefiting from the use of polarities and focussing. We described how these concepts can contribute to the description of the $\text{DPLL}(\mathcal{T})$ procedure for SMT-solving [NOT06] as a specific strategy for the bottom-up proof-search process specified by the sequent calculus. For this we extended the focussed sequent calculus LKF into $\text{LK}^p(\mathcal{T})$, equipped with the ability to polarise atoms *on-the-fly* and call a decision procedure specific to the background theory $\mathcal{T}$.

We then described the implementation of a proof-search engine called PSYCHE [Psy] whose architecture is based on a kernel that interacts with plugins to be programmed via a specific API. This allows the implementation and experimentation of various reasoning techniques (among which $\text{DPLL}(\mathcal{T})$) and heuristics, without worrying about breaking the correctness of PSYCHE's output: this is guaranteed to be correct by the architecture, which develops a new variant of the LCF style [GMW79].

10.2 Impact of this dissertation on the development of PSYCHE, and further work

In conclusion of this dissertation we proffer two main directions in which the material of this dissertation will be developed and integrated to the next releases of PSYCHE.

The first one is a rather major change in the kernel of PSYCHE: instead of implementing bottom-up proof-search in the particular focussed sequent calculus called $LK^P(\mathcal{T})$, which is specific to classical logic, PSYCHE 3.0 will implement proof-search in the abstract focussed sequent calculus **LAF** developed in Part II of this dissertation. This will allow the kernel to be decomposed into smaller components: the main module will have much fewer rules to implement than in $LK^P(\mathcal{T})$, taking advantage of the “big-step presentation” of focussing; the decomposition of formulae into smaller formulae, given by the specific relation $\Vdash$ of **LAF**, will be moved to a specific module where the inductive syntax of formulae is implemented; another module will implement typing contexts with their notion of context extension that crucially determines which logic is being implemented, etc.

The advantages of modularising the kernel in this way are numerous:

- It will allow PSYCHE to run on different instances of **LAF**, thus handling different systems and logics;
- PSYCHE will then be equipped with proof-terms, which may be used as compact representations of proofs in memory (e.g. in the memoisation table); this will also allow the extraction of programs from proofs (in different logics);
- the code will also be simpler to understand and formally prove correct, as most of the specifications describing the roles of each component have already been identified in Part II of this dissertation, with most of the theorems already formalised in Coq [GL14];

In retrospect, this next move in the development of PSYCHE is also what motivated the detailed study of **LAF**, at the cost of presenting it in a rather technical way. However, theoretical work still needs to be done before this new basis for PSYCHE’s implementation replaces the current one. Indeed, as described in details in Section 8.3, it is not clear how PSYCHE can take advantage of rules that are admissible for specific instances of **LAF** but not generically (e.g. specific forms of cuts, on-the-fly polarisation rules, etc); more importantly, in order to supersede the current implementation, **LAF** needs to be generalised into a system $LAF(\mathcal{T})$ that can call a decision procedure for a theory $\mathcal{T}$. What conditions are required of such procedures for cut-elimination to work, etc, remains to be identified.

The second direction for further development is exploiting the machinery for quantifiers that has newly been introduced with the release of PSYCHE 2.0 on 20th September 2014.

As briefly described in Section 9.5, this machinery involves meta-variables which are introduced when breaking an existential quantifier sitting on top of a formula to be proved, and eigenvariables which are introduced when breaking a universal quantifier. Dependencies between them are recorded in the proof-search, so as to avoid the production of incorrect instances for meta-variables, from which no actual proof could be re-constituted.

On the note of dependencies, Skolemisation is often described as the transformation of a formula $\forall x_1 \dots \forall x_n \exists y A$ to be *refuted* (by tableaux methods, resolution, etc...) into the formula $\forall x_1 \dots \forall x_n \left\{ \text{sk}_y(x_1, \dots, x_n) / y \right\} A$, where sk_y is a (new) *Skolem symbol* (specific to y). The

correctness of this transformation is often justified by semantical models involving the axiom of choice, and Skolemisation is often applied as a pre-processing step before refutational methods are applied. In bottom-up proof-search, Skolemisation occurs *on-the-fly* when the universal quantifier of $\exists x_1 \dots \exists x_n \forall y A^\perp$ is broken, and the skolem symbol $\mathbf{sk}_y$ is merely the eigenvariable Y introduced for y ; the fact that the Skolem symbol is *applied* to $x_1, \dots, x_n$ (or in proof-search, to their corresponding meta-variables $?X_1, \dots, ?X_n$) is a mere implementation trick to record the dependencies between eigenvariables and meta-variables: writing $Y(?X_1, \dots, ?X_n)$ simply records that $?X_1, \dots, ?X_n$ were introduced before Y and any correct instances for them cannot mention Y . This is a smart implementation of the dependencies in the case of pure first-order logic, inasmuch first-order unification will rule out incorrect instances “for free” thanks to the `occurs_check`.

As PSYCHE aims at working modulo theories, it is no longer clear that this specific implementation of dependencies will be as appropriate when another algorithm than first-order unification is run to close branches. A dual implementation of dependencies would consist for instance in recording, whenever a meta-variable $?X$ is introduced, the eigenvariables that existed at that point, among which any correct instance for $?X$ would need to find its free variables. This is for instance the choice in Coq [Coq]: instead of recording the dependencies that are disallowed (as in Skolemisation), one records the dependencies that are allowed. To avoid making any commitment on that choice of implementation, PSYCHE 2.0 is modular in the data-structure that implements dependencies.

Similarly, PSYCHE’s kernel is agnostic in regard of the *constraints* imposed on the instantiation of meta-variables by closing branches: proving a sequent $\Gamma \vdash \Gamma'$ mentioning meta-variables should output (if successful) on constraint σ on these meta-variables, which in pure first-order logic could simply be a *first-order unifier*, but in other theories could be of a different nature (one could think of convex polytops for arithmetic, for instance). A branching rule such as

$$(\wedge^-) \frac{\Gamma \vdash A, \Delta \quad \Gamma \vdash B, \Delta}{\Gamma \vdash A \wedge B, \Delta}$$

should eventually produce the *meet* $\sigma \wedge \sigma'$ of the two constraints σ and σ' returned by the recursive calls of the proof-search function on the two premisses, something we can write as

$$(\wedge^-) \frac{\Gamma \vdash A, \Delta \Rightarrow \sigma \quad \Gamma \vdash B, \Delta \Rightarrow \sigma'}{\Gamma \vdash A \wedge B, \Delta \Rightarrow \sigma \wedge \sigma'}$$

The meet should of course exist (otherwise we should try to find other ways to prove the premisses), in other words $\sigma \wedge \sigma'$ should not be the unsatisfiable constraint $\perp$. There we see an algebraic structure like a semi-lattice appear as the natural concept to spell out the abstract specifications of how constraints work. To avoid the independent exploration of branches before realising they produce incompatible constraints, PSYCHE 2.0 implements the propagation of constraints from one branch to the next, which we could write as

$$(\wedge^-) \frac{\sigma_0 \Rightarrow \Gamma \vdash A, \Delta \Rightarrow \sigma \quad \sigma \Rightarrow \Gamma \vdash B, \Delta \Rightarrow \sigma'}{\sigma_0 \Rightarrow \Gamma \vdash A \wedge B, \Delta \Rightarrow \sigma'}$$

where $\sigma_0 \geq \sigma \geq \sigma'$ for the semi-lattice ordering, breaking the symmetry of the introduction rule for conjunction depending on which of the two branches is actually explored first.

A paper with Damien Rouhling is currently being written on such systems of constraint propagation, which still allow proof-search to backtrack, with persistent data-structures for

constraints. Besides justifying the implementation of PSYCHE 2.0 and identifying the algebraic specifications that an implemented constraint module should satisfy, further development should investigate the connections with *constraint tableaux* [GH03] and *constraint logic programming* [JM94].

Now, what to do with the quantifier features of Psyche 2.0?

First, test them on standard benchmarks and possibly add a TPTP parser to PSYCHE.

Secondly, it was shown in [Far13] how to simulate, as proof-search in $LK^p(\emptyset)$ with quantifiers, the techniques of *clause tableaux* and strong or weak *connection tableaux* (see e.g. [RV01]) for pure first-order logic (hence the empty theory $\emptyset$). How to simulate resolution is also known. So turning these simulations into the implementation of plugins for PSYCHE is the obvious next step, which would allow us to run tests and compare the plugins with other implementations.

Thirdly, we have an approach for mixing first-order reasoning with theories (or, said differently, perform instantiations in presence of a theory); with this we can:

- Investigate to what extent the standard *triggers-based* mechanisms of SMT-solvers for instantiations, can be described in our setting; in particular, Dross [DCKP12] provides theoretical foundations for the use of triggers:

Intuitively, an existential formula $\exists x[k]A$ with a *trigger* $[k]$ allows, as only instantiations of x , those which turn k into a term that is already *known*, as if the notation represented a formula $\exists x(\text{known}(k) \wedge A)$ together with a proof-search policy that forces to prove the left branch $\text{known}(k)$ before the branch on formula A is explored. This strongly suggests a focussing approach in which the predicate $\text{known}(_)$ is positive, so that in the above case $\text{known}(k)$ has to immediately be proved by an axiom. This should be formalised in our focussed calculi.

- More generally look at the various extensions of $DPLL(\mathcal{T})$, in particular the systems with full first-order logic and/or equality as developed by e.g. [Bau00, BT08, BT11], and compare them to $LK^p(\mathcal{T})$ with quantifiers and meta-variables. Note that the abstract setting of LAF might be appropriate for equality itself: since we have abstracted away from connectives, it may be the case that sequent calculi with equality just fit as LAF instances.

Bibliography

- [ADH⁺12] Z. M. Ariola, P. Downen, H. Herbelin, K. Nakata, and A. Saurin. Classical call-by-need sequent calculi: The unity of semantic artifacts. In T. Schrijvers and P. Thiemann, editors, *Proc. of the 11th Int. Symp. Functional and Logic Programming (FLOPS'12)*, volume 7294 of *LNCS*, pages 32–46. Springer-Verlag, 2012. [49](#)
- [AF97] Z. M. Ariola and M. Felleisen. The call-by-need lambda calculus. *J. Funct. Programming*, 7(3):265–301, 1997. [49](#)
- [AFG⁺11] M. Armand, G. Faure, B. Grégoire, C. Keller, L. Théry, and B. Werner. A modular integration of SAT/SMT solvers to Coq through proof witnesses. In *Proc. of the 1st Int. Conf. on Certified Programs and Proofs (CPP'11)*, volume 7086 of *LNCS*, pages 135–150. Springer, 2011. [152](#), [169](#)
- [AH03] Z. M. Ariola and H. Herbelin. Minimal classical logic and control operators. In J. C. M. Baeten, J. K. Lenstra, J. Parrow, and G. J. Woeginger, editors, *Proc. of the 30th Intern. Col. on Automata, Languages and Programming (ICALP)*, volume 2719 of *LNCS*, pages 871–885. Springer-Verlag, 2003. [22](#), [27](#)
- [AH08] Z. M. Ariola and H. Herbelin. Control reduction theories: the benefit of structural substitution. *J. Funct. Programming*, 18(3):373–419, 2008. [27](#)
- [AHS09] Z. M. Ariola, H. Herbelin, and A. Sabry. A type-theoretic foundation of delimited continuations. *Higher-Order and Symbolic Computation*, 22(3):233–273, 2009. online from 2007. [49](#)
- [AHS11] Z. M. Ariola, H. Herbelin, and A. Saurin. Classical call-by-need and duality. In C. L. Ong, editor, *Proc. of the 10th Int. Conf. on Typed Lambda Calculus and Applications (TLCA'11)*, volume 6690 of *LNCS*, pages 27–44. Springer-Verlag, 2011. [49](#)
- [And92] J. M. Andreoli. Logic programming with focusing proofs in linear logic. *J. Logic Comput.*, 2(3):297–347, 1992. [1](#), [71](#), [73](#), [149](#)
- [AP89] J.-M. Andreoli and R. Pareschi. Logic programming with sequent systems, a linear logic approach. In P. Schroeder-Heister, editor, *Proc. of the Int. Work. on Extensions of Logic Programming*, LNCS, pages 1–30. Springer-Verlag, 1989. [1](#), [71](#), [73](#), [149](#)

- [Bar84] H. P. Barendregt. *The Lambda-Calculus, its syntax and semantics*. Studies in Logic and the Foundation of Mathematics. Elsevier, 1984. Second edition. 16, 17, 20, 61, 67
- [Bar91] H. P. Barendregt. Introduction to generalized type systems. *J. Funct. Programming*, 1(2):125–154, 1991. 116
- [Bar92] H. P. Barendregt. Lambda calculi with types. In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum, editors, *Hand. Log. Comput. Sci.*, volume 2, chapter 2, pages 117–309. Oxford University Press, 1992. 149
- [Bau00] P. Baumgartner. FDPLL - a first order Davis-Putnam-Longeman-Loveland procedure. In *Proc. of the 17th Int. Conf. on Automated Deduction (CADE'00)*, volume 1831 of *LNCS*, pages 200–219. Springer-Verlag, 2000. 182
- [BB96] F. Barbanera and S. Berardi. A symmetric lambda-calculus for classical program extraction. *Inform. and Comput.*, 125(2):103–117, 1996. 6, 29, 55
- [BBP11] J. C. Blanchette, S. Böhme, and L. C. Paulson. Extending Sledgehammer with SMT solvers. In *Automated Deduction*, volume 6803 of *LNCS*, pages 116–130. Springer-Verlag, 2011. 152, 169
- [BCP11] F. Besson, P.-E. Cornilleau, and D. Pichardie. Modular SMT proofs for fast reflexive checking inside Coq. In J.-P. Jouannaud and Z. Shao, editors, *Certified Programs and Proofs*, volume 7086 of *LNCS*, pages 151–166. Springer-Verlag, 2011. 152, 169
- [BFM⁺13] F. Bobot, J.-C. Filliâtre, C. Marché, G. Melquiond, and A. Paskevich. The Why3 platform 0.81, 2013. Tutorial and Reference Manual. <http://hal.inria.fr/hal-00822856> 150
- [BG01] H. Barendregt and H. Geuvers. Proof-assistants using dependent type systems. In Robinson and Voronkov [RV01], pages 1149–1238. 149
- [BGL12] A. Bernadet and S. Graham-Lengrand. A simple presentation of the effective topos. Technical report, Laboratoire d’informatique de l’École Polytechnique - CNRS, France, 2012. Available at <http://hal.archives-ouvertes.fr/hal-00844250> 6
- [BGL13] A. Bernadet and S. Graham-Lengrand. Non-idempotent intersection types and strong normalisation. *Logic. Methods Comput. Science*, 9(4), 2013. 6
- [BL11a] A. Bernadet and S. Lengrand. Complexity of strongly normalising λ -terms via non-idempotent intersection types. In M. Hofmann, editor, *Proc. of the 14th Int. Conf. on Foundations of Software Science and Computation Structures (FOSACS'11)*, volume 6604 of *LNCS*. Springer-Verlag, 2011. 7
- [BL11b] A. Bernadet and S. Lengrand. Filter models: non-idempotent intersection types, orthogonality and polymorphism. In M. Bezem, editor, *Proc. of the 20th Annual Conf. of the European Association for Computer Science Logic (CSL'11)*, LIPIcs. Schloss Dagstuhl LCI, 2011. 6, 51, 52, 54, 55

- [BL11c] A. Bernadet and S. Lengrand. Filter models: non-idempotent intersection types, orthogonality and polymorphism - long version. Technical report, LIX, CNRS-INRIA-Ecole Polytechnique, 2011. Available at <http://hal.archives-ouvertes.fr/hal-00600070/en/> 6
- [BMS10] D. Baelde, D. Miller, and Z. Snow. Focused inductive theorem proving. In J. Giesl and R. Hähnle, editors, *Proc. of the 5th Int. Joint Conf. on Automated Reasoning (IJCAR'10)*, volume 6173 of *LNCS*, pages 278–292. Springer-Verlag, 2010. 3, 4
- [BT08] P. Baumgartner and C. Tinelli. The model evolution calculus as a first-order DPLL method. *Artificial Intelligence*, 172(4-5):591–632, 2008. 182
- [BT11] P. Baumgartner and C. Tinelli. Model evolution with equality modulo built-in theories. In N. Bjørner and V. Sofronie-Stokkermans, editors, *Proc. of the 23rd Int. Conf. on Automated Deduction (CADE'11)*, volume 6803 of *LNCS*, pages 85–100. Springer-Verlag, 2011. 182
- [CD78] M. Coppo and M. Dezani-Ciancaglini. A new type assignment for lambda-terms. *Arch. Math. Log.*, 19:139–156, 1978. 7
- [CF58] H. B. Curry and R. Feys. *Combinatory Logic*, volume I. North-Holland, 1958. 1, 13, 14, 18
- [CH00] P.-L. Curien and H. Herbelin. The duality of computation. In *Proc. of the 5th ACM SIGPLAN Int. Conf. on Functional Programming (ICFP'00)*, pages 233–243. ACM Press, 2000. 3, 5, 29, 44, 73, 179
- [Chu41] A. Church. *The Calculi of Lambda Conversion*. Princeton University Press, 1941. 16
- [CMM10] P.-L. Curien and G. Munch-Maccagnoni. The duality of computation under focus. In C. S. Calude and V. Sassone, editors, *Theoretical Computer Science*, volume 323 of *IFIP Advances in Information and Communication Technology*, pages 165–181. Springer-Verlag, 2010. 57, 71
- [Coq] The Coq Proof Assistant. <http://coq.inria.fr/> 5, 6, 87, 149, 169, 181
- [Cro04] T. Crolard. A formulae-as-types interpretation of subtractive logic. *J. Logic Comput.*, 14(4):529–570, 2004. 46, 49, 70
- [CS09] J. R. B. Cockett and L. Santocanale. On the word problem for $\Sigma\Pi$ -categories, and the properties of two-way communication. In E. Grädel and R. Kahle, editors, *Proc. of the 18th Annual Conf. of the European Association for Computer Science Logic (CSL'09)*, volume 5771 of *LNCS*, pages 194–208. Springer-Verlag, 2009. 37
- [dC05] D. de Carvalho. Intersection types for light affine lambda calculus. *ENTCS*, 136:133–152, 2005. 7
- [dC09] D. de Carvalho. Execution time of lambda-terms via denotational semantics and intersection types. *CoRR*, abs/0905.4251, 2009. 7

- [DCKP12] C. Dross, S. Conchon, J. Kanig, and A. Paskevich. Reasoning with triggers. In P. Fontaine and A. Goel, editors, *10th Int. Work. on Satisfiability Modulo Theories, SMT 2012*, volume 20 of *EPiC Series*, pages 22–31. EasyChair, 2012. 182
- [DF89] O. Danvy and A. Filinski. A functional abstraction of typed contexts. Technical Report 89/12, DIKU, University of Copenhagen, 1989. 49
- [DF90] O. Danvy and A. Filinski. Abstracting control. In *Proc. of the 1990 ACM Conf. on LISP and functional programming*, pages 151–160. ACM Press, 1990. 49
- [DGHP99] M. D’Agostino, D. M. Gabbay, R. Hähnle, and J. Posegga. Kluwer Academic Publishers, 1999. 149, 167
- [DJS95] V. Danos, J.-B. Joinet, and H. Schellinx. LKQ and LKT: sequent calculi for second order logic based upon dual linear decompositions of classical implication. In J.-Y. Girard, Y. Lafont, and L. Regnier, editors, *Proc. of the Work. on Advances in Linear Logic*, volume 222 of *London Math. Soc. Lecture Note Ser.*, pages 211–224. Cambridge University Press, 1995. 1, 44, 71, 73
- [DJS97] V. Danos, J.-B. Joinet, and H. Schellinx. A new deconstructive logic: Linear logic. *J. of Symbolic Logic*, 62(3):755–807, 1997. 1, 44, 71, 73
- [DK00] V. Danos and J.-L. Krivine. Disjunctive tautologies as synchronisation schemes. In P. Clote and H. Schwichtenberg, editors, *Proc. of the 9th Annual Conf. of the European Association for Computer Science Logic (CSL’00)*, volume 1862 of *LNCS*, pages 292–301. Springer-Verlag, 2000. 2, 3, 51, 53, 121, 179
- [DL07] R. Dyckhoff and S. Lengrand. Call-by-value λ -calculus and LJQ. *J. Logic Comput.*, 17:1109–1134, 2007. 5
- [DLL62] M. Davis, G. Logemann, and D. W. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5(7):394–397, 1962. 4, 6, 151
- [DP60] M. Davis and H. Putnam. A computing procedure for quantification theory. *J. of the ACM Press*, 7(3):201–215, 1960. 4, 6, 151
- [DP04] K. Došen and Z. Petrić. *Proof-theoretical Coherence*. King’s College Publications, 2004. 37
- [Far13] M. Farooque. *Automated reasoning techniques as proof-search in sequent calculus*. PhD thesis, Ecole Polytechnique, 2013. 6, 73, 149, 150, 153, 156, 161, 162, 163, 172, 177, 182
- [Fel87] M. Felleisen. *The Calculi of λ -v-CS Conversion: A Syntactic Theory of Control and State in Imperative Higher-Order Programming Languages*. PhD thesis, Department of Computer Science, Indiana University, Bloomington, Indiana, 1987. 3, 24
- [FGL13] M. Farooque and S. Graham-Lengrand. Sequent calculi with procedure calls. Technical report, Laboratoire d’informatique de l’École Polytechnique - CNRS, Parsifal - INRIA Saclay, France, 2013. Available at <http://hal.archives-ouvertes.fr/hal-00779199> 6, 149, 156, 169

- [FGLM13] M. Farooque, S. Graham-Lengrand, and A. Mahboubi. A bisimulation between DPLL(T) and a proof-search strategy for the focused sequent calculus. In A. Momigliano, B. Pientka, and R. Pollack, editors, *Proc. of the 2013 Int. Work. on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2013)*. ACM Press, 2013. 6, 149, 153, 161, 162, 163, 169, 172, 173
- [Fil89] A. Filinski. Declarative continuations and categorical duality. Master’s thesis, DIKU, Computer Science Department, University of Copenhagen, 1989. DIKU Rapport 89/11. 29
- [Fis72] M. J. Fischer. Lambda calculus schemata. In *Proc. of the ACM Conf. on Proving Assertions about Programs*, pages 104–109. SIGPLAN Notices, Vol. 7, No 1 and SIGACT News, No 14, 1972. 41, 44, 48
- [FL11] M. Farooque and S. Lengrand. A sequent calculus with procedure calls. Technical report, Laboratoire d’informatique de l’École Polytechnique - CNRS, Parsifal - INRIA Saclay, France, 2011. Available at <http://hal.archives-ouvertes.fr/hal-00690577> 6
- [FLM12a] M. Farooque, S. Lengrand, and A. Mahboubi. Simulating the DPLL(T) procedure in a sequent calculus with focusing. Technical report, Laboratoire d’informatique de l’École Polytechnique - CNRS, Microsoft Research - INRIA Joint Centre, Parsifal & TypiCal - INRIA Saclay, France, 2012. Available at <http://hal.inria.fr/hal-00690392> 6, 149
- [FLM12b] M. Farooque, S. Lengrand, and A. Mahboubi. Two simulations about DPLL(T). Technical report, Laboratoire d’informatique de l’École Polytechnique - CNRS, Microsoft Research - INRIA Joint Centre, Parsifal & TypiCal - INRIA Saclay, France, 2012. Available at <http://hal.archives-ouvertes.fr/hal-00690044> 6, 149
- [FP06] C. Fürmann and D. Pym. Order-enriched categorical models of the classical sequent calculus. *J. Pure Appl. Algebra*, 204(1):21–78, 2006. 37
- [FP13] J.-C. Filliâtre and A. Paskevich. Why3 - where programs meet provers. In M. Felleisen and P. Gardner, editors, *ESOP’13 22nd European Symposium on Programming*, volume 7792 of *LNCS*, pages 125–128. Springer-Verlag, 2013. 150
- [FR94] A. Fleury and C. Retoré. The mix rule. *Math. Structures in Comput. Sci.*, 4(2):273–285, 1994. 35
- [Fre79] G. Frege. *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Verlag von Louis Nebert, 1879 15
- [Gen35] G. Gentzen. Investigations into logical deduction. In *Gentzen collected works*, pages 68–131. Ed M. E. Szabo, North Holland, (1969), 1935. 1, 16, 22, 29, 30, 34, 73, 133, 149
- [GGL14] D. Galmiche and S. Graham-Lengrand. Special issue on computational logic in honour of Roy Dyckhoff. *J. Logic Comput.*, 2014. 5

- [GH03] M. Giese and R. Hähnle. Tableaux + constraints. In M. C. Mayer and F. Pirri, editors, *Proc. of the 16th Int. Conf. on Automated Reasoning with Analytic Tableaux and Related Methods (Tableaux'03)*, volume 2796 of *LNCS*, pages 37–42. Springer-Verlag, 2003. 182
- [Gie00] M. Giese. Proof search without backtracking using instance streams, position paper. In P. Baumgartner and H. Zhang, editors, *3rd Int. Work. on First-Order Theorem Proving (FTP), St. Andrews, Scotland, TR 5/2000 Univ. of Koblenz*, pages 227–228, 2000. 176
- [Gir72] J.-Y. Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Thèse d'état, Université Paris 7, 1972. 6, 51, 52, 55, 56
- [Gir87] J.-Y. Girard. Linear logic. *Theoret. Comput. Sci.*, 50(1):1–101, 1987. 2, 51, 68, 73, 74, 121, 149
- [Gir91] J.-Y. Girard. A new constructive logic: Classical logic. *Math. Structures in Comput. Sci.*, 1(3):255–296, 1991. 73
- [GL08] M. Gabbay and S. Lengrand. The lambda-context calculus. volume 196 of *ENTCS*, pages 19–35, 2008. Revision from the Second Int. Work. on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2007) (was *A(nother) NEW Calculus of Contexts*). 7
- [GL09] M. Gabbay and S. Lengrand. The λ -context calculus. *Inform. and Comput.*, 207(12):1369–1400, 2009. 7
- [GL13] S. Graham-Lengrand. Psyche: a proof-search engine based on sequent calculus with an LCF-style architecture. In D. Galmiche and D. Larchey-Wendling, editors, *Proc. of the 22nd Int. Conf. on Automated Reasoning with Analytic Tableaux and Related Methods (Tableaux'13)*, volume 8123 of *LNCS*, pages 149–156. Springer-Verlag, 2013. 6, 150
- [GL14] S. Graham-Lengrand. Polarities & focussing: a journey from realisability to automated reasoning – Coq proofs of Part II, 2014. <http://www.lix.polytechnique.fr/~lengrand/Work/HDR/> 5, 121, 128, 140, 180
- [GMW79] M. Gordon, R. Milner, and C. Wadsworth. *Edinburgh LCF: a mechanized logic of computation*, volume 78 of *LNCS*. Springer-Verlag, 1979. 150, 167, 169, 179
- [Gri90] T. G. Griffin. A formulae-as-type notion of control. In P. Hudak, editor, *17th Annual ACM Symp. on Principles of Programming Languages (POPL'90)*, pages 47–58. ACM Press, 1990. 3, 13, 24
- [GTL89] J.-Y. Girard, P. Taylor, and Y. Lafont. *Proofs and Types*, volume 7 of *Cambridge Tracts in Theoret. Comput. Sci.* Cambridge University Press, 1989. 22
- [Her05] H. Herbelin. *C'est maintenant qu'on calcule: au coeur de la dualité*. Thèse d'habilitation à diriger des recherches, Université Paris 11, 2005. 29

- [HG08] H. Herbelin and S. Ghilezan. An approach to call-by-name delimited continuations. In G. C. Necula and P. Wadler, editors, *Proc. of the 35th Annual ACM Symp. on Principles of Programming Languages (POPL'08)*, pages 383–394. ACM Press, 2008. 49
- [Hil28] D. Hilbert. Die Grundlagen der Mathematik. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 6(1):65–85, 1928. 15
- [Hil31] D. Hilbert. Die Grundlegung der elementaren Zahlenlehre. *Mathematische Annalen*, 104:485–494, 1931. 4, 107
- [How80] W. A. Howard. The formulae-as-types notion of construction. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays on Combinatory Logic, Lambda Calculus, and Formalism*, pages 479–490. Academic Press, 1980. Reprint of a manuscript written 1969. 1, 13, 16, 18, 133
- [HS97] M. Hofmann and T. Streicher. Continuation models are universal for $\lambda\mu$ -calculus. In *Proc. of the 12th Annual IEEE Symp. on Logic in Computer Science*, pages 387–397. IEEE Computer Society Press, 1997. 41, 44, 48, 49
- [Hyl82] J. Hyland. The effective topos. In A. Troelstra and D. V. Dalen, editors, *The L.E.J. Brouwer Centenary Symposium*, pages 165–216. North Holland Publishing Company, 1982. 6
- [HZ09] H. Herbelin and S. Zimmermann. An operational account of call-by-value minimal and classical λ -calculus in “natural deduction” form. In P. Curien, editor, *Proc. of the 9th Int. Conf. on Typed Lambda Calculus and Applications (TLCA'09)*, volume 5608 of *LNCS*, pages 142–156. Springer-Verlag, 2009. 49
- [Isa] The Isabelle theorem prover. <http://isabelle.in.tum.de/> 169
- [JM94] J. Jaffar and M. J. Maher. Constraint logic programming: a survey. *J. Logic Programming*, 19–20, Supplement 1(0):503 – 581, 1994. Special Issue: Ten Years of Logic Programming. 182
- [Joh36] I. Johansson. Der Minimalkalkül, ein reduzierter intuitionistischer Formalismus. *Compositio Math.*, 4:119–136, 1936. 15
- [KL08] K. Kikuchi and S. Lengrand. Strong normalisation of cut-elimination that simulates β -reduction. In R. Amadio, editor, *Proc. of the 11th Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS'08)*, volume 4962 of *LNCS*, pages 380–394. Springer-Verlag, 2008. 5
- [Kle45] S. Kleene. On the interpretation of intuitionistic number theory. *J. of Symbolic Logic*, 10:109–124, 1945. 2
- [Kri71] J.-L. Krivine. *Introduction to axiomatic set theory*. Dordrecht, Reidel, 1971. 9
- [Kri01] J.-L. Krivine. Typed lambda-calculus in classical Zermelo-Fraenkel set theory. *Arch. Math. Log.*, 40(3):189–205, 2001. 2, 3, 51, 53, 121, 179

- [Lam07] F. Lamarche. Exploring the Gap between Linear and Classical Logic. *Theory and Applications of Categories*, 18(17):473–535, 2007 37
- [Lau02] O. Laurent. *Etude de la polarisation en logique*. Thèse de doctorat, Université Aix-Marseille II, 2002. 1, 73
- [LC09] S. Lescuyer and S. Conchon. Improving Coq propositional reasoning using a lazy CNF conversion scheme. In *Proc. of the 7th Int. Conf. on Frontiers of combining systems (FroCoS'09)*, pages 287–303. Springer-Verlag, 2009 152
- [LDM11] S. Lengrand, R. Dyckhoff, and J. McKinna. A focused sequent calculus framework for proof search in Pure Type Systems. *Logic. Methods Comput. Science*, 7(1), 2011. 6, 149
- [Len03] S. Lengrand. Call-by-value, call-by-name, and strong normalization for the classical sequent calculus. volume 86(4) of *ENTCS*. Elsevier, 2003. Revision from the 3rd Int. Work. on Reduction Strategies in Rewriting and Programming (WRS'03). 5, 29, 34
- [Len06] S. Lengrand. *Normalisation & Equivalence in Proof Theory & Type Theory*. PhD thesis, Université Paris 7 & University of St Andrews, 2006. 7, 9, 40
- [Len08] S. Lengrand. Termination of lambda-calculus with the extra call-by-value rule known as assoc. Technical report, LIX, CNRS-INRIA-Ecole Polytechnique, 2008. Available at <http://hal.inria.fr/inria-00292029> 5
- [LM08] S. Lengrand and A. Miquel. Classical F_ω , orthogonality and symmetric candidates. *Ann. Pure Appl. Logic*, 153:3–20, 2008. 3, 6, 51, 58, 60, 121
- [LM09] C. Liang and D. Miller. Focusing and polarization in linear, intuitionistic, and classical logics. *Theoret. Comput. Sci.*, 410(46):4747–4768, 2009. 1, 3, 69, 71, 74, 75, 95, 96, 102, 104, 149, 152, 153, 179
- [LM11] C. Liang and D. Miller. A focused approach to combining logics. *Ann. Pure Appl. Logic*, 162(9):679–697, 2011. 95, 96, 97, 98, 104
- [LQdF05] O. Laurent, M. Quatrini, and L. T. de Falco. Polarized and focalized linear and classical proofs. *Ann. Pure Appl. Logic*, 134(2-3):217–264, 2005. 1, 73
- [LS86] J. Lambek and P. J. Scott. *Introduction to Higher Order Categorical Logic*. Cambridge University Press, 1986. 22
- [LS05] F. Lamarche and L. Straßburger. Constructing free boolean categories. In P. Panangaden, editor, *Proc. of the 20th Annual IEEE Symp. on Logic in Computer Science*, pages 209–218. IEEE Computer Society Press, 2005. 37
- [Miq09] A. Miquel. Relating classical realizability and negative translation for existential witness extraction. In P. Curien, editor, *Proc. of the 9th Int. Conf. on Typed Lambda Calculus and Applications (TLCA'09)*, volume 5608 of *LNCs*, pages 188–202. Springer-Verlag, 2009. 3, 6, 52, 60, 64

- [Miq11] A. Miquel. Existential witness extraction in classical realizability and via a negative translation. *Logic. Methods Comput. Science*, 7(2), 2011. [3](#), [6](#), [52](#), [60](#), [64](#), [179](#)
- [ML82] P. Martin-Löf. Constructive mathematics and computer programming. In *Proc. of the Sixth Int. Congress for Logic, Methodology, and Philosophy of Science*, pages 153–175. North-Holland, 1982. [13](#)
- [ML84] P. Martin-Löf. *Intuitionistic Type Theory*. Number 1 in Studies in Proof Theory, Lecture Notes. Bibliopolis, 1984. [13](#)
- [MM09] G. Munch-Maccagnoni. Focalisation and classical realisability. In E. Grädel and R. Kahle, editors, *Proc. of the 18th Annual Conf. of the European Association for Computer Science Logic (CSL'09)*, volume 5771 of *LNCS*, pages 409–423. Springer-Verlag, 2009. [1](#), [3](#), [6](#), [44](#), [45](#), [51](#), [54](#), [57](#), [71](#), [73](#), [84](#), [87](#), [121](#), [179](#)
- [MM13] G. Munch-Maccagnoni. *Syntax and Models of a Non-Associative Composition of Programs and Proofs*. PhD thesis, Université Paris Diderot – Paris 7, 2013. [57](#), [71](#)
- [MMZ⁺01] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *DAC*, pages 530–535. ACM Press, 2001. [173](#)
- [MNPS91] D. Miller, G. Nadathur, F. Pfenning, and A. Scedrov. Uniform proofs as a foundation for logic programming. *Ann. Pure Appl. Logic*, 51:125–157, 1991. [1](#), [2](#), [149](#)
- [Mog89] E. Moggi. Computational lambda-calculus and monads. In *Proc. of the 4th Annual IEEE Symp. on Logic in Computer Science*, pages 14–23. IEEE Computer Society Press, 1989. [37](#), [40](#)
- [MP08] S. McLaughlin and F. Pfenning. Imogen: Focusing the polarized inverse method for intuitionistic propositional logic. In I. Cervesato, H. Veith, and A. Voronkov, editors, *Proc. of the 15th Int. Conf. on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR'08)*, volume 5330 of *LNCS*, pages 174–181. Springer-Verlag, 2008. [3](#), [4](#)
- [MV05] P.-A. Mellies and J. Vouillon. Recursive polymorphic types and parametricity in an operational framework. In P. Panangaden, editor, *Proc. of the 20th Annual IEEE Symp. on Logic in Computer Science*, pages 82–91. IEEE Computer Society Press, 2005. [51](#)
- [Nig09] V. Nigam. *Exploiting non-canonicity in the sequent calculus*. PhD thesis, Ecole Polytechnique, 2009. Available at <https://tel.archives-ouvertes.fr/pastel-00005487/> [93](#), [94](#)
- [NM10] V. Nigam and D. Miller. A framework for proof systems. *J. of Automated Reasoning*, 45(2):157–188, 2010. [93](#), [94](#)
- [NOT06] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(*T*). *J. of the ACM Press*, 53(6):937–977, 2006. [4](#), [151](#), [152](#), [157](#), [158](#), [163](#), [173](#), [179](#)

- [NPW02] T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, volume 2283 of *LNCS*. Springer-Verlag, 2002. 169
- [Par92] M. Parigot. $\lambda\mu$ -calculus: An algorithmic interpretation of classical natural deduction. In A. Voronkov, editor, *Proc. of the Int. Conf. on Logic Programming and Automated Reasoning (LPAR'92)*, volume 624 of *LNCS*, pages 190–201. Springer-Verlag, 1992. 3, 25, 30
- [Par97] M. Parigot. Proofs of strong normalisation for second order classical natural deduction. *J. of Symbolic Logic*, 62(4):1461–1479, 1997. 3, 51, 55, 56, 179
- [PB12] L. C. Paulson and J. C. Blanchette. Three years of experience with Sledgehammer, a practical link between automatic and interactive theorem provers. In G. Sutcliffe, S. Schulz, and E. Ternovska, editors, *IWIL 2010*, volume 2 of *EPiC Series*, pages 1–11. EasyChair, 2012. 169
- [Pit03] A. M. Pitts. Nominal logic, a first order theory of names and binding. *Inform. and Control*, 186:165–193, 2003. 7, 91
- [Plo75] G. D. Plotkin. Call-by-name, call-by-value and the lambda-calculus. *Theoret. Comput. Sci.*, 1:125–159, 1975. 3, 37, 38, 39
- [Pol04] E. Polonovski. Strong normalization of $\lambda\mu\tilde{\mu}$ -calculus with explicit substitutions. In I. Walukiewicz, editor, *Proc. of the 7th Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS'04)*, volume 2987 of *LNCS*, pages 423–437. Springer-Verlag, 2004. 34
- [PSI] The PSI project. 2009-2013. <http://www.lix.polytechnique.fr/~lengrand/PSI>. 6, 149, 150
- [Psy] Psyche: the Proof-Search factorY for Collaborative HEuristics. <http://www.lix.polytechnique.fr/~lengrand/Psyche> 3, 4, 6, 167, 177, 179
- [Rea00] S. Read. Harmony and autonomy in classical logic. *J. of Philosophical Logic*, 29(2):123–154, 2000. 81
- [Rea10] S. Read. General-elimination harmony and the meaning of the logical constants. *J. of Philosophical Logic*, 39(5):557–576, 2010. 81
- [Rey72] J. C. Reynolds. Definitional interpreters for higher-order programming languages. In *Proc. of the ACM annual Conf.*, pages 717–740, 1972. 3, 24, 39, 44, 51
- [Roc05] J. Rocheteau. lambda- μ -calculus and duality: Call-by-name and call-by-value. In J. Giesl, editor, *Proc. of the 16th Int. Conf. on Rewriting Techniques and Applications (RTA'05)*, volume 3467 of *LNCS*, pages 204–218. Springer-Verlag, 2005. 29
- [RV01] J. A. Robinson and A. Voronkov, editors. *Handbook of Automated Reasoning (in 2 volumes)*. Elsevier and The MIT Press, 2001. 2, 182, 184
- [Sau05] A. Saurin. Separation with streams in the lambda μ -calculus. In P. Panangaden, editor, *Proc. of the 20th Annual IEEE Symp. on Logic in Computer Science*, pages 356–365. IEEE Computer Society Press, 2005. 49

- [Sau08] A. Saurin. On the relations between the syntactic theories of lambda-mu-calculi. In *Proc. of the 17th Annual Conf. of the European Association for Computer Science Logic (CSL'08)*, volume 5213 of *LNCS*, pages 154–168. Springer-Verlag, 2008. [49](#)
- [Sau10a] A. Saurin. A hierarchy for delimited continuations in call-by-name. In C. L. Ong, editor, *Proc. of the 13th Int. Conf. on Foundations of Software Science and Computation Structures (FOSSACS'10)*, volume 6014 of *LNCS*, pages 374–388. Springer-Verlag, 2010. [49](#)
- [Sau10b] A. Saurin. Standardization and böhm trees for lambda μ -calculus. In M. Blume, N. Kobayashi, and G. Vidal, editors, *Proc. of the 10th Int. Symp. Functional and Logic Programming (FLOPS'10)*, volume 6009 of *LNCS*, pages 134–149. Springer-Verlag, 2010. [49](#)
- [Sau10c] A. Saurin. Typing streams in the lambda μ -calculus. *ACM Trans. on Comput. Logic*, 11(4), 2010. [49](#)
- [Sau12] A. Saurin. Böhm theorem and böhm trees for the $\lambda\mu$ -calculus. *Theoret. Comput. Sci.*, 435:106–138, 2012. [49](#)
- [Sch50] K. Schütte. Beweistheoretische Erfassung der unendlichen Induktion in der Zahlentheorie. *Mathematische Annalen*, 122:369–389, 1950. [4](#), [107](#)
- [Sel01] P. Selinger. Control categories and duality: on the categorical semantics of the $\lambda\mu$ -calculus. *Math. Structures in Comput. Sci.*, 11(2):207–260, 2001. [3](#), [42](#), [49](#), [72](#), [73](#), [179](#)
- [SR98] T. Streicher and B. Reus. Classical logic, continuation semantics and abstract machines. *J. Funct. Programming*, 8(6):543–572, 1998. [49](#)
- [Str11] L. Straßburger. *Towards a Theory of Proofs of Classical Logic*. Habilitation thesis, Université Denis Diderot – Paris 7, 2011. [22](#), [37](#)
- [SU06] M. H. B. Sørensen and P. Urzyczyn. *Lectures on the Curry-Howard Isomorphism*. Studies in Logic and the Foundations of Mathematics. Elsevier, 2006. [14](#)
- [SW00] C. Strachey and C. P. Wadsworth. Continuations: A mathematical semantics for handling fulljumps. *Higher-Order and Symbolic Computation*, 13:135–152, 2000. [3](#), [24](#)
- [Tai67] W. W. Tait. Intensional interpretations of functionals of finite type I. *J. of Symbolic Logic*, 32:198–212, 1967. [51](#)
- [Tai75] W. W. Tait. A realizability interpretation of the theory of species. In *Logic Colloquium*, volume 453 of *LNM*, pages 240–251. Springer-Verlag, 1975. [51](#), [52](#), [55](#), [56](#)
- [Ten78] N. Tennant. *Natural logic*. Edinburgh University Press, 1978. [81](#)
- [Ter03] Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoret. Comput. Sci.* Cambridge University Press, 2003. [9](#)

- [Tin07] C. Tinelli. An abstract framework for satisfiability modulo theories. In N. Olivetti, editor, *Proc. of the 16th Int. Conf. on Automated Reasoning with Analytic Tableaux and Related Methods (Tableaux'07)*, volume 4548 of *LNCS*. Springer-Verlag, 2007. Invited talk, available at <http://ftp.cs.uiowa.edu/pub/tinelli/talks/TABLEAUX-07.pdf>. 150
- [TS00] A. S. Troelstra and H. Schwichtenberg. *Basic Proof Theory*. Cambridge University Press, 2000. 10, 30, 34, 35, 66
- [Twe] The Twelf Project. <http://twelf.org> 149
- [Urb00] C. Urban. *Classical Logic and Computation*. PhD thesis, University of Cambridge, 2000. 5, 29, 30
- [VO02] J. VAN OOSTEN. Realizability: a historical essay. *Math. Structures in Comput. Sci.*, 12:239–263, 2002. 2
- [Wad03] P. Wadler. Call-by-value is dual to call-by-name. In *Proc. of the 8th ACM SIGPLAN Int. Conf. on Functional programming (ICFP'03)*, volume 38(9), pages 189–201. ACM Press, 2003. 29, 31, 44, 45, 47, 70
- [Web11] T. Weber. SMT solvers: New oracles for the HOL theorem prover. *International Journal on Software Tools for Technology Transfer (STTT)*, 13(5):419–429, 2011. 152, 169
- [Zei08a] N. Zeilberger. Focusing and higher-order abstract syntax. In G. C. Necula and P. Wadler, editors, *Proc. of the 35th Annual ACM Symp. on Principles of Programming Languages (POPL'08)*, pages 359–369. ACM Press, 2008. 3, 6, 74, 77, 79, 87, 88, 89, 179
- [Zei08b] N. Zeilberger. On the unity of duality. *Ann. Pure Appl. Logic*, 153(1-3):66–96, 2008. 3, 6, 74, 77, 79, 87, 88, 89, 179
- [Zei09] N. Zeilberger. *The Logical Basis of Evaluation Order and Pattern-Matching*. PhD thesis, Carnegie Mellon University, 2009. 133
- [Zei10] N. Zeilberger. Polarity and the logic of delimited continuations. In J.-P. Jouannaud, editor, *Proc. of the 25th Annual IEEE Symp. on Logic in Computer Science*, pages 219–227. IEEE Computer Society Press, 2010. 77, 79

List of Figures

1	Simply-typed λ -calculus	17
2	Natural Deduction for minimal logic - $\text{NJ} \Rightarrow$	17
3	Semantics of the simply-typed λ -calculus in a CCC	19
4	$(\mathbf{I}, \mathbf{K}, \mathbf{S})$ -combinators as λ -terms	20
5	The proof system corresponding to the simply-typed $\lambda\mu$ -calculus	26
6	Typing system for $\mathbf{L}$	31
7	A proof of LEM	32
8	CBN and CBV reduction in System $\mathbf{L}$ (first attempt)	44
9	CBN and CBV reduction in System $\mathbf{L}$	45
10	System F	52
11	Semantics of expressions and formulae	62
12	Rewrite system for polarised System $\mathbf{L}$	70
13	LKF	75
14	Positive decomposition relation	78
15	Negative decomposition relation	79
16	Big-step LKF, v1	79
17	Big-step LKF, v2	79
18	Big-step LKF, v3	80
19	Big-step LKF, v4	80
20	Big-step LKF, v5	81
21	Decomposition with patterns	82
22	Typing for the pattern-matching calculus	83
23	LAF	93
24	Decomposition relation for LAF_{K1}	96
25	Decomposition relation for LAF_{K2}	100
26	Decomposition relation for LAF_J	101
27	LAF	113
28	Decomposition relation for LAF_{K1}	114

29	Free labels	139
30	Renaming	139
31	Cut-elimination	141
32	Typing a syntactic abstract machine	142
33	Substitution	144
34	System $LK^p(\mathcal{T})$	155
35	Elementary $DPLL(\mathcal{T})$	157
36	Examples of elementary $DPLL(\mathcal{T})$ runs	158
37	Decomposition relation for LAF_{K1p}	164
38	$LAF(\mathcal{T})$	166

Appendix A

Basic definitions for categories

DEFINITION 130 (Category) A *category* is the combination of

- a class of elements called *objects*, denoted $A, B, \dots$
- for every pair of objects A and B , a class $\text{hom}(A, B)$ of elements called *morphisms from A to B* ; the expression $f: A \longrightarrow B$ denotes that f is a morphism from A to B ;
- for every object A , a morphism Id_A called *identity*;
- for every objects A, B, C , a binary operation called *composition* mapping every $f: A \longrightarrow B$ and $g: B \longrightarrow C$ to a morphism $f \cdot g: A \longrightarrow C$

such that the following properties holds

- composition is associative ($(f \cdot g) \cdot h = f \cdot (g \cdot h)$)
- identities are units for composition ($\text{Id}_A \cdot f = f \cdot \text{Id}_A = f$).

Given a category, we often use *diagrams* to represent a collection of objects -represented as vertices- and morphisms -represented as labelled arrows between vertices. Using morphism composition, each path between any two given vertices unambiguously represents a morphism. A diagram *commutes* when for each pair of vertices A and B , all paths from A to B represent equal morphisms.

Two objects A and B are *isomorphic* when there are two morphisms $f: A \longrightarrow B$ and $g: B \longrightarrow A$ such that $f \cdot g = \text{Id}_A$ and $g \cdot f = \text{Id}_B$. ※

Standard examples of categories are: the categories of sets and functions (objects are sets and morphisms from A to B are functions from A to B), the category of sets and relations (objects are sets and morphisms from A to B are relations from A to B), etc. Groups form a particular kind of categories (where there is only one object, and elements of the group are the morphisms from that object to itself). Partially ordered sets form another particular kind of categories (where there is at most one morphism between any two objects), etc.

DEFINITION 131 (Cartesian Closed Category)

A *Cartesian Closed Category* (CCC), is a category such that

- there is an object, denoted 1 and called *terminal object*, such that, for each object A , there is a unique morphism $1_A: A \rightarrow 1$;
- for every two objects A and B , there is an object, denoted $A \times B$ and called the *product* of A and B , together with two morphisms $\pi_1: A \times B \rightarrow A$ and $\pi_2: A \times B \rightarrow B$, called the first and second *projections*, satisfying the following property:
for every object C and morphisms $f_1: C \rightarrow A$ and $f_2: C \rightarrow B$, there is a unique $f: C \rightarrow A \times B$, denoted $\langle f_1, f_2 \rangle$, such that the following diagram commutes

$$\begin{array}{ccccc} & & C & & \\ & f_1 \swarrow & \downarrow f & \searrow f_2 & \\ A & \xleftarrow{\pi_1} & A \times B & \xrightarrow{\pi_2} & B \end{array}$$

- for every two objects A and B , there is an object, denoted B^A and called the *exponential* of A and B , together with a morphism $\text{eval}: B^A \times A \rightarrow B$, satisfying the following property:
for every object X and morphism $g: X \times A \rightarrow B$ there is a unique $f: X \rightarrow B^A$, denoted Λg , such that the following diagram commutes

$$\begin{array}{ccc} X \times A & & \\ \downarrow \langle \Lambda g, \text{Id}_A \rangle & \searrow g & \\ B^A \times A & \xrightarrow{\text{eval}} & B \end{array}$$

We choose the convention that products are associative to the left, i.e. $(A \times B) \times C$ can be abbreviated as $A \times B \times C$. A family of morphisms $\pi_{i/n}: A_1 \times \cdots \times A_n \rightarrow A_i$, for $1 \leq i \leq n$, can be defined by composing the two projections in the obvious way:

$$\begin{aligned} \pi_{1/1} &:= \text{Id}_{A_1} \\ \pi_{n/n} &:= \pi_2 && \text{when } 1 < n \\ \pi_{p/n} &:= \pi_1 \cdot \pi_{p/n-1} && \text{when } p < n \end{aligned}$$

※

REMARK 83 One can quickly check that the terminal object, products and exponentials are unique up to isomorphism (i.e. two objects satisfying the property of the terminal object, or the product / exponential object for a given A and B , are isomorphic); hence the notations 1 , $A \times B$, B^A .

※