

A Framework for Specifying, Prototyping, and Reasoning about
Computational Systems

SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Andrew Jude Gacek

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Gopalan Nadathur, Advisor
September, 2009

© Andrew Jude Gacek 2009

Acknowledgments

Many people have supported me during the development of this thesis and I owe them all a debt of gratitude.

Firstly, I would like to thank my advisor Gopalan Nadathur for his patience and guidance which have played a significant part in my development as a researcher. His willingness to share his opinions on everything from academic life to playing squash has helped me to develop a perspective and to have fun while doing this. I look forward to continuing my interactions with him far into the future.

I am grateful to Dale Miller for sharing with me an excitement for research and an appreciation of the uncertainty that precedes understanding. I have never met anybody else who so enjoys when things seem amiss, because he knows that a new perspective will eventually emerge and bring clarity.

This thesis has been heavily influenced by the time I have spent working with Alwen Tiu, David Baelde, Zach Snow, and Xiaochu Qi. Understanding their work has given me a deeper understanding of my own research and its role in the bigger picture. I am thankful for the time I have had with each and every one of them.

I have been inspired in my studies by my friends Mike Whalen and Jared Davis. Their intelligence, drive, and curiosity are remarkable and have challenged me to work harder so that I may hope to be considered their equals.

I also want to thank my committee members Eric Van Wyk, Mats Heimdahl, and Wayne Richter for their time and for their guidance in my research career.

Finally, I am thankful to the many people who have supported me long before this thesis began. I want especially to thank my wife, Ann, for her patience, understanding, and love, and my parents for their never-ending encouragement and support. To the rest of my family and friends: I thank you all!

Work on this thesis has been partially funded by the NSF Grants CCR-0429572 and CCF-0917140. Support has also been received from a research contract from Boston Scientific and from funds provided by the Institute of Technology and the Department of Computer Science and Engineering at the University of Minnesota. Opinions, findings, and conclusions or recommendations expressed in this thesis should be understood as mine. In particular, they do not necessarily reflect the views of the National Science Foundation.

Abstract

A major motivation for formal systems such as programming languages and logics is that they support the ability to perform computations in a safe, secure, and understandable way. A considerable amount of effort has consequently been devoted to developing tools and techniques for structuring and analyzing such systems. It is natural to imagine that research in this setting might draw benefits from its own labor. In particular, one might expect the study of formal systems to be conducted with the help of languages and logics designed for such study. There are, however, significant problems that must be solved before such a possibility can be made a practical reality. One such problem arises from the fact that formal systems often have to treat objects such as formulas, proofs, programs, and types that have an inherent binding structure. In this context, it is necessary to provide a flexible and logically precise treatment of related notions such as the equality of objects under the renaming of bound variables and substitution that respects the scopes of binders; there is considerable evidence that if such issues are not dealt with in an intrinsic and systematic way, then they can overwhelm any relevant reasoning tasks. For a logic to be useful in this setting, it must also support rich capabilities such as those for inductive reasoning over computations that are described by recursion over syntax.

This thesis concerns the development of a framework that facilitates the design and analysis of formal systems. Specifically, this framework is intended to provide 1) a specification language which supports the concise and direct description of a system based on its informal presentation, 2) a mechanism for animating the specification language so that descriptions written in it can quickly and effectively be turned into prototypes of the systems they are about, and 3) a logic for proving properties of descriptions provided in the specification language and thereby of the systems they encode. A defining characteristic of the proposed framework is that it is based on two separate but closely intertwined logics. One of these is a specification logic that facilitates the description of computational structure while the other is a logic that exploits the special characteristics of the specification logic to support reasoning about the computational behavior of systems that are described using it. Both logics embody a natural treatment of binding structure by using the λ -calculus as a means for representing objects and by incorporating special mechanisms for working with such structure. By using this technique, they lift the treatment of binding from the object language into the domain of the relevant meta logic, thereby allowing the specification or analysis components to focus on the more essential logical aspects of the systems that are encoded.

One focus of this thesis is on developing a rich and expressive reasoning logic that is of use within the described framework. This work exploits a previously developed capability of definitions for embedding recursive specifications into the reasoning logic; this notion

of definitions is complemented by a device for a case-analysis style reasoning over the descriptions they encode. Use is also made of a special kind of judgment called a generic judgment for reflecting object language binding into the meta logic and thereby for reasoning about such structure. Existing methods have, however, had a shortcoming in how they combine these two devices. Generic judgments lead to the introduction of syntactic objects called nominal constants into formulas and terms. The manner in which such objects are introduced often ensures that they satisfy certain properties which are necessary to take note of in the reasoning process. Unfortunately, this has heretofore not been possible to do. To overcome this problem, we introduce a special binary relation between terms called *nominal abstraction* and show this can be combined with definitions to encode the desired properties. The treatment of definitions is further enriched by endowing them with the capability of being interpreted inductively or co-inductively. The resulting logic is shown to be consistent and examples are presented to demonstrate its richness and usefulness in reasoning tasks.

This thesis is also concerned with the practical application of the logical machinery it develops. Specifically, it describes an interactive, tactic-style theorem prover called Abella that realizes the reasoning logic. Abella embodies the use of lemmas in proofs and also provides intuitively well-motivated tactics for inductive and co-inductive reasoning. The idea of reasoning using two-levels of logic is exploited in this context. This form of reasoning, pioneered by McDowell and Miller, embeds the specification logic explicitly into the reasoning logic and then reasons about particular specifications through this embedding. The usefulness of this approach is demonstrated by showing that general properties can be proved about the specification logic and then used as lemmas to simplify the overall reasoning process. We use these ideas together with Abella to develop several interesting and challenging proofs. The examples considered include ones in the recently proposed POPLmark challenge and a formalization of Girard's proof of strong normalization for the simply-typed λ -calculus. We also explore the notion of adequacy that relates theorems proved using Abella to the properties of the object systems that are ultimately of primary interest.

Contents

List of Figures	vii
1 Introduction	1
1.1 A Specification, Prototyping, and Reasoning Framework	1
1.2 An Illustration of the Application of the Framework	4
1.3 The Contributions of this Thesis	8
1.4 Overview of the Thesis	11
2 A Logic for Specifying Computational Systems	14
2.1 The Syntax and Semantics of the Logic	15
2.2 Properties of the Specification Logic	18
2.3 Example Encoding in the Specification Logic	20
2.4 Adequacy of Encodings in the Specification Logic	22
3 A Logic for Reasoning About Specifications	26
3.1 A Logic with Generic Quantification	28
3.1.1 The Basic Syntax	28
3.1.2 Generic Judgments and ∇ -quantification	29
3.1.3 A Sequent Calculus Presentation of the Core Logic	31
3.2 Characterizing Occurrences of Nominal Constants	33
3.2.1 Substitutions and their Interaction with Nominal Constants	34
3.2.2 Nominal Abstraction	38
3.2.3 Proof Rules for Nominal Abstraction	40
3.2.4 Computing Complete Sets of Nominal Abstraction Solutions	42
3.3 Definitions, Induction, and Co-induction	45
3.4 A Pattern-Based Form for Definitions	48
3.5 Examples	54
3.5.1 Properties of ∇ and Freshness	55
3.5.2 Polymorphic Type Generalization	56
3.5.3 Arbitrarily Cascading Substitutions	57
4 Some Properties of the Meta-logic	59
4.1 Consistency of the Meta-logic	59
4.2 Adequacy of Encodings and Theorems in the Meta-logic	71

5	An Interactive Theorem Prover for the Meta-logic	77
5.1	A Framework for Using Lemmas	78
5.2	An Annotation Based Scheme for Induction	80
5.3	Extensions to the Basic Scheme for Induction	83
5.3.1	Induction on a Predicate in the Scope of Generic Quantifiers	83
5.3.2	Induction in the Presence of Additional Premises	84
5.3.3	Delayed Applications of the Induction Hypothesis	85
5.3.4	Nested Inductions	86
5.4	An Annotation Based Scheme for Co-induction	89
6	A Two-level Logic Approach to Reasoning	92
6.1	Encoding the Specification Logic	93
6.1.1	Encoding the Syntax of the Specification Logic	94
6.1.2	Encoding the Semantics of the Specification Logic	95
6.1.3	Some Provable Properties of the Specification Logic	98
6.2	Formalizing Properties of the Specification Logic	100
6.3	An Example of the Two-level Logic Reasoning Approach	101
6.4	Architecture of a Two-level Logic Based Theorem Prover	103
6.5	Adequacy for the Two-level Logic Approach to Reasoning	105
6.5.1	Adequacy of Encoding of the Specification Logic	106
6.5.2	Adequacy of Type Preservation Example	112
7	Applications of The Framework	113
7.1	Type-uniqueness for the Simply-typed λ -calculus	115
7.2	The POPLmark Challenge	117
7.3	Path Equivalence for λ -terms	122
7.4	Conversion between HOAS and de Bruijn Notation	127
7.5	Formalizing Tait-Style Proofs for Strong Normalization	130
7.5.1	Typing and One-step Reduction	132
7.5.2	The Logical Relation	133
7.5.3	Arbitrary Cascading Substitutions and Freshness Results	135
7.5.4	The Final Result	138
8	Related Work	139
8.1	First-order Representations	140
8.1.1	Named Representation	140
8.1.2	Nameless Representation	141
8.1.3	Locally Nameless Representation	142
8.2	Nominal Representations	144
8.3	Higher-order Representations	146
8.3.1	Hybrid	146
8.3.2	Twelf	147

8.3.3	Delphin	150
8.3.4	Tac	151
9	Conclusion and Future Work	154
9.1	More Permissive Stratification Conditions for Definitions	155
9.2	Context Inversion Properties	156
9.3	Types and Explicit Typing	157
9.4	Alternate Specification Logics	158
9.5	Focusing and Proof Search	158
9.6	An Integrated Framework	158
	Bibliography	160

List of Figures

1.1	Evaluation in the simply-typed λ -calculus	5
1.2	Typing in the simply-typed λ -calculus	5
1.3	A Horn clause-like encoding of evaluation and typing	6
2.1	Derivation rules for the hH^2 logic	18
2.2	hH^2 specification of evaluation and typing	21
3.1	The core rules of \mathcal{G}	32
3.2	Nominal abstraction rules	40
3.3	A variant of $\triangleright\mathcal{L}$ based on <i>CSNAS</i>	40
3.4	Introduction rules for atoms whose predicate is defined as $\forall\vec{x}. p \vec{x} \triangleq B p \vec{x}$.	45
3.5	The induction left and co-induction right rules	47
3.6	Introduction rules for a pattern-based definition \mathcal{D}	50
3.7	Induction rule for pattern-based definitions	53
4.1	An evaluation relation for untyped λ -terms	73
4.2	An encoding of the evaluation relation in Figure 4.1	73
5.1	Transition diagrams for two different processes	89
6.1	Second-order hereditary Harrop logic in \mathcal{G}	96
6.2	<i>prog</i> clauses for simply-typed λ -calculus	102
7.1	Potential <i>ctx</i> definition without nominal abstraction	116
7.2	Algorithmic subtyping rules for System F_{\prec}	118
7.3	Specification of algorithmic subtyping for System F_{\prec}	120
7.4	Tree form of $\lambda x.(x(\lambda y.y))$	122
7.5	Specification of paths through λ -terms	123
7.6	Specification of translation between HOAS and de Bruijn notation	127
7.7	Specification of typing and one-step reduction	131
9.1	Typing judgment directly within \mathcal{G}	156

Chapter 1

Introduction

In this thesis we are interested in developing a framework for mechanizing the specification and prototyping of formal systems and also the process of reasoning about the properties of such systems based on their specifications. The formal systems that are of interest to us are ones that concern computation: for instance, they might characterize evaluation and typing in a programming language, provability in a logic, or behavior in a concurrency system. Formal systems of these kinds typically manipulate syntactically complex objects such as formulas, proofs, and programs. Mechanized specification and reasoning about such systems has proven difficult to achieve through the use of traditional tools and techniques [ABF⁺05]. We propose a framework here which overcomes these difficulties and, through this process, brings the benefits of automation and computer-aided verification to bear on the development of these types of systems. In particular, this thesis proposes a framework that facilitates the development of such systems by providing 1) a specification language which supports the concise and direct description of a system based on its informal presentation, 2) a mechanism for animating the specification language so that descriptions written in it can quickly and effectively be turned into prototypes of the systems they are about, and 3) a logic for proving properties of descriptions provided in the specification language and thereby of the systems they encode.

1.1 A Framework for Specification, Prototyping, and Reasoning

The formal systems that we would like to specify and reason about are all characterized by the fact that they are based on syntactic expressions and their behavior is determined by the structure of these expressions. For brevity we will refer to such systems simply as

computational systems. A popular approach to describing such systems starts by describing various possible judgments over the syntax of the systems. Then rule schemas are presented where each schema allows a judgment to be formed from other judgments, often in a compositional manner. Finally, instances of these rules schemas are chained together into a *derivation* where each premise judgment of a rule instance is the consequence judgment of another rule instance. A judgment is said to hold if and only if it is the final conclusion judgment of derivation. Thus one can understand the behavior of a system by studying the rule schemas for forming judgments about the system. This approach to describing a computational system is known as structural operational semantics [Plo81].

Structural operational semantics descriptions have a logical flavor in that one simply describes a few declarative rules for manipulating syntax and these are orchestrated together to reach larger conclusions about the behavior of the system. The framework we propose allows for such descriptions to be formally specified via a *specification logic* similar to the logic of Horn clauses. We call such an encoding of a computational system into this logic a *specification*. More specifically, the system syntax is encoded as specification logic terms, judgments are encoded as specification logic atomic formulas, and rules are encoded as richer specification logic formulas. Derivations of atomic formulas within the specification logic then correspond to derivations in structural operational semantics descriptions. Thus we can study a wide variety of computational systems via a study of the specification logic.

In order to interact with computational systems, our proposed framework supports prototyping based on the system specification. This prototyping is driven directly by the formal specification, by giving a computational interpretation of the specification logic in the same sense that Prolog provides a computational interpretation to the logic of Horn clauses. This eliminates the need for the framework user to manually develop a prototype based on the specification, thus avoiding a source of potential errors. Also, as the specification evolves this ensures that the prototype remains faithful to the current specification.

The specification of a computational system consists of local rules about the system behavior, but one is often interested in global properties of the system. For example, pro-

programming language designers often describe the rules for evaluation and typing judgments for a language and then prove properties which relate the two judgments together such as that the evaluation judgment preserves the typing judgment. Such properties ensure that the language is well-behaved relative to programmers' expectations. In order to prove these properties about a structural operational semantics description one must be able to analyze the ways in which derivations may be formed. In the example of proving that evaluation preserves typing, one may inductively analyze the possible forms that a derivation of an evaluation judgment may have and for each possibility argue that the typing judgment for the evaluated term can be restructured into a typing judgment for the term which results from the evaluation.

The proposed framework allows for reasoning over structural operational semantics descriptions via a *meta-logic*. The meta-logic contains mechanisms such as induction and co-induction which are essential to sophisticated reasoning. The meta-logic also contains a mechanism called *definitions* which allows one to connect atomic judgments to descriptions of behavior in a "closed world" fashion. Thus, it allows for both positive reasoning, *i.e.*, showing that a judgment holds, and negative reasoning, *i.e.*, analyzing why a judgment holds. This allows one to easily carry out the case analysis-like reasoning described in the example of typing and evaluation.

We refer to this second logic as a meta-logic because, in our approach, we use it to encode the entire specification logic, rather than to encode each specification independently. We then reason about particular specifications by reasoning about their descriptions in the specification logic. This style of reasoning was pioneered by McDowell and Miller [MM02] and is called the *two-level logic approach to reasoning*. One of its benefits is that it allows us to reason over specifications exactly as they are written and used in prototyping. Another is that it allows properties of the specification logic to be formally proven once and for all in the meta-logic and then used freely during reasoning. In practice, many tedious substitution lemmas proven about particular specifications are subsumed by these more general properties of the specification logic.

A pervasive issue in the computational systems of interest is dealing with the binding structure of syntactic objects. For example, to develop a programming language we need to formalize the rules for binding local variables which requires a systematic way 1) to associate variable occurrences with their binders, 2) to treat objects which differ only in the name of bound variables as being identical, and 3) to realize a logically correct notion of capture-avoiding substitution which respects the binding structure of objects. Our proposed framework addresses all of these issues by mapping the binding structure of objects into the abstraction mechanism of the meta-language, *i.e.*, the specification logic during specification and the meta-logic during reasoning. This is called a *higher-order abstract syntax* representation [MN87, PE88]. In this way, the meta-language notion of binding describes how variable occurrences are associated to the binder, the meta-language notion of equality provides a way to identify objects differing only in the names of bound variables, and meta-language function application and reduction realize capture-avoiding substitution.

1.2 An Illustration of the Application of the Framework

Throughout this thesis we will use the example of the simply-typed λ -calculus [Chu41, Bar84]. This is a compact example which highlights many of the essential difficulties involved in specifying, prototyping, and reasoning about a computational system with binding. Anytime we use such a system as the focus of study we shall refer to it as the *object language* or the *object logic*.

The syntax of the simply-typed λ -calculus is made up of two classes of expressions called *types* and *pre-terms* which are defined, respectively, by the following grammar rules.

$$a ::= i \mid a \rightarrow a \qquad t ::= x \mid (\lambda x : a. t) \mid (t t)$$

Here x is variable occurrence and in the expression $(\lambda x : a. t)$ the x is to be considered bound within the expression t . We assume the standard notions of binding including free and bound variables, equivalence under renaming of bound variables, and a notion of capture-avoiding

$$\frac{}{(\lambda x:a. r) \Downarrow (\lambda x:a. r)} \qquad \frac{m \Downarrow (\lambda x:a. r) \quad r[x := n] \Downarrow v}{(m \ n) \Downarrow v}$$

Figure 1.1: Evaluation in the simply-typed λ -calculus

$$\frac{x : a \in \Gamma}{\Gamma \vdash x : a} \qquad \frac{\Gamma, x : a \vdash r : b}{\Gamma \vdash (\lambda x:a. r) : a \rightarrow b} \quad x \notin \text{dom}(\Gamma) \qquad \frac{\Gamma \vdash m : a \rightarrow b \quad \Gamma \vdash n : a}{\Gamma \vdash (m \ n) : b}$$

Figure 1.2: Typing in the simply-typed λ -calculus

substitution denoted by $t[x := s]$. Note, however, that when one formally specifies this system within a framework, these notions will need to be dealt with somehow. We will denote types using variables named a, b, c , and d , pre-terms using variables named m, n, r, s, t , and v , and object language variables using x, y , and z .

We define a notion of big-step call-by-name weak reduction which we call simply *evaluation*. This is denoted by the judgment $t \Downarrow v$ which can be read as “ t evaluates to v .” The rules for forming derivations of this judgment are presented in Figure 1.1.

We define a notion of typing via the judgment $\Gamma \vdash t : a$ which can be read as “ t has type a relative to the context Γ .” Here Γ is called a *typing context* and is described by the following grammar.

$$\Gamma ::= \cdot \mid \Gamma, x : a$$

We will write a context of the form $\cdot, x_1 : a_1, \dots, x_n : a_n$ simply as $x_1 : a_1, \dots, x_n : a_n$. We define $\text{dom}(x_1 : a_1, \dots, x_n : a_n)$ as $\{x_1, \dots, x_n\}$. In $\Gamma, x : a$ we require that $x \notin \text{dom}(\Gamma)$. We satisfy this restriction by renaming bound variables as needed. The rules for forming derivations of the typing judgment are presented in Figure 1.2. If t is a pre-term such that there exists a type a for which $\Gamma \vdash t : a$ holds, then we call t a *term*.

We can now think of encoding the simply-typed λ -calculus into our specification logic. This begins with the constructors i and arr for representing the base and arrow types. We also use the constructors app and abs for representing applications and abstractions. Using

$$\begin{aligned}
& \forall a, m. [\text{eval } (\text{abs } a \ m) \ (\text{abs } a \ m)] \\
& \forall m, a, r, n, v. [\text{eval } m \ (\text{abs } a \ r) \supset \text{eval } (r \ n) \ v \supset \text{eval } (\text{app } m \ n) \ v] \\
& \forall m, a, b, n. [\text{of } m \ (\text{arr } a \ b) \supset \text{of } n \ a \supset \text{of } (\text{app } m \ n) \ b] \\
& \forall a, r, b. [(\forall x. \text{of } x \ a \supset \text{of } (r \ x) \ b) \supset \text{of } (\text{abs } a \ r) \ (\text{arr } a \ b)]
\end{aligned}$$

Figure 1.3: A Horn clause-like encoding of evaluation and typing

a higher-order abstract syntax encoding there is no constructor for variables, and instead the *abs* constructor takes two arguments: 1) the type of the abstracted variable and 2) a specification logic abstraction representing the body. For example, the object language term $(\lambda x : i. (\lambda y : i. x))$ is denoted by $(\text{abs } i \ (\lambda x. \text{abs } i \ (\lambda y. x)))$ where these latter λ s are specification logic abstractions.

We introduce the specification logic predicates *eval* and *of* for representing evaluation and typing judgments respectively. Assuming a Horn clause-like specification logic, the rules for forming evaluation and typing judgments are encoded into the specification logic formulas shown in Figure 1.3. This specification uses various features of the specification logic which go beyond simple Horn clauses such as function application for realizing capture-avoiding substitution, universal quantification to avoid explicit side-conditions, and specification logic hypotheses for representing typing contexts. The complete details of this specification are presented in Chapter 2. For now it is sufficient to appreciate that the structural operational semantics description of the simply-typed λ -calculus can be encoded very directly into the specification logic. Moreover, a Prolog-like operational interpretation of proof search for the specification logic yields a prototype for our specification.

Returning to the original structural operational semantics description of the simply-typed λ -calculus for the moment, let us think of proving some global property of the system. One such property of interest is that evaluation preserves the type of a term, called the *type*

preservation property. Let us consider how such a property can be proved in an informal, mathematical setting. We might proceed by first showing the auxiliary properties of typing judgments that are contained in the following two lemmas.

Lemma 1.2.1. *If $\Gamma \vdash t : a$ and Γ' is a permutation of Γ , then $\Gamma' \vdash t : a$. Moreover, the derivations have the same height.*

Proof. The proof is by induction on the height of the derivation of $\Gamma \vdash t : a$. □

Lemma 1.2.2. *If $\Gamma, x : a \vdash t : b$ and $\Gamma \vdash s : a$ then $\Gamma \vdash t[x := s] : b$.*

Proof. The proof is by induction on the height of the derivation of $\Gamma, x : a \vdash t : b$. In the case where t is an abstraction we use Lemma 1.2.1 to permute the assumption $x : a$ to the end of the context. □

We can now state and prove the main property of interest.

Theorem 1.2.3. *If $t \Downarrow v$ and $\vdash t : a$ then $\vdash v : a$.*

Proof. The proof is by induction on the height of the derivation of $t \Downarrow v$.

Base case. If the derivation has height one then it must end with the following.

$$\overline{(\lambda x : a. r) \Downarrow (\lambda x : a. r)}$$

Then $t = v$ and the result is trivial.

Inductive case. If the derivation has height greater than one, then it must end with the following.

$$\frac{m \Downarrow (\lambda x : b. r) \quad r[x := n] \Downarrow v}{(m \ n) \Downarrow v}$$

Here $t = (m \ n)$ and we have shorter derivations of $m \Downarrow (\lambda x : b. r)$ and $r[x := n] \Downarrow v$. By assumption we know that $\vdash (m \ n) : a$ holds which means it has a derivation which must end with

$$\frac{\vdash m : c \rightarrow a \quad \vdash n : c}{\vdash (m \ n) : a}$$

for some type c . Now we can apply the inductive hypothesis to $m \Downarrow (\lambda x : b. r)$ and $\vdash m : c \rightarrow a$ to obtain a derivation of $\vdash (\lambda x : b. r) : c \rightarrow a$. Then it must be that $b = c$ and this derivation ends with the following rule.

$$\frac{x : b \vdash r : a}{\vdash (\lambda x : b. r) : b \rightarrow a}$$

By Lemma 1.2.2 we have a derivation of $\vdash r[x := n] : a$. Finally, we use the inductive hypothesis on $r[x := n] \Downarrow v$ and this typing judgment to conclude $\vdash v : a$. \square

Our objective is to carry out the style of reasoning described above in a formalized, computer-supported way. The framework that we will develop in this thesis will support such an ability. The key to doing this is designing a meta-logic for reasoning directly about the specification logic and, in this particular instance, the descriptions of evaluation and typing that have been encoded in it. The meta-logic that we will describe will allow the specification logic to be encoded as a definition in it, which then leads to the ability to reason, within the meta-logic, about the structure of specification logic derivations. Since these derivations have a close correspondence to the structural operational semantics derivations, a reasoning process very similar to that in Theorem 1.2.3 can be carried out within the meta-logic. Moreover, Lemmas 1.2.2 and 1.2.1 turn out to be instances of more general properties of the specification logic, and thus one can essentially obtain these results for free.

1.3 The Contributions of this Thesis

The framework that we are interested in developing in this thesis is characterized by a specification logic, a meta-logic, and an integration of these logics in a way that supports the two-level logic approach to reasoning. We shall base our specification logic on the intuitionistic theory of higher-order hereditary Harrop formulas [MNPS91]. This theory, which supports higher-order abstract syntax, underlies the λ Prolog programming language [NM88] and descriptions written in it can be animated using the Teyjus system [GHN⁺08,

Qi09]. Our focus in this work is on developing the meta-logic and the two-level logic approach to reasoning and on demonstrating their practical usefulness.

The starting point for our work will be a variant of the meta-logic called $FO\lambda^{\Delta\mathbb{N}}$ described by McDowell and Miller [MM00] that also supports the notion of higher-order abstract syntax. Further, our work will be inspired by the two-level logic approach to reasoning also described by McDowell and Miller [MM02]; from one perspective, we will mainly be strengthening the foundations of this approach and demonstrating how that it can be exploited effectively in practice. The specific realization of the two-level logic approach to reasoning in the work of McDowell and Miller is based on $FO\lambda^{\Delta\mathbb{N}}$ together with the same specification logic that we will be using in our framework. One of the most significant components of $FO\lambda^{\Delta\mathbb{N}}$ is a definition mechanism which allows one to reason about “closed” descriptions of systems. Thus, one can use the logic to perform case analysis-like reasoning about the behavior of an encoded system. This definition mechanism is based on earlier work on closed-world reasoning by many others, but most notably by Schroeder-Heister [SH93], Eriksson [Eri91], and Girard [Gir92]. The $FO\lambda^{\Delta\mathbb{N}}$ logic includes within it a mechanism for induction on natural numbers. Tiu extended this capability in the meta-logic Linc to a more general one that allows definitions themselves to be treated inductively and co-inductively. The co-inductive treatment was initially limited, but Tiu and Momigliano have subsequently developed the logic Linc⁻ which removes these limitations [TM09].

McDowell and Miller’s original meta-logic has also evolved in another way: the idea of generic judgments has been added to it to provide a better treatment of binding structure in higher-order abstract syntax representations than that afforded by the universal judgments originally used for this purpose. More specifically, Miller and Tiu introduced a new quantifier called ∇ which provides an elegant way to decompose higher-order abstract syntax representations by mapping term-level binding structure into a closely related proof-level binding structure. However, the original treatment of the ∇ quantifier interacted poorly with inductive and co-inductive reasoning. This has motivated Tiu to develop the logic LG^ω which refines the treatment of this new quantifier by including certain structural rules

for it [Tiu06].

This thesis makes contributions to the setting described above by further strengthening the meta-logic, by using it to develop an actual computer-based system for reasoning about specifications, and by demonstrating the benefit of the overall framework through actual reasoning applications. We discuss these contributions in more detail below.

1. We define a meta-logic called \mathcal{G} which improves on previous logics such as Linc and LG^ω . These other logics allow one to decompose higher-order abstract syntax by introducing ∇ -quantified variables into the structure of terms. These variables act like proof-level binders and allowed one to reason about the binding structure of objects without explicitly selecting variable names. However, these logics do not have any way to analyze the structure of terms with respect to the occurrences of such proof-level bound variables, a task which is common to almost all reasoning about binding structure. The meta-logic \mathcal{G} rectifies this situation by providing a generalization of the notion of equality which allows for exactly the type of analysis described. This generalized notion of equality behaves well with respect to definitions, induction, and co-induction. We establish consistency and more generally the cut-elimination property for \mathcal{G} , and we find that this meta-theory is a natural and pleasing extension of the meta-theory of previous logics. These contributions are the contents of Chapters 3 and 4.
2. The two-level logic approach had previously not been implemented and, hence, tested and the Linc logic had received only a partial implementation in a system called Bedwyr [BGM⁺06]. This thesis develops, for the first time, a complete realization of the reasoning component of the proposed framework. In particular, it develops a system called Abella that implements the meta-logic \mathcal{G} and supports the two-level logic approach to reasoning. Abella greatly extends the capabilities of Bedwyr by incorporating full inductive and co-inductive reasoning capabilities. Experiments with Abella have largely verified the effectiveness of the framework it supports, and this

aspect of our work has consequently contributed significantly to demonstrating the practicality of the two-level logic approach to reasoning. The discussion of Abella and its architecture is the content of Chapter 5.

3. We use Abella to expose a methodology of proof construction within the proposed framework which has a close correspondence with traditional pencil-and-paper proofs. We formally prove part of this correspondence through adequacy results for our two-level logic approach, and we demonstrate how to prove the full correspondence between the two-level logic approach to reasoning and traditional pencil-and-paper proofs. Finally, through concrete examples, we showcase the expressive power of the meta-logic \mathcal{G} and the practical benefits of the two-level logic approach to reasoning. These contributions are the contents of Chapters 6 and 7.

We note that the work described in this thesis has already contributed to the tools and techniques used by other researchers. The Abella system, that has been freely distributed, has been downloaded and experimented with by several researchers. It has also been used in at least one instance to verify a paper-and-pencil proof in a research paper [TM08].

1.4 Overview of the Thesis

In Chapter 2 we present the specification logic used in our proposed framework. We prove properties of this logic which make it a good basis for reasoning about object systems. We then encode the example of the simply-typed λ -calculus within the specification logic and prove the type preservation property via this encoding. The reasoning techniques used in this proof motivate some of the design of the meta-logic \mathcal{G} . We pick up on the specification logic again when we discuss the two-level approach to reasoning in Chapter 6.

Chapter 3 introduces the meta-logic \mathcal{G} and its various features including an extended notion of equality, a definition mechanism for encoding specifications, and induction and co-induction capabilities. We show how the extended notion of equality can be combined with the definition mechanism to produce a useful way of describing certain objects which

occur frequently when reasoning over higher-order abstract syntax descriptions. Finally, we provide examples which highlight the expressiveness of the new extended notion of equality. The contents of this chapter and the next also appear in [GMN08a, GMN09].

We develop the meta-theory of the meta-logic \mathcal{G} in Chapter 4. The primary result of this chapter is the proof of cut-elimination which we use to prove other useful properties relative to our meta-logic. We discover here that there is a nice (meta-theoretic) modularity to our use of an extended notion of equality as the basis for endowing \mathcal{G} with richer capabilities than the logics it builds on. In particular, we are able to reuse in this chapter much of the meta-theory already developed for Linc^- [TM09], thereby greatly reducing the effort that is needed for proving properties such as cut-elimination.

In Chapter 5 we describe the Abella system and its architecture. We describe the role of lemmas and lemma-like hypotheses during proof construction, and we show how the induction and co-induction rules of \mathcal{G} can be presented to the user in a very natural way.

Chapter 6 brings together the specification logic and the meta-logic to develop the two-level logic approach to reasoning. In particular, this chapter describes how the specification logic can be embedded in the meta-logic and what benefit this has towards formalizing the properties of the specification logic. We reconsider the example of the simply-typed λ -calculus and using the two-level logic approach to reasoning we provide a very short and elegant proof of type preservation. Finally, we show that our encoding of the specification logic is adequate subject to some minor conditions.

Using the two-level logic approach to reasoning and its embodiment in the Abella theorem prover we present larger applications of our framework in Chapter 7. These applications are intended to highlight the strengths and weaknesses of the two-level logic approach to reasoning. They include examples such as the POPLmark challenge [ABF⁺05] and Girard's proof of strong normalization for the simply-typed λ -calculus.

In Chapter 8 we compare our framework against other approaches to specifying, prototyping, and reasoning about computational systems with binding.

We conclude this thesis in Chapter 9 and discuss various avenues of future work. These

range from foundational extensions which would increase the expressive power of the meta-logic to more implementation oriented extensions which would better facilitate the reasoning process.

Chapter 2

A Logic for Specifying Computational Systems

The primary requirement of a specification logic within the framework that we want to develop is that it allow for a transparent encoding of the kinds of formal systems that are of interest to us. In particular, such an encoding should cover both the objects manipulated within the formal system and the rules by which they are manipulated. In the context of our work, we are particularly concerned with the representation of objects that incorporate a variable binding structure. A logically precise encoding of such structure plays an important role in the overall treatment of the relevant computational systems. An encoding that has this character usually requires the treatment of concepts related to binding, such as equality under bound variable renaming and capture-avoiding substitution. If these aspects are not dealt with in a systematic way within the specification logic, they can overwhelm the process of constructing encodings and can make the subsequent process of reasoning about specifications unnecessarily complex. We therefore seek a specification logic which incorporates a flexible and sophisticated treatment of variable binding structure and which also builds in the related binding notions.

In this chapter we introduce the specification logic of second-order hereditary Harrop formulas, abbreviated hH^2 . This logic is essentially a restriction of the logic of higher-order hereditary Harrop formulas [MNPS91] that underlies the language λ Prolog [NM88]. The hH^2 logic can be seen as an extension of the Horn clause logic, the logic that underlies Prolog, with devices for representing, examining, and manipulating objects with binding structure. In particular, hH^2 allows for a higher-order abstract syntax representation of objects with binding structure [MN87, PE88]. Thus issues of variable renaming and capture-avoiding substitution are taken care of once and for all in the specification

logic, leaving particular specifications free to focus on the more essential aspects of the system they encode. Furthermore, like the logic of higher-order hereditary Harrop formulas that it derives from, hH^2 admits an operational semantics which allows specifications to be animated automatically thus yielding quick prototypes of the computational systems they encode.

In this chapter we formally define the hH^2 logic, describe its operational semantics, state and prove properties of the logic, and demonstrate its use through a concrete example.

2.1 The Syntax and Semantics of the Logic

Following Church [Chu40], terms in hH^2 are constructed using abstraction and application from constants and bound variables. All terms are typed using a monomorphic typing system. The provability relation concerns well-formed terms of the distinguished type o that are also called formulas. Logic is introduced by including special constants representing the propositional connectives \top , \wedge , \vee , \supset and, for every type τ that does not contain o , the constants \forall_τ and \exists_τ of type $(\tau \rightarrow o) \rightarrow o$. We do not allow any other constants or variables to have a type containing the type o . The binary propositional connectives are written as usual in infix form and the expressions $\forall_\tau x.B$ and $\exists_\tau x.B$ abbreviate the formulas $\forall_\tau \lambda x.B$ and $\exists_\tau \lambda x.B$, respectively. Type subscripts will be omitted from quantified formulas when they can be inferred from the context or are not important to the discussion. We also use a shorthand for iterated quantification: if Q is a quantifier, we will often abbreviate $Qx_1 \dots Qx_n.P$ to $Qx_1, \dots, x_n.P$ or simply $Q\vec{x}.P$. We consider the scope of λ -binders (and therefore quantifiers) as extending as far right as possible. We further assume that \supset is right associative and has lower precedence than \wedge and \vee . For example, $\forall x.t_1 \supset t_2 \supset t_3 \wedge t_4$ should be read as $\forall x.(t_1 \supset (t_2 \supset (t_3 \wedge t_4)))$.

We restrict our attention to two classes of formulas in hH^2 described by the following

grammar.

$$\begin{aligned}
 G &::= \top \mid A \mid A \supset G \mid \forall_{\tau}x.G \mid \exists_{\tau}x.G \mid G \wedge G \mid G \vee G \\
 D &::= A \mid G \supset D \mid \forall_{\tau}x.D
 \end{aligned}$$

Here A denotes an atomic formula. The formulas denoted by G are called *goals* and represent the conclusions we can infer in the logic. A notable restriction on implication in goal formulas is that the left hand side must be an atomic formula. Formulas denoted by D are called *definite clauses* and represent the hypotheses we can assume in the logic. Notice that disjunctions and existentials are not allowed in definite formulas because they represent indefinite knowledge. For simplicity, we also disallow conjunction, but the effect of conjunctions can be recovered by using a set of clauses in place of a single clause. The order of a formula is the depth of implications which are nested to the left of other implications. Our restriction on implication means goal formulas are at most first-order and definite clauses are at most second-order. It is precisely this restriction which carves out the logic of second-order hereditary Harrop formulas from the larger logic of higher-order hereditary Harrop formulas. Finally, by using logical equivalences we can percolate universal quantifiers to the top, to rewrite all definite clauses to be of the form $\forall x_1 \dots \forall x_n.(G_1 \supset \dots \supset G_m \supset A)$ where n and m may both be zero. In the future we will assume all definite clauses are in this form.

The semantics of hH^2 are formalized by means of a proof-theoretic presentation of what it means for a goal to follow from a set of definite clauses. Specifically, we will be concerned with the derivation of *sequents* of the form $\Sigma : \Delta \vdash G$ where Δ is a list of D -formulas, G is a G -formula, and Σ is a set of variables called eigenvariables. For such a sequent to be well-formed, we require that the formulas in $\Delta \cup \{G\}$ must be constructed using only the logical and non-logical constants of the language and the eigenvariables in Σ . This well-formedness condition is guaranteed for every sequent considered in a derivation by ensuring that we try to construct derivations only for well-formed ones at the top-level and by the use of typing judgments of the form $\Sigma \vdash t : \tau$ in rules that introduce new terms when these

rules are interpreted in a proof search direction. The meaning of this typing judgment, that we do not explicitly formalize here, is the following: for it to hold, the term t must have the type τ and it must also be constructed using only the non-logical constants and the eigenvariables in Σ .

The rules for constructing proofs for such sequents are presented in Figure 2.1. The GENERIC rule introduces an eigenvariable when read in a proof search direction. There is a freshness side-condition associated with this eigenvariable: c must not already be in Σ . Note that for this to be possible, we must assume that there is an unlimited supply of eigenvariables of each type. In the INSTANCE rule t is required to be a term such that $\Sigma \vdash t : \tau$ holds. Similarly, in the BACKCHAIN rule for each term $t_i \in \vec{t}$ we must have $\Sigma \vdash t_i : \tau_i$ where τ_i is the type of the quantified variable x_i . An important property to note about these rules is that if we use them to search for a proof of the sequent $\Delta \vdash G$, then all the intermediate sequents that we will encounter will have the form $\Delta, \mathcal{L} \vdash G'$ for some G -formula G' and some list of atomic formulas \mathcal{L} . Thus the initial context Δ is global, and only atomic formulas are added to the context during proof construction.

In presenting sequents in later parts of this thesis, we shall occasionally omit writing the signature. We will do this only when either the identity of the signature is irrelevant to the discussion or when it can be inferred from the context.

The rules of hH^2 admit a simple proof search procedure: given a sequent $\Delta \vdash G$ we decompose the goal G until we reach an atomic formula at which point we backchain and attempt to prove the resulting goals. This is, in fact, a manifestation of the *uniform proofs* property that hH^2 inherits from the parent logic of higher-order hereditary Harrop formulas [MNPS91]. The resulting procedure is non-deterministic since we have a choice when the goal is a disjunction, an existential, or an atomic formula (we can choose which clause to backchain on). The non-determinism induced by existentials can be handled using the standard notion of instantiable variables and unification while the non-determinism of the OR and BACKCHAIN rules can be handled using depth-first search complemented with backtracking. Computations described by hH^2 are included within those corresponding to

$$\begin{array}{c}
\frac{}{\Sigma : \Delta \vdash \top} \text{TRUE} \\
\\
\frac{\Sigma : \Delta \vdash G_1}{\Sigma : \Delta \vdash G_1 \vee G_2} \text{OR}_1 \quad \frac{\Sigma : \Delta \vdash G_2}{\Sigma : \Delta \vdash G_1 \vee G_2} \text{OR}_2 \\
\\
\frac{\Sigma : \Delta \vdash G_1 \quad \Sigma : \Delta \vdash G_2}{\Sigma : \Delta \vdash G_1 \wedge G_2} \text{AND} \quad \frac{\Sigma : \Delta \vdash G[t/x]}{\Sigma : \Delta \vdash \exists x.G} \text{INSTANCE} \\
\\
\frac{\Sigma : \Delta, A \vdash G}{\Sigma : \Delta \vdash A \supset G} \text{AUGMENT} \quad \frac{\Sigma \cup \{c:\tau\} : \Delta \vdash G[c/x]}{\Sigma : \Delta \vdash \forall_{\tau}x.G} \text{GENERIC} \\
\\
\frac{\Sigma : \Delta \vdash G_1[\vec{t}/\vec{x}] \quad \cdots \quad \Sigma : \Delta \vdash G_m[\vec{t}/\vec{x}]}{\Sigma : \Delta \vdash A} \text{BACKCHAIN}
\end{array}$$

where $\forall \vec{x}.(G_1 \supset \cdots \supset G_m \supset A') \in \Delta$ and $A'[\vec{t}/\vec{x}] = A$

Figure 2.1: Derivation rules for the hH^2 logic

λ Prolog and can therefore be compiled and executed efficiently, *e.g.*, by the Teyjus system [GHN⁺08, Qi09].

2.2 Properties of the Specification Logic

We will eventually encode object logic judgments into specification logic judgments. By doing this, we enable ourselves to use properties of the specification logic in proving properties of the object logic. Therefore in this section we enumerate the various properties of the hH^2 logic which may be useful in such reasoning. The proofs of these properties will be based on induction over the *height* of a derivation, a notion we define now.

Definition 2.2.1. *The height of a derivation Π , denoted by $\text{ht}(\Pi)$, is 1 if Π has no premise derivations and is $\max\{\text{ht}(\Pi_i) + 1\}_{i \in 1..n}$ if Π has the premise derivations $\{\Pi_i\}_{i \in 1..n}$.*

The *monotonicity property* of hH^2 states that the eigenvariables and the context of a sequent can always be expanded while preserving provability.

Lemma 2.2.2. *Let $\Sigma : \Delta \vdash G$ be a well-formed sequent, let Δ' be a list of definite clauses such that $\Delta \subseteq \Delta'$, and let Σ' be a set of eigenvariables such that $\Sigma \subseteq \Sigma'$ and Σ' contains all the eigenvariables of Δ' . If $\Sigma : \Delta \vdash G$ has a derivation then $\Sigma' : \Delta' \vdash G$ has a derivation. Moreover, the height of the derivation does not increase.*

Proof. Induction on the height of the derivation of $\Sigma : \Delta \vdash G$. □

The *instantiation property* states that a eigenvariable c which arises from a use of the GENERIC rule can always be instantiated with a particular value while preserving provability. As a result, our use of eigenvariables to denote universal quantification in hH^2 is well justified.

Lemma 2.2.3. *Let c be a variable not in Σ . If $\Sigma \cup \{c : \tau\} : \Delta \vdash G$ has a derivation then for all terms t such that $\Sigma \vdash t : \tau$ there is a derivation of $\Sigma : \Delta[t/c] \vdash G[t/c]$. Moreover, the height of the derivation does not increase.*

Proof. Induction on the height of the derivation of $\Sigma \cup \{c : \tau\} : \Delta \vdash G$. □

Finally, the *cut admissibility property* says that the assumption of an atomic formula can be discharged if the atomic formula is itself provable.

Lemma 2.2.4. *If $\Sigma : \Delta, A \vdash G$ and $\Sigma : \Delta \vdash A$ then $\Sigma : \Delta \vdash G$.*

Proof. Induction on the height of the derivation of $\Sigma : \Delta, A \vdash G$. There are two interesting cases. The first case is when G is $A' \supset G'$ in which case we must apply the monotonicity property to move from $\Sigma : \Delta, A, A' \vdash G$ to $\Sigma : \Delta, A', A \vdash G$. The other case is when the BACKCHAIN rule selects A , in which case the derivation of $\Sigma : \Delta \vdash A$ can be substituted. □

2.3 Example Encoding in the Specification Logic

We now take the example of evaluation and typing for the simply-typed λ -calculus from Section 1.2, and we encode it into the specification logic. We introduce the specification logic types tp and tm for representing types and pre-terms respectively in the simply-typed λ -calculus. Types in the simply-typed λ -calculus will be mapped to specification logic terms constructed from the constants i and arr of types tp and $tp \rightarrow tp \rightarrow tp$, respectively. Pre-terms in the simply-typed λ -calculus will be mapped to specification logic terms constructed from the constants app and abs of types $tm \rightarrow tm \rightarrow tm$ and $tp \rightarrow (tm \rightarrow tm) \rightarrow tm$, respectively. Notice that the second argument of abs is expected to be an abstraction over tm in the specification logic. Finally, we will have two constants of and $eval$ of types $tm \rightarrow tp \rightarrow o$ and $tm \rightarrow tm \rightarrow o$, respectively, which denote typing and evaluation, respectively. The clauses for these predicates are presented in Figure 2.2. Here and in the future we use the convention that tokens given by capital letters denote variables that are implicitly universally quantified over the entire formula. In the second clause for evaluation, R is an abstraction in the specification logic and thus the built-in notion of β -reduction means that $(R N)$ realizes capture-avoiding substitution of N in for the bound variable in R . For the typing judgment, we do not keep an explicit context of typing assumptions, instead relying on the specification logic context. This is reflected in the rule for typing abstractions where we use the \forall quantifier to create a fresh eigenvariable and we assume that this eigenvariable has the proper type while we derive a typing assignment for the body of the abstraction. In this way, we avoid having an explicit base case for typing. Next, when we reason about this specification we will be able to exploit this encoding of the typing context.

Using this encoding, we can now repeat the proof of type preservation and leverage on the properties we have shown of the hH^2 logic. Let Δ be the clauses from Figure 2.2.

Theorem 2.3.1. *If $\Delta \vdash eval\ e\ v$ holds and $\Delta \vdash of\ e\ t$ holds then $\Delta \vdash of\ v\ t$ holds.*

Proof. By induction on the height of the derivation of $\Delta \vdash eval\ e\ v$. We proceed by cases on

$$\begin{aligned}
& eval (abs A M) (abs A M) \\
& eval M (abs A R) \supset eval (R N) V \supset eval (app M N) V \\
& of M (arr A B) \supset of N A \supset of (app M N) B \\
& (\forall x.of x A \supset of (R x) B) \supset of (abs A R) (arr A B)
\end{aligned}$$

Figure 2.2: hH^2 specification of evaluation and typing

the derivation of $\Delta \vdash eval e v$. This judgment must have been derived by backchaining on one of the clauses for *eval*. If it was by the first clause, then $e = v$ and the case is complete. Otherwise it was by the second clause so e must be $(app m n)$ for some m and n and we have shorter derivations of $\Delta \vdash eval m (abs a r)$ and $\Delta \vdash eval (r n) v$ for some a and r . By similarly examining the derivation $\Delta \vdash of (app m n) t$ we must have derivations of $\Delta \vdash of m (arr b t)$ and $\Delta \vdash of n b$ for some b . Applying the inductive hypothesis to $\Delta \vdash eval m (abs a r)$ and $\Delta \vdash of m (arr b t)$ we have $\Delta \vdash of (abs a r) (arr b t)$. This derivation could only result if $a = b$ and we have a derivation of $\Delta \vdash \forall x[of x a \supset of (r x) t]$ and thus a derivation of $\Delta, of c a \vdash of (r c) t$ for some eigenvariable c . Now we can apply the instantiation property of our specification logic to get a derivation of $\Delta, of n a \vdash of (r n) t$. Next we apply the cut property with our derivation of $\Delta \vdash of n a$ to get $\Delta \vdash of (r n) t$. Finally, we apply the inductive hypothesis again to $\Delta \vdash of (r n) t$ and $\Delta \vdash eval (r n) v$ to get $\Delta \vdash of v t$ which completes the proof. \square

It is important to note in this proof that we did not have to prove a type substitution property for the object logic. Instead, the object logic inherited this property from the more general instantiation and cut properties of the specification logic. Also, induction over the height of specification logic derivations corresponded with induction over the height of object logic derivations. Thus we can reason about computational systems through their encoding in the specification logic with little overhead cost.

2.4 Adequacy of Encodings in the Specification Logic

A tacit assumption in the example we considered in the previous section is that the specification of pre-terms, types, typing, and evaluation are all faithful representations of the corresponding concepts in the object logic. This kind of property of encodings is referred to as the *adequacy* property. We must, of course, prove such a property before we can derive benefit from it. To do this, we need to prove that there is a bijection between components of the object logic and their specification logic representations and that this bijection preserves properties of relevance in the two systems. With specific reference to the example encoding we have considered, we have to show that each object in the simply-typed λ -calculus has a unique representation in the specification logic, and each representation in the specification logic corresponds to a unique object in the simply-typed λ -calculus. We show below how such arguments are typically carried out.

To simplify the argument we will assume an implicit mapping between bound variables in the object language and bound variables in the specification language, and between free variables in the object language and eigenvariables in the specification language. A more rigorous treatment of adequacy would make this mapping explicit [Fel89].

We define the bijections ϕ_{tp} , ϕ_{tm} , ϕ_{eval} , ϕ_{ctx} , and ϕ_{of} which are used to map types, pre-terms, evaluation judgments, typing contexts, and typing judgments in the simply-typed λ -calculus to their corresponding representations in the specification logic. We will omit the subscripts on ϕ when they can be inferred from context. The proofs that these mappings are bijective are always by straightforward induction on the size of terms or strong induction on the height of derivations.

Types in the simply-typed λ -calculus map to terms of type tp in the specification logic. We formalize this mapping as follows.

$$\phi(i) = i \qquad \phi(a \rightarrow b) = arr \phi(a) \phi(b)$$

This function is clearly a bijection.

Next we define the mapping between α -equivalence classes of pre-terms in the object

logic and terms of type tm in the specification logic.

$$\phi(x) = x \quad \phi(m \ n) = \text{app } \phi(m) \ \phi(n) \quad \phi(\lambda x : a. r) = \text{abs } \phi(a) \ (\lambda x. \phi(r))$$

In the last rule for this mapping note that the λ within the ϕ is that of the simply-typed λ -calculus while the one outside of ϕ is from the specification logic. This mapping is clearly bijective under the assumption that α -convertible terms in the specification logic are considered to be identical.

Let Δ be the clauses from Figure 2.2. Then derivations of evaluation judgments in the simply typed λ -calculus correspond to derivations of the sequent $\Delta \vdash \text{eval } e \ v$ in the specification logic as follows. First consider the translation of evaluation for abstractions:

$$\begin{aligned} \phi \left(\frac{}{\lambda x : a. t \Downarrow \lambda x : a. t} \right) &= \frac{}{\Delta \vdash \text{eval } \phi(\lambda x : a. t) \ \phi(\lambda x : a. t)} \\ &= \frac{}{\Delta \vdash \text{eval } (\text{abs } \phi(a) \ (\lambda x. \phi(t))) \ (\text{abs } \phi(a) \ (\lambda x. \phi(t)))} \end{aligned}$$

Here and in the future we propagate the mapping ϕ to make it clear that the specification logic inference rules are well-formed. In this case, the right-hand inference rule an instance of the BACKCHAIN rule over the clause for evaluating abstractions.

The translation for evaluations of applications is the following.

$$\begin{aligned} \phi \left(\frac{m \Downarrow \lambda x : a. r \quad r[x := n] \Downarrow v}{m \ n \Downarrow v} \right) &= \frac{\phi \left(\frac{m \Downarrow \lambda x : a. r}{m \Downarrow \lambda x : a. r} \right) \quad \phi \left(\frac{r[x := n] \Downarrow v}{r[x := n] \Downarrow v} \right)}{\Delta \vdash \text{eval } \phi(m \ n) \ \phi(v)} \\ &= \frac{\frac{\Delta \vdash \text{eval } \phi(m) \ (\text{abs } \phi(a) \ (\lambda x. \phi(r)))}{\phi(\cdot)} \quad \frac{\Delta \vdash \text{eval } (\phi(r)[\phi(n)/x]) \ \phi(v)}{\phi(\cdot)}}{\Delta \vdash \text{eval } (\text{app } \phi(m) \ \phi(n)) \ \phi(v)} \end{aligned}$$

In the final formula, we make use of the automatic β -conversion in the specification logic where $(\lambda x. \phi(r)) \ \phi(n) = \phi(r)[\phi(n)/x]$, and we use the following compositional property of

the bijection for terms.

$$\phi(r[x := n]) = \phi(r)[\phi(n)/x]$$

This equation relates the substitution of the simply-typed λ -calculus on the left with the substitution in the specification logic on the right. The proof of this equality is by induction on the structure of r . Thus the inference rule on the right-hand side above is a proper instance of the BACKCHAIN rule over the clause for evaluating applications. The inverse of the ϕ mapping is defined in the natural way and thus ϕ is a bijection.

Finally, we look at derivations of typing judgments in the simply-typed λ -calculus and we map these to derivations of sequents of the form $\Delta, \mathcal{L} \vdash \text{of } e \ t$ where \mathcal{L} is a list of atomic formulas of the form $\text{of } x_1 \ a_1, \dots, \text{of } x_k \ a_k$ where each x_i is a unique eigenvariable. We first define the following bijection between a list of typing assumptions Γ from the simply-typed λ -calculus and a list of atomic formulas of the form described for \mathcal{L} .

$$\phi(x_1 : a_1, \dots, x_k : a_k) = \text{of } x_1 \ \phi(a_1), \dots, \text{of } x_k \ \phi(a_k)$$

Given this, we can define the mapping for typing variables as follows.

$$\phi\left(\frac{}{\Gamma \vdash x_i : a_i}\right) = \frac{}{\Delta, \phi(\Gamma) \vdash \text{of } x_i \ \phi(a_i)}$$

If the typing derivation within the ϕ is correct then it must be that $x_i : a_i \in \Gamma$. Thus the right-hand side is an instance of the BACKCHAIN rule on the clause $\text{of } x_i \ \phi(a_i)$ which is in $\phi(\Gamma)$.

The typing rule for applications is mapped in the expected way:

$$\begin{aligned} & \phi\left(\frac{\begin{array}{c} \vdots \\ \Gamma \vdash m : a \rightarrow b \quad \Gamma \vdash n : a \\ \vdots \end{array}}{\Gamma \vdash m \ n : b}\right) \\ &= \frac{\phi\left(\frac{\begin{array}{c} \vdots \\ \Gamma \vdash m : a \rightarrow b \\ \vdots \end{array}\right) \quad \phi\left(\frac{\begin{array}{c} \vdots \\ \Gamma \vdash n : a \\ \vdots \end{array}\right)}{\Delta, \phi(\Gamma) \vdash \text{of } \phi(m \ n) \ \phi(b)}}{\Delta, \phi(\Gamma) \vdash \text{of } \phi(m) \ (\text{arr } \phi(a) \ \phi(b)) \quad \Delta, \phi(\Gamma) \vdash \text{of } \phi(n) \ \phi(a)} \\ &= \frac{\phi(\cdot) \quad \phi(\cdot)}{\Delta, \phi(\Gamma) \vdash \text{of } (\text{app } \phi(m) \ \phi(n)) \ \phi(b)} \end{aligned}$$

For mapping the abstraction typing rule, we need to be mindful of the variable naming restriction and how this is realized in the specification logic. Suppose we want to define the following mapping.

$$\phi \left(\frac{\Gamma, x : a \vdash r : b}{\Gamma \vdash (\lambda x : a. r) : a \rightarrow b} \right)$$

Here we assume that x does not appear in Γ so that the naming restriction is satisfied. We map this to the following specification logic derivation.

$$\frac{\frac{\frac{\phi(\cdot)}{\Delta, \phi(\Gamma), \text{of } x \phi(a) \vdash \text{of } \phi(r) \phi(b)}}{\Delta, \phi(\Gamma) \vdash \text{of } x \phi(a) \supset \text{of } \phi(r) \phi(b)} \text{ AUGMENT}}{\Delta, \phi(\Gamma) \vdash \forall x. [\text{of } x \phi(a) \supset \text{of } ((\lambda x. \phi(r)) x) \phi(b)]} \text{ GENERIC}}{\Delta, \phi(\Gamma) \vdash \text{of } (\text{abs } \phi(a) (\lambda x. \phi(r))) (\text{arr } \phi(a) \phi(b))} \text{ BACKCHAIN}$$

In the **GENERIC** rule we overload notation to let x be the eigenvariable we select. Since it does not appear in Γ it will not appear in $\phi(\Gamma)$, and thus the freshness side-condition on the **GENERIC** rule is satisfied. In fact, the naming restriction in the object logic matches up with the freshness side-condition in the specification logic exactly as needed.

The inverse of the ϕ mapping for typing judgments can be defined in the expected way, and thus ϕ is a bijection. This concludes the proof of adequacy for our specification. In the future we will omit such arguments since our specifications are often transparent encodings of the systems they represent.

Chapter 3

A Logic for Reasoning About Specifications

In this chapter we present the meta-logic \mathcal{G} . This logic allows for encoding descriptions of computational systems and for reasoning over those descriptions. The logic includes traditional reasoning devices such as case analysis, induction, and co-induction as well as new devices specifically designed for working with higher-order abstract syntax.

The relevant history of \mathcal{G} begins with the meta-logic $FO\lambda^{\Delta\mathbb{N}}$ developed by McDowell and Miller for the purposes of inductive reasoning over higher-order abstract syntax descriptions [MM02, MM00]. This logic contains a definition mechanism which allows one to specify and reason about closed-world descriptions, *i.e.*, allows one to form judgments and to perform case analysis on them. This definition mechanism is based on earlier work on closed-world reasoning by many others, but most notably by Schroeder-Heister [SH93], Eriksson [Eri91], and Girard [Gir92]. The primary contribution of $FO\lambda^{\Delta\mathbb{N}}$ was the recognition that definitions provided a way of encoding higher-order abstract syntax descriptions in such a way that does not conflict with inductive reasoning. In particular, $FO\lambda^{\Delta\mathbb{N}}$ allowed for natural number induction, and so many reasoning tasks could be naturally encoded. More recently, Tiu [Tiu04] developed the meta-logic Linc which extends the mechanism of definitions to integrate notions of generalized induction and co-induction over the structure of definitions. These more general notions are present in \mathcal{G} as well.

Another central advancement in the development of logics for reasoning over higher-order abstract syntax descriptions was the recognition that one needed a way to reflect the binding structure of terms into the structure of proofs. This was realized in earlier logics by using universal judgments. However, this kind of correspondence was always an uneasy one and the mismatch became explicit when it was necessary to use case analysis arguments over

binding structure as must be done, for example, in bisimilarity proofs associated with π -calculus models of concurrent systems. The desire to provide a logically precise and cleaner treatment led to the development of the ∇ -quantifier and the associated generic judgment by Miller and Tiu in the meta-logic $FO\lambda^{\Delta\nabla}$ [MT05]. Tiu later refined this notion in the meta-logic LG^ω so that ∇ -quantifier behaved well with respect to inductive reasoning [Tiu06]. This interpretation of the ∇ -quantifier is present in \mathcal{G} , and in this context it can be understood as quantifying over fresh names.

The meta-logic \mathcal{G} is a continuation of the research surrounding inductive reasoning and higher-order abstract syntax descriptions. In particular, it extends the notion of equality in the logic to one which can describe the binding structure of terms relative to the proof context in which they occur. This turns out to be essential to describing the structure of terms which are generated during inductive reasoning over higher-order abstract syntax descriptions. Moreover, \mathcal{G} identifies how this extended notion of equality can be integrated with the definition mechanism to allow a succinct description of such objects.

The presentation of \mathcal{G} is divided into three parts. First, Section 3.1 contains the core of the logic including generic quantification. Then Section 3.2 introduces the extended notion of equality known as *nominal abstraction* and rules for treating this notion within the logic. Finally, Section 3.3 presents rules for treating fixed-points in the logic including mechanisms for induction and co-induction. Although the logical features of \mathcal{G} are described in their entirety in the first three sections, it is sometimes convenient to use an alternative presentation for fixed-point definitions. This form, which uses patterns to distinguish different cases in the structure of the atom being defined, is introduced in Section 3.4 and is elaborated as an interpretation of the basic form of definitions that uses nominal abstractions explicitly. Rules for treating this alternative form of fixed-points are presented and proven to be admissible. Finally, Section 3.5 provides some small examples to illustrate the expressive power of the logic.

3.1 A Logic with Generic Quantification

In this section we present the core logic underlying \mathcal{G} . This logic is obtained by extending an intuitionistic and predicative subset of Church’s Simple Theory of Types with a treatment of generic judgments. The encoding of generic judgments is based on the quantifier called ∇ (pronounced nabla) introduced by Miller and Tiu [MT05] and further includes the structural rules associated with this quantifier in the logic LG^ω described by Tiu [Tiu06].

3.1.1 The Basic Syntax

Following Church [Chu40], terms are constructed from constants and variables using abstraction and application. All terms are assigned types using a monomorphic typing system; these types also constrain the set of well-formed expressions in the expected way. The collection of types includes o , a type that corresponds to propositions. Well-formed terms of this type are also called formulas. Two terms are considered to be equal if one can be obtained from the other by a sequence of applications of the α -, β - and η -conversion rules, *i.e.*, the λ -conversion rules. This notion of equality is henceforth assumed implicitly wherever there is a need to compare terms. Logic is introduced by including special constants representing the propositional connectives \top , \perp , \wedge , \vee , \supset and, for every type τ that does not contain o , the constants \forall_τ and \exists_τ of type $(\tau \rightarrow o) \rightarrow o$. The binary propositional connectives are written as usual in infix form and the expressions $\forall_\tau x.B$ and $\exists_\tau x.B$ abbreviate the formulas $\forall_\tau \lambda x.B$ and $\exists_\tau \lambda x.B$, respectively. Type subscripts will be omitted from quantified formulas when they can be inferred from the context or are not important to the discussion. We also use a shorthand for iterated quantification: if \mathcal{Q} is a quantifier, we will often abbreviate $\mathcal{Q}x_1 \dots \mathcal{Q}x_n.P$ to $\mathcal{Q}x_1, \dots, x_n.P$ or simply $\mathcal{Q}\vec{x}.P$. We consider the scope of λ -binders (and therefore quantifiers) as extending as far right as possible. We further assume that \supset is right associative and has lower precedence than \wedge and \vee . For example, $\forall x.t_1 \supset t_2 \supset t_3 \wedge t_4$ should be read as $\forall x.(t_1 \supset (t_2 \supset (t_3 \wedge t_4)))$.

The usual inference rules for the universal quantifier can be seen as equating it to the

conjunction of all of its instances: that is, this quantifier is treated extensionally. There are several situations where one wishes to treat an expression such as “ $B(x)$ holds for all x ” as a statement about the existence of a uniform argument for every instance rather than the truth of a particular property for each instance [MT05]; such situations typically arise when one is reasoning about the binding structure of formal objects represented using the *λ-tree syntax* [Mil00] version of *higher-order abstract syntax* [PE88]. The ∇ -quantifier serves to encode judgments that have this kind of a “generic” property associated with them. Syntactically, this quantifier corresponds to including a constant ∇_τ of type $(\tau \rightarrow o) \rightarrow o$ for each type τ not containing o .¹ As with the other quantifiers, $\nabla_\tau x.B$ abbreviates $\nabla_\tau \lambda x.B$ and the type subscripts are often suppressed for readability.

3.1.2 Generic Judgments and ∇ -quantification

Sequents in intuitionistic logic can be written as

$$\Sigma : B_1, \dots, B_n \longrightarrow B_0 \quad (n \geq 0)$$

where Σ is the “global signature” for the sequent that contains the *eigenvariables* (*i.e.*, variables associated to the $\exists\mathcal{L}$ and $\forall\mathcal{R}$ inference rules) relevant to the sequent proof. We shall think of Σ in this prefix position as an operator that binds each of the variables it contains and that has the rest of the sequent as its scope. To treat the ∇ -quantifier, the $FO\lambda^{\Delta\nabla}$ logic [MT05] extends the notion of a judgment from just a formula to a formula paired with a “local signature.” Thus, sequents within this logic are written more elaborately as

$$\Sigma : \sigma_1 \triangleright B_1, \dots, \sigma_n \triangleright B_n \longrightarrow \sigma_0 \triangleright B_0,$$

where each $\sigma_0, \dots, \sigma_n$ is a list of variables that are bound locally in the formula adjacent to it. Such local signatures correspond to a proof-level encoding of binding that is expressed within formulas through the ∇ -quantifier. In particular, the judgment $x_1, \dots, x_n \triangleright B$ and

¹ We may choose to allow ∇ -quantification at fewer types in particular applications; such a restriction may be useful in adequacy arguments for reasons we discuss later.

the formula $\nabla x_1 \cdots \nabla x_n.B$ for $n \geq 0$ have the same proof-theoretic force. In keeping with this observation, we shall refer to a judgment of the form $\sigma \triangleright B$ as a *generic judgment*.

As part of a generalization of sequents that bases them on generic judgments rather than on formulas, we need to define when two such judgments are equal: this is necessary for describing at least the initial and cut inference rules. The $FO\lambda^{\Delta\nabla}$ logic [MT05] uses a simple form of equality for this purpose. It deems two generic judgments of the form $x_1, \dots, x_n \triangleright B$ and $y_1, \dots, y_m \triangleright C$ to be equal exactly when the λ -terms $\lambda x_1 \dots \lambda x_n.B$ and $\lambda y_1 \dots \lambda y_m.C$ are λ -convertible; notice that this necessarily implies that $n = m$. An equality notion is also needed in formulating an induction rule. Unfortunately, the simple form of equality present in $FO\lambda^{\Delta\nabla}$ leads to a rather weak version of such a rule. To overcome this difficulty, Tiu proposed the addition to the logic of two natural “structural” identities between generic judgments. These identities are the ∇ -*strengthening rule* $\nabla x.F = F$, provided x is not free in F , and the ∇ -*exchange rule* $\nabla x \nabla y.F = \nabla y \nabla x.F$. In its essence, the LG^ω proof system [Tiu06] is obtained from $FO\lambda^{\Delta\nabla}$ by strengthening its notion of equality based on λ -conversion through the addition of these two structural rules for ∇ .

The move from the weaker logic $FO\lambda^{\Delta\nabla}$ to the stronger logic LG^ω involves an ontological commitment and has a proof-theoretic consequence.

At the ontological level, the strengthening rule implies that every type at which one is willing to use ∇ -quantification is non-empty and, in fact, contains an unbounded number of members. For example, the formula $\exists_\tau x.\top$ is always provable, even if there are no closed terms of type τ because this formula is equivalent to $\nabla_\tau y.\exists_\tau x.\top$, which is provable. Similarly, for any given $n \geq 1$, the following formula is provable

$$\exists_\tau x_1 \dots \exists_\tau x_n. \left[\bigwedge_{1 \leq i, j \leq n, i \neq j} x_i \neq x_j \right].$$

At the proof-theoretic level, an acceptance of the strengthening and exchange rules means that the length of a local context and the order of variables within it are unimportant. For example, a sequent that contains the generic judgments $x_1, \dots, x_n \triangleright B$ and $y_1, \dots, y_m \triangleright C$ can be rewritten (assuming $n \geq m$) using α -conversion and strengthening into the judgments

$z_1, \dots, z_n \triangleright B'$ and $z_1, \dots, z_n \triangleright C'$ where B' and C' are equal to B and C modulo variable renamings. In this fashion, all local bindings in a sequent can be made to involve the same variables, and, hence, the local bindings can be seen as a global binding over a sequent that contains formulas and not generic judgments. The resulting sequent-level variable bindings will be represented by specially designated *nominal constants*. Notice, however, that each of these nominal “constants” has as its scope only a single formula. Thus, we must distinguish the same nominal constant when it appears in two different formulas and we should treat judgments as being equal if they are identical up to permutations of these constants.

3.1.3 A Sequent Calculus Presentation of the Core Logic

The logic \mathcal{G} inherits from LG^ω the shift from a local to a global scope in the treatment of the ∇ -quantifier. In particular, we assume that the collection of constants is partitioned into the set \mathcal{C} of nominal constants and the set \mathcal{K} of usual, non-nominal constants. We assume the set \mathcal{C} contains an infinite number of nominal constants for each type at which ∇ quantification is permitted. We define the *support* of a term (or formula), written $\text{supp}(t)$, as the set of nominal constants appearing in it. A permutation of nominal constants is a type-preserving bijection π from \mathcal{C} to \mathcal{C} such that $\{x \mid \pi(x) \neq x\}$ is finite. We denote the application of such a permutation to a term or formula t by $\pi.t$ and define this as follows:

$$\begin{aligned} \pi.a &= \pi(a), \text{ if } a \in \mathcal{C} & \pi.c &= c, \text{ if } c \notin \mathcal{C} \text{ is atomic} \\ \pi.(\lambda x.M) &= \lambda x.(\pi.M) & \pi.(M N) &= (\pi.M) (\pi.N) \end{aligned}$$

We extend the notion of equality between terms to encompass also the application of permutations to nominal constants appearing in them. Specifically, we write $B \approx B'$ to denote the fact that there is a permutation π such that B λ -converts to $\pi.B'$. Using the observations that permutations are invertible and composable and that λ -convertibility is an equivalence relation, it is easy to see that \approx is also an equivalence relation.

The rules defining the core of \mathcal{G} are presented in Figure 3.1. Sequents in this logic have the form $\Sigma : \Gamma \longrightarrow C$ where Γ is a multiset and the signature Σ contains all the free variables of Γ and C . We use expressions of the form $B[t/x]$ in the quantifier rules

$$\begin{array}{c}
\frac{B \approx B'}{\Sigma : \Gamma, B \longrightarrow B'} \textit{id} \quad \frac{\Sigma : \Gamma \longrightarrow B \quad \Sigma : B, \Delta \longrightarrow C}{\Sigma : \Gamma, \Delta \longrightarrow C} \textit{cut} \quad \frac{\Sigma : \Gamma, B, B \longrightarrow C}{\Sigma : \Gamma, B \longrightarrow C} \textit{c}\mathcal{L} \\
\\
\overline{\Sigma : \Gamma, \perp \longrightarrow C} \perp\mathcal{L} \quad \overline{\Sigma : \Gamma \longrightarrow \top} \top\mathcal{R} \\
\\
\frac{\Sigma : \Gamma, B \longrightarrow C \quad \Sigma : \Gamma, D \longrightarrow C}{\Sigma : \Gamma, B \vee D \longrightarrow C} \vee\mathcal{L} \quad \frac{\Sigma : \Gamma \longrightarrow B_i}{\Sigma : \Gamma \longrightarrow B_1 \vee B_2} \vee\mathcal{R}, i \in \{1, 2\} \\
\\
\frac{\Sigma : \Gamma, B_i \longrightarrow C}{\Sigma : \Gamma, B_1 \wedge B_2 \longrightarrow C} \wedge\mathcal{L}, i \in \{1, 2\} \quad \frac{\Sigma : \Gamma \longrightarrow B \quad \Sigma : \Gamma \longrightarrow C}{\Sigma : \Gamma \longrightarrow B \wedge C} \wedge\mathcal{R} \\
\\
\frac{\Sigma : \Gamma \longrightarrow B \quad \Sigma : \Gamma, D \longrightarrow C}{\Sigma : \Gamma, B \supset D \longrightarrow C} \supset\mathcal{L} \quad \frac{\Sigma : \Gamma, B \longrightarrow C}{\Sigma : \Gamma \longrightarrow B \supset C} \supset\mathcal{R} \\
\\
\frac{\Sigma, \mathcal{K}, \mathcal{C} \vdash t : \tau \quad \Sigma : \Gamma, B[t/x] \longrightarrow C}{\Sigma : \Gamma, \forall_\tau x. B \longrightarrow C} \forall\mathcal{L} \quad \frac{\Sigma, h : \Gamma \longrightarrow B[h \vec{c}/x]}{\Sigma : \Gamma \longrightarrow \forall x. B} \forall\mathcal{R}, h \notin \Sigma, \text{supp}(B) = \{\vec{c}\} \\
\\
\frac{\Sigma, h : \Gamma, B[h \vec{c}/x] \longrightarrow C}{\Sigma : \Gamma, \exists x. B \longrightarrow C} \exists\mathcal{L}, h \notin \Sigma, \text{supp}(B) = \{\vec{c}\} \quad \frac{\Sigma, \mathcal{K}, \mathcal{C} \vdash t : \tau \quad \Sigma : \Gamma \longrightarrow B[t/x]}{\Sigma : \Gamma \longrightarrow \exists_\tau x. B} \exists\mathcal{R} \\
\\
\frac{\Sigma : \Gamma, B[a/x] \longrightarrow C}{\Sigma : \Gamma, \nabla x. B \longrightarrow C} \nabla\mathcal{L}, a \notin \text{supp}(B) \quad \frac{\Sigma : \Gamma \longrightarrow B[a/x]}{\Sigma : \Gamma \longrightarrow \nabla x. B} \nabla\mathcal{R}, a \notin \text{supp}(B)
\end{array}$$

Figure 3.1: The core rules of \mathcal{G}

to denote the result of substituting the term t for x in the formula B . Note that such a substitution must be done carefully, making sure to rename bound variables in B to avoid capture of variables appearing in t . In the $\nabla\mathcal{L}$ and $\nabla\mathcal{R}$ rules, a denotes a nominal constant of an appropriate type. In the $\exists\mathcal{L}$ and $\forall\mathcal{R}$ rule we use raising [Mil92] to encode the dependency of the quantified variable on the support of B ; the expression $(h \vec{c})$ in which h is a fresh eigenvariable is used in these two rules to denote the (curried) application of h to the constants appearing in the sequence \vec{c} . The $\forall\mathcal{L}$ and $\exists\mathcal{R}$ rules make use of judgments of the form $\Sigma, \mathcal{K}, \mathcal{C} \vdash t : \tau$. These judgments enforce the requirement that the expression

t instantiating the quantifier in the rule is a well-formed term of type τ constructed from the eigenvariables in Σ and the constants in $\mathcal{K} \cup \mathcal{C}$. Notice that in contrast the $\forall\mathcal{R}$ and $\exists\mathcal{L}$ rules seem to allow for a dependency on only a restricted set of nominal constants. However, this asymmetry is not significant: Corollary 4.1.5 in Section 4.1 will tell us that the dependency expressed through raising in the latter rules can be extended to any number of nominal constants that are not in the relevant support set without affecting the provability of sequents.

Equality modulo λ -conversion is built into the rules in Figure 3.1, and also into later extensions of this logic, in a fundamental way: in particular, proofs are preserved under the replacement of formulas in sequents by ones to which they λ -convert. A more involved observation is that we can replace a formula B in a sequent by another formula B' such that $B \approx B'$ without affecting the provability of the sequent or even the very structure of the proof. For the core logic, this observation follows from the form of the *id* rule and the fact that permutations distribute over logical structure. We shall prove this property explicitly for the full logic in Chapter 4.

3.2 Characterizing Occurrences of Nominal Constants

We are interested in adding to our logic the capability of characterizing occurrences of nominal constants within terms and also of analyzing the structure of terms with respect to such occurrences. For example, we may want to define a predicate called *name* that holds of a term exactly when that term is a nominal constant. Similarly, we might need to identify a binary relation called *fresh* that holds between two terms just in the case that the first term is a nominal constant that does not occur in the second term. Towards supporting such possibilities, we define in this section a special binary relation called *nominal abstraction* and then present proof rules that incorporate an understanding of this relation into the logic. A formalization of these ideas requires a careful treatment of substitution. In particular, this operation must be defined to respect the intended formula-level scope of nominal constants. We begin our discussion with an elaboration of this aspect.

3.2.1 Substitutions and their Interaction with Nominal Constants

The following definition reiterates a common view of substitutions in logical contexts.

Definition 3.2.1. *A substitution is a type preserving mapping from variables to terms that is the identity at all but a finite number of variables. The domain of a substitution is the set of variables that are not mapped to themselves and its range is the set of terms resulting from applying it to the variables in its domain. We write a substitution as $\{t_1/x_1, \dots, t_n/x_n\}$ where x_1, \dots, x_n is a list of variables that contains the domain of the substitution and t_1, \dots, t_n is the value of the map on these variables. The support of a substitution θ , written as $\text{supp}(\theta)$, is the set of nominal constants that appear in the range of θ . The restriction of a substitution θ to the set of variables Σ , written as $\theta \upharpoonright \Sigma$, is a mapping that is like θ on the variables in Σ and the identity everywhere else.*

A substitution essentially calls for the replacement of variables by their associated terms in any context to which it is applied. A complicating factor in our setting is that nominal constants can appear in the terms that are to replace particular variables. A substitution may be determined relative to one formula in a sequent but may then have to be applied to other formulas in the same sequent. In doing this, we have to take into account the fact that the scopes of the implicit quantifiers over nominal constants are restricted to individual formulas. Thus, the logically correct application of a substitution should be accompanied by a renaming of these constants in the term being substituted into so as to ensure that they are not confused with the ones appearing in the range of the substitution.

Definition 3.2.2. *The ordinary application of a substitution θ to a term B is denoted by $B[\theta]$ and corresponds to the replacement of the variables in B by the terms that θ maps them to, making sure, as usual, to avoid accidental binding of the variables appearing in the range of θ . More precisely, if $\theta = \{t_1/x_1, \dots, t_n/x_n\}$, then $B[\theta]$ is the term $(\lambda x_1 \dots \lambda x_n. B) t_1 \dots t_n$; this term is, of course, considered to be equal to any other term that it λ -converts to. By contrast, the nominal capture avoiding application of θ to B is written as $B[\![\theta]\!]$ and is defined as follows. Assuming that π is a permutation of nomi-*

nal constants that maps those appearing in $\text{supp}(B)$ to ones not appearing in $\text{supp}(\theta)$, let $B' = \pi.B$. Then $B[\theta] = B'[\theta]$.

The notation $B[\theta]$ generalizes the one used in the quantifier rules in Figure 3.1. The definition of the nominal capture avoiding application of a substitution is ambiguous in that we do not uniquely specify the permutation to be used. We resolve this ambiguity by deeming as acceptable *any* permutation that avoids conflicts. As a special instance of the lemma below, we see that for any given formula B and substitution θ , all the possible values for $B[\theta]$ are equivalent modulo the \approx relation. Moreover, as we show in Chapter 4, formulas that are equivalent under \approx are interchangeable in the contexts of proofs.

Lemma 3.2.3. *If $t \approx t'$ then $t[\theta] \approx t'[\theta]$.*

Proof. Let t be λ -convertible to $\pi_1.t'$, let $t[\theta] = (\pi_2.t)[\theta]$ where $\text{supp}(\pi_2.t) \cap \text{supp}(\theta) = \emptyset$, and let $t'[\theta]$ be λ -convertible to $(\pi_3.t')[\theta]$ where $\text{supp}(\pi_3.t') \cap \text{supp}(\theta) = \emptyset$. Then we define a function π partially by the following rules:

1. $\pi(c) = \pi_2.\pi_1.\pi_3^{-1}(c)$ if $c \in \text{supp}(\pi_3.t')$ and
2. $\pi(c) = c$ if $c \in \text{supp}(\theta)$.

Since $\text{supp}(\pi_3.t') \cap \text{supp}(\theta) = \emptyset$, these rules are not contradictory, *i.e.*, this (partial) function is well-defined. The range of the first rule is $\text{supp}(\pi_2.\pi_1.\pi_3^{-1}.\pi_3.t') = \text{supp}(\pi_2.\pi_1.t') = \text{supp}(\pi_2.t)$ which is disjoint from the range of the second rule, $\text{supp}(\theta)$. Since the mapping in each rule is determined by a permutation, these rules together define a one-to-one partial mapping that can be extended to a bijection on \mathcal{C} . We take any such extension to be the complete definition of π that must therefore be a permutation.

To prove that $t[\theta] \approx t'[\theta]$ it suffices to show that $(\pi_2.t)[\theta]$ is λ -convertible to $\pi.((\pi_3.t')[\theta])$. We do this by induction on the structure of t' under the further assumption that t λ -converts to $\pi_1.t'$. Suppose t' is an abstraction. Then, it is easy to see that $(\pi_2.t)[\theta]$ λ -converts to $\lambda x.((\pi_2.s)[\theta])$ and $\pi.((\pi_3.t')[\theta])$ λ -converts to $\lambda x.(\pi.((\pi_3.s')[\theta]))$ for some choice of variable x and terms s and s' such that s' is structurally less complex than t' and s λ -converts

to $\pi_1.s'$. But then, by the induction hypothesis, $(\pi_2.s)[\theta]$ λ -converts to $\pi.((\pi_3.s')[\theta])$ and hence $(\pi_2.t)[\theta]$ is λ -convertible to $\pi.((\pi_3.t')[\theta])$. A similar and, in fact, simpler argument can be provided in the case where t' is an application. If t' is a nominal constant c then $(\pi_2.t)[\theta]$ must be λ -convertible to $(\pi_2.\pi_1.c)[\theta] = \pi_2.\pi_1.c$. Also, $\pi.((\pi_3.t')[\theta])$ must be λ -convertible to $\pi.\pi_3.c$. Further, in this case the first rule for π applies which means $\pi.\pi_3.c = \pi_2.\pi_1.\pi_3^{-1}.\pi_3.c = \pi_2.\pi_1.c$. Thus $(\pi_2.t)[\theta]$ is again λ -convertible to $\pi.((\pi_3.t')[\theta])$. Finally, suppose t' is a variable x . In this case t must be λ -convertible to x so that we must show $x[\theta]$ λ -converts to $\pi.(x[\theta])$. If x does not have a binding in θ then both terms are equal. Alternatively, if $x[\theta] = s$ then $\pi.s = s$ by the second rule for π and so the two terms are again equal. Thus $(\pi_2.t)[\theta]$ λ -converts to $\pi.((\pi_3.t')[\theta])$, as is required. \square

The nominal capture avoiding application of substitutions turns out to be the dominant notion in the analysis of provability. For this reason, when we speak of the application of a substitution in an unqualified way, we shall mean the nominal capture avoiding form of this notion.

We shall need to consider the composition of substitutions later in this section. The definition of this notion must also pay attention to the presence of nominal constants.

Definition 3.2.4. *Given a substitution θ and a permutation π of nominal constants, let $\pi.\theta$ denote the substitution that is obtained by replacing each t/x in θ with $(\pi.t)/x$. Given any two substitutions θ and ρ , let $\theta \circ \rho$ denote the substitution that is such that $B[\theta \circ \rho] = B[\theta][\rho]$. In this context, the nominal capture avoiding composition of θ and ρ is written as $\theta \bullet \rho$ and defined as follows. Let π be a permutation of nominal constants such that $\text{supp}(\pi.\theta)$ is disjoint from $\text{supp}(\rho)$. Then $\theta \bullet \rho = (\pi.\theta) \circ \rho$.*

The notation $\theta \circ \rho$ in the above definition represents the usual composition of θ and ρ and can, in fact, be given in an explicit form based on these substitutions. Thus, $\theta \bullet \rho$ can also be presented in an explicit form. Notice that our definition of nominal capture avoiding composition is, once again, ambiguous because it does not fix the permutation to be used, accepting instead any one that satisfies the constraints. However, as before, this

ambiguity is harmless. To understand this, we first extend the notion of equivalence under permutations to substitutions.

Definition 3.2.5. *Two substitutions θ and ρ are considered to be permutation equivalent, written $\theta \approx \rho$, if and only if there is a permutation of nominal constants π such that $\theta = \pi.\rho$. This notion of equivalence may also be parameterized by a set of variables Σ as follows: $\theta \approx_{\Sigma} \rho$ just in the case that $\theta \uparrow \Sigma \approx \rho \uparrow \Sigma$.*

It is easy to see that all possible choices for $\theta \bullet \rho$ are permutation equivalent and that if $\varphi_1 \approx \varphi_2$ then $B[\varphi_1] \approx B[\varphi_2]$ for any term B . Thus, if our focus is on provability, the ambiguity in Definition 3.2.4 is inconsequential by a result to be established in Chapter 4. As a further observation, note that $B[\theta \bullet \rho] \approx B[\theta][\rho]$ for any B . Hence our notion of nominal capture avoiding composition of substitutions is sensible.

The composition operation can be used to define an ordering relation between substitutions:

Definition 3.2.6. *Given two substitutions ρ and θ , we say ρ is less general than θ , notated as $\rho \leq \theta$, if and only if there exists a σ such that $\rho \approx \theta \bullet \sigma$. This relation can also be parameterized by a set of variables: ρ is less general than θ relative to Σ , written as $\rho \leq_{\Sigma} \theta$, if and only if $\rho \uparrow \Sigma \leq \theta \uparrow \Sigma$.*

The notion of generality between substitutions that is based on nominal capture avoiding composition has a different flavor from that based on the traditional form of substitution composition. For example, if a is a nominal constant, the substitution $\{a/x\}$ is strictly less general than $\{a/x, y'a/y\}$ relative to Σ for any Σ which contains x and y . To see this, note that we can compose the latter substitution with $\{(\lambda z.y)/y'\}$ to obtain the former, but the naive attempt to compose the former with $\{y'a/y\}$ yields $\{b/x, y'a/y\}$ where b is a nominal constant distinct from a . In fact, the “most general” solution relative to Σ containing $\{a/x\}$ will be $\{a/x\} \cup \{z'a/z \mid z \in \Sigma \setminus \{x\}\}$.

3.2.2 Nominal Abstraction

The nominal abstraction relation allows implicit formula-level bindings represented by nominal constants to be moved into explicit abstractions over terms. The following notation is useful for defining this relationship.

Notation 3.2.7. *Let t be a term, let c_1, \dots, c_n be distinct nominal constants that possibly occur in t , and let y_1, \dots, y_n be distinct variables not occurring in t and such that, for $1 \leq i \leq n$, y_i and c_i have the same type. Then we write $\lambda c_1 \dots \lambda c_n.t$ to denote the term $\lambda y_1 \dots \lambda y_n.t'$ where t' is the term obtained from t by replacing c_i by y_i for $1 \leq i \leq n$.*

There is an ambiguity in the notation introduced above in that the choice of variables y_1, \dots, y_n is not fixed. However, this ambiguity is harmless: the terms that are produced by acceptable choices are all equivalent under a renaming of bound variables.

Definition 3.2.8. *Let $n \geq 0$ and let s and t be terms of type $\tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \tau$ and τ , respectively; notice, in particular, that s takes n arguments to yield a term of the same type as t . Then the expression $s \triangleright t$ is a formula that is referred to as a nominal abstraction of degree n or simply as a nominal abstraction. The symbol \triangleright is used here in an overloaded way in that the degree of the nominal abstraction it participates in can vary. The nominal abstraction $s \triangleright t$ of degree n is said to hold just in the case that s λ -converts to $\lambda c_1 \dots \lambda c_n.t$ for some nominal constants c_1, \dots, c_n .*

Clearly, nominal abstraction of degree 0 is the same as equality between terms based on λ -conversion, and we will therefore use $=$ to denote this relation in that situation. In the more general case, the term on the left of the operator serves as a pattern for isolating occurrences of nominal constants. For example, the relation $(\lambda x.x) \triangleright t$ holds exactly when t is a nominal constant.

The symbol \triangleright corresponds, at the moment, to a mathematical relation that holds between pairs of terms as explicated by Definition 3.2.8. We now overload this symbol by treating it also as a binary predicate symbol of \mathcal{G} . In the next subsection we shall add

inference rules to make the mathematical understanding of \triangleright coincide with its syntactic use as a predicate in sequents. It is, of course, necessary to be able to determine when we mean to use \triangleright in the mathematical sense and when as a logical symbol. When we write an expression such as $s \triangleright t$ without qualification, this should be read as a logical formula whereas if we say that “ $s \triangleright t$ holds” then we are referring to the abstract relation from Definition 3.2.8. We might also sometimes use an expression such as “ $(s \triangleright t)[\theta]$ holds.” In this case, we first treat $s \triangleright t$ as a formula to which we apply the substitution θ in a nominal capture avoiding way to get a (syntactic) expression of the form $s' \triangleright t'$. We then read \triangleright in the mathematical sense, interpreting the overall expression as the assertion that “ $s' \triangleright t'$ holds.” Note in this context that $s \triangleright t$ constitutes a single formula when read syntactically and hence the expression $(s \triangleright t)[\theta]$ is, in general, *not* equivalent to the expression $s[\theta] \triangleright t[\theta]$.

In the proof-theoretic setting, nominal abstraction will be used with terms that contain free occurrences of variables for which substitutions can be made. The following definition is relevant to this situation.

Definition 3.2.9. *A substitution θ is said to be a solution to the nominal abstraction $s \triangleright t$ just in the case that $(s \triangleright t)[\theta]$ holds.*

Solutions to a nominal abstraction can be used to provide rich characterizations of the structures of terms. For example, consider the nominal abstraction $(\lambda x. \text{fresh } x T) \triangleright S$ in which T and S are variables and *fresh* is a binary predicate symbol. Any solution to this problem requires that S be substituted for by a term of the form *fresh* $a R$ where a is a nominal constant and R is a term in which a does not appear, *i.e.*, a must be “fresh” to R .

An important property of solutions to a nominal abstraction is that these are preserved under permutations to nominal constants. We establish this fact in the lemma below; this lemma will be used later in showing the stability of the provability of sequents with respect to the replacement of formulas by ones they are equivalent to modulo the \approx relation.

Lemma 3.2.10. *Suppose $(s \triangleright t) \approx (s' \triangleright t')$. Then $s \triangleright t$ and $s' \triangleright t'$ have exactly the same solutions. In particular, $s \triangleright t$ holds if and only if $s' \triangleright t'$ holds.*

$$\frac{\{\Sigma\theta : \Gamma[\theta] \longrightarrow C[\theta] \mid \theta \text{ is a solution to } (s \succeq t)\}_\theta}{\Sigma : \Gamma, s \succeq t \longrightarrow C} \succeq \mathcal{L} \quad \frac{}{\Sigma : \Gamma \longrightarrow s \succeq t} \succeq \mathcal{R}, s \succeq t \text{ holds}$$

Figure 3.2: Nominal abstraction rules

$$\frac{\{\Sigma\theta : \Gamma[\theta] \longrightarrow C[\theta] \mid \theta \in \text{CSNAS}(\Sigma, s, t)\}_\theta}{\Sigma : \Gamma, s \succeq t \longrightarrow C} \succeq \mathcal{L}_{\text{CSNAS}}$$

Figure 3.3: A variant of $\succeq \mathcal{L}$ based on CSNAS

Proof. We prove the particular result first. It suffices to only show it in the forward direction since \approx is symmetric. Let π be the permutation such that the expression $s' \succeq t'$ λ -converts to $\pi.(s \succeq t)$. Now suppose $s \succeq t$ holds since s λ -converts to $\lambda\vec{c}.t$. Then s' will λ -convert to $\lambda(\pi.\vec{c}).t'$ where $\pi.\vec{c}$ is the result of applying π to each element in the sequence \vec{c} . Thus $s' \succeq t'$ holds.

For the general result it again suffices to show it in one direction, *i.e.*, that all the solutions of $s \succeq t$ are solutions to $s' \succeq t'$. Let θ be a substitution such that $(s \succeq t)[\theta]$ holds. By Lemma 3.2.3, $(s \succeq t)[\theta] \approx (s' \succeq t')[\theta]$. Thus by the particular result from the first half of this proof, $(s' \succeq t')[\theta]$ holds. \square

3.2.3 Proof Rules for Nominal Abstraction

We now add the left and right introduction rules for \succeq that are shown in Figure 3.2 to link its use as a predicate symbol to its mathematical interpretation. The expression $\Sigma\theta$ in the $\succeq \mathcal{L}$ rule denotes the application of a substitution $\theta = \{t_1/x_1, \dots, t_n/x_n\}$ to the signature Σ that is defined to be the signature that results from removing from Σ the variables $\{x_1, \dots, x_n\}$ and then adding every variable that is free in any term in $\{t_1, \dots, t_n\}$. Notice also that in the same inference rule the operator $[\theta]$ is applied to a multiset of formulas in the natural way: $\Gamma[\theta] = \{B[\theta] \mid B \in \Gamma\}$. Note that the $\succeq \mathcal{L}$ rule has an *a priori* unspecified number of premises that depends on the number of substitutions that are solutions to the

relevant nominal abstraction. If $s \supseteq t$ expresses an unsatisfiable constraint, meaning that it has no solutions, then the premise of $\supseteq\mathcal{L}$ is empty and the rule provides an immediate proof of its conclusion.

The $\supseteq\mathcal{L}$ and $\supseteq\mathcal{R}$ rules capture nicely the intended interpretation of nominal abstraction. However, there is an obstacle to using the former rule in derivations: this rule has an infinite number of premises any time the nominal abstraction $s \supseteq t$ has a solution. We can overcome this difficulty by describing a rule that includes only a few of these premises but in such way that their provability ensures the provability of all the other premises. Since the provability of $\Gamma \longrightarrow C$ implies the provability of $\Gamma[\theta] \longrightarrow C[\theta]$ for any θ (a property established formally in Chapter 4), if the first sequent is a premise of an occurrence of the $\supseteq\mathcal{L}$ rule, the second does not need to be used as a premise of that same rule occurrence. Thus, we can limit the set of premises to be considered if we can identify with any given nominal abstraction a (possibly finite) set of solutions from which any other solution can be obtained through composition with a suitable substitution. The following definition formalizes the idea of such a “covering set.”

Definition 3.2.11. *A complete set of nominal abstraction solutions (CSNAS) of s and t on Σ is a set S of substitutions such that*

1. *each $\theta \in S$ is a solution to $s \supseteq t$, and*
2. *for every solution ρ to $s \supseteq t$, there exists a $\theta \in S$ such that $\rho \leq_{\Sigma} \theta$.*

We denote any such set by $CSNAS(\Sigma, s, t)$.

Using this definition we present an alternative version of $\supseteq\mathcal{L}$ in Figure 3.3. Note that if we can find a finite complete set of nominal abstraction solutions then the number of premises to this rule will be finite.

Theorem 3.2.12. *The rules $\supseteq\mathcal{L}$ and $\supseteq\mathcal{L}_{CSNAS}$ are inter-admissible.*

Proof. Suppose we have the following arbitrary instance of $\supseteq\mathcal{L}$ in a derivation:

$$\frac{\{\Sigma\theta : \Gamma[\theta] \longrightarrow C[\theta] \mid \theta \text{ is a solution to } (s \supseteq t)\}_{\theta}}{\Sigma : \Gamma, s \supseteq t \longrightarrow C} \supseteq\mathcal{L}$$

This rule can be replaced with a use of $\triangleright\mathcal{L}_{CSNAS}$ instead if we could be certain that, for each $\rho \in CSNAS(\Sigma, s, t)$, it is the case that $\Sigma\rho : \Gamma[\rho] \longrightarrow C[\rho]$ is included in the set of premises of the shown rule instance. But this must be the case: by the definition of $CSNAS$, each such ρ is a solution to $s \triangleright t$.

In the other direction, suppose we have the following arbitrary instance of $\triangleright\mathcal{L}_{CSNAS}$.

$$\frac{\{\Sigma\theta : \Gamma[\theta] \longrightarrow C[\theta] \mid \theta \in CSNAS(\Sigma, s, t)\}_\theta}{\Sigma : \Gamma, s \triangleright t \longrightarrow C} \triangleright\mathcal{L}_{CSNAS}$$

To replace this rule with a use of the $\triangleright\mathcal{L}$ rule instead, we need to be able to construct a derivation of $\Sigma\rho : \Gamma[\rho] \longrightarrow C[\rho]$ for each ρ that is a solution to $s \triangleright t$. By the definition of $CSNAS$, we know that for any such ρ there exists a $\theta \in CSNAS(\Sigma, s, t)$ such that $\rho \leq_\Sigma \theta$, *i.e.*, such that there exists a σ for which $\rho \uparrow \Sigma \approx (\theta \uparrow \Sigma) \bullet \sigma$. Since we are considering the application of these substitutions to a sequent all of whose eigenvariables are contained in Σ , we can drop the restriction on the substitutions and suppose that $\rho \approx \theta \bullet \sigma$. Now, we shall show in Chapter 4 that if a sequent has a derivation then the result of applying a substitution to it in a nominal capture-avoiding way produces a sequent that also has a derivation. Using this observation, it follows that $\Sigma\theta\sigma : \Gamma[\theta][\sigma] \longrightarrow C[\theta][\sigma]$ has a proof. But this sequent is permutation equivalent to $\Sigma\rho : \Gamma[\rho] \longrightarrow C[\rho]$ which must, again by a result established explicitly in Chapter 4, also have a proof. \square

Theorem 3.2.12 allows us to choose which of the left rules we wish to consider in any given context. We shall assume the $\triangleright\mathcal{L}$ rule in the formal treatment in the rest of this thesis, leaving the use of the $\triangleright\mathcal{L}_{CSNAS}$ rule to practical applications of the logic.

3.2.4 Computing Complete Sets of Nominal Abstraction Solutions

For the $\triangleright\mathcal{L}_{CSNAS}$ rule to be useful, we need an effective way to compute restricted complete sets of nominal abstraction solutions. We show here that the task of finding such complete sets of solutions can be reduced to that of finding complete sets of unifiers (*CSU*) for higher-order unification problems [Hue75]. In the straightforward approach to finding a solution

to a nominal abstraction $s \triangleright t$, we would first identify a substitution θ that we apply to $s \triangleright t$ to get $s' \triangleright t'$ and we would subsequently look for nominal constants to abstract from t' to get s' . To relate this problem to the usual notion of unification, we would like to invert this order: in particular, we would like to consider all possible ways of abstracting over nominal constants first and only later think of applying substitutions to make the terms equal. The difficulty with this second approach is that we do not know which nominal constants might appear in t' until after the substitution is applied. However, there is a way around this problem. Given the nominal abstraction $s \triangleright t$ of degree n , we first consider substitutions for the variables occurring in it that introduce n new nominal constants in a completely general way. Then we consider all possible ways of abstracting over the nominal constants appearing in the altered form of t and, for each of these cases, we look for a complete set of unifiers.

The idea described above is formalized in the following definition and associated theorem. We use the notation $CSU(s, t)$ in them to denote an arbitrary but fixed selection of a complete set of unifiers for the terms s and t .

Definition 3.2.13. *Let s and t be terms of type $\tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \tau$ and τ , respectively. Let c_1, \dots, c_n be n distinct nominal constants disjoint from $\text{supp}(s \triangleright t)$ such that, for $1 \leq i \leq n$, c_i has the type τ_i . Let Σ be a set of variables and for each $h \in \Sigma$ of type τ' , let h' be a distinct variable not in Σ that has type $\tau_1 \rightarrow \dots \rightarrow \tau_n \rightarrow \tau'$. Let $\sigma = \{h' \ c_1 \ \dots \ c_n / h \mid h \in \Sigma\}$ and let $s' = s[\sigma]$ and $t' = t[\sigma]$. Let*

$$C = \bigcup_{\vec{a}} CSU(\lambda \vec{b}. s', \lambda \vec{b}. \lambda \vec{a}. t')$$

where $\vec{a} = a_1, \dots, a_n$ ranges over all selections of n distinct nominal constants from $\text{supp}(t) \cup \{\vec{c}\}$ such that, for $1 \leq i \leq n$, a_i has type τ_i and \vec{b} is some corresponding listing of all the nominal constants in s' and t' that are not included in \vec{a} . Then we define

$$S(\Sigma, s, t) = \{\sigma \bullet \rho \mid \rho \in C\}$$

The use of the substitution σ above represents another instance of the application of the general technique of raising that allows certain variables (the h variables in this definition)

whose substitution instances might depend on certain nominal constants (c_1, \dots, c_n here) to be replaced by new variables of higher type (the h' variables) whose substitution instances are not allowed to depend on those nominal constants. This technique was previously used in the $\exists\mathcal{L}$ and $\forall\mathcal{R}$ rules presented in Section 3.1.

Theorem 3.2.14. *$S(\Sigma, s, t)$ is a complete set of nominal abstraction solutions for $s \succeq t$ on Σ .*

Proof. First note that $\text{supp}(\sigma) \cap \text{supp}(s \succeq t) = \emptyset$ and thus $(s \succeq t)[\theta]$ is equal to $(s' \succeq t')$. Now we must show that every element of $S(\Sigma, s, t)$ is a solution to $s \succeq t$. Let $\sigma \bullet \rho \in S(\Sigma, s, t)$ be an arbitrary element where σ is as in Definition 3.2.13, ρ is from $CSU(\lambda\vec{b}.s', \lambda\vec{b}.\lambda\vec{a}.t')$, and $s' = s[\sigma]$ and $t' = t[\sigma]$. By the definition of CSU we know $(\lambda\vec{b}.s' = \lambda\vec{b}.\lambda\vec{a}.t')[\rho]$. This means $(s' = \lambda\vec{a}.t')[\rho]$ holds and thus $(s' \succeq t')[\rho]$ holds. Rewriting s' and t' in terms of s and t this means $(s \succeq t)[\sigma][\rho]$. Thus $\sigma \bullet \rho$ is a solution to $s \succeq t$.

In the other direction, we must show that if θ is a solution to $s \succeq t$ then there exists $\sigma \bullet \rho \in S(\Sigma, s, t)$ such that $\theta \leq_{\Sigma} \sigma \bullet \rho$. Let θ be a solution to $s \succeq t$. Then we know $(s \succeq t)[\theta]$ holds. The substitution θ may introduce some nominal constants which are abstracted out of the right-hand side when determining equality, so let us call these the *important* nominal constants. Let $\sigma = \{h' \ c_1 \ \dots \ c_n/h \mid h \in \Sigma\}$ be as in Definition 3.2.13 and let π' be a permutation which maps the important nominal constants of θ to nominal constants from c_1, \dots, c_n . This is possible since n nominal constants are abstract from the right-hand side and thus there are at most n important nominal constants. Then let $\theta' = \pi'.\theta$, so that $(s \succeq t)[\theta']$ holds and it suffices to show that $\theta' \leq_{\Sigma} \sigma \bullet \rho$. Note that all we have done at this point is to rename the important nominal constants of θ so that they match those introduced by σ . Now we define $\rho' = \{\lambda c_1 \dots \lambda c_n.r/h' \mid r/h \in \theta'\}$ so that $\theta' = \sigma \bullet \rho'$. Thus $(s \succeq t)[\sigma][\rho']$ holds. By construction, σ shares no nominal constants with s and t , thus we know $(s' \succeq t')[\rho']$ where $s' = s[\sigma]$ and $t' = t[\sigma]$. Also by construction, ρ' contains no interesting nominal constants and thus $(s' = \lambda\vec{a}.t')[\rho]$ holds for some nominal constants \vec{a} taken from $\text{supp}(t) \cup \{\vec{c}\}$. If we let \vec{b} be a listing of all nominal constants in s' and t' but not in \vec{a} , then $(\lambda\vec{b}.s' = \lambda\vec{b}.\lambda\vec{a}.t')[\rho]$ holds. At this point the inner equality has no nominal

$$\frac{\Sigma : \Gamma, B \ p \ \vec{t} \longrightarrow C}{\Sigma : \Gamma, p \ \vec{t} \longrightarrow C} \text{ def}\mathcal{L} \qquad \frac{\Sigma : \Gamma \longrightarrow B \ p \ \vec{t}}{\Sigma : \Gamma \longrightarrow p \ \vec{t}} \text{ def}\mathcal{R}$$

Figure 3.4: Introduction rules for atoms whose predicate is defined as $\forall \vec{x}. p \ \vec{x} \triangleq B \ p \ \vec{x}$

constants and thus the substitution ρ can be applied without renaming: $(\lambda \vec{b}.s' = \lambda \vec{b}.\lambda \vec{a}.t')[\rho']$ holds. By the definition of CSU , there must be a $\rho \in CSU(\lambda \vec{b}.s', \lambda \vec{b}.\lambda \vec{a}.t')$ such that $\rho' \leq \rho$. Thus $\sigma \bullet \rho' \leq_{\Sigma} \sigma \bullet \rho$ as desired. \square

3.3 Definitions, Induction, and Co-induction

The sequent calculus rules presented in Figure 3.1 treat atomic judgments as fixed, unanalyzed objects. We now add the capability of defining such judgments by means of formulas, possibly involving other predicates. In particular, we shall assume that we are given a fixed, finite set of *clauses* of the form $\forall \vec{x}. p \ \vec{x} \triangleq B \ p \ \vec{x}$ where p is a predicate constant that takes a number of arguments equal to the length of \vec{x} . Such a clause is said to define p and the entire collection of clauses is called a *definition*. The expression B , called the *body* of the clause, must be a term that does not contain p or any of the variables in \vec{x} and must have a type such that $B \ p \ \vec{x}$ has type o . Definitions are also restricted so that a predicate is defined by at most one clause. The intended interpretation of a clause $\forall \vec{x}. p \ \vec{x} \triangleq B \ p \ \vec{x}$ is that the atomic formula $p \ \vec{t}$, where \vec{t} is a list of terms of the same length and type as the variables in \vec{x} , is true if and only if $B \ p \ \vec{t}$ is true. This interpretation is realized by adding to the calculus the rules $\text{def}\mathcal{L}$ and $\text{def}\mathcal{R}$ shown in Figure 3.4 for unfolding predicates on the left and the right of sequents using their defining clauses.

A definition can have a recursive structure. For example, in the clause $\forall \vec{x}. p \ \vec{x} \triangleq B \ p \ \vec{x}$, the predicate p can appear free in $B \ p \ \vec{x}$. In this setting, the meanings of predicates are intended to be given by any one of the fixed points that can be associated with the definition. Such an interpretation may not always be sensible. In particular, without further restrictions, the resulting proof system may not be consistent. There are two constraints

that suffice to ensure consistency. First, the body of a clause must not contain any nominal constants. This restriction can be justified from another perspective as well: as we see in Chapter 4, it helps in establishing that \approx is a provability preserving equivalence between formulas. Second, definitions should be *stratified* so that clauses, such as $a \triangleq (a \supset \perp)$, in which a predicate has a negative dependency on itself, are forbidden. While such stratification can be enforced in different ways, we use a simple approach to doing this in this thesis. This approach is based on associating with each predicate p a natural number that is called its *level* and that is denoted by $\text{lvl}(p)$. This measure is then extended to arbitrary formulas by the following definition.

Definition 3.3.1. *Given an assignment of levels to predicates, the function lvl is extended to all formulas in λ -normal form as follows:*

1. $\text{lvl}(p \bar{t}) = \text{lvl}(p)$
2. $\text{lvl}(\perp) = \text{lvl}(\top) = \text{lvl}(s \supseteq t) = 0$
3. $\text{lvl}(B \wedge C) = \text{lvl}(B \vee C) = \max(\text{lvl}(B), \text{lvl}(C))$
4. $\text{lvl}(B \supset C) = \max(\text{lvl}(B) + 1, \text{lvl}(C))$
5. $\text{lvl}(\forall x.B) = \text{lvl}(\nabla x.B) = \text{lvl}(\exists x.B) = \text{lvl}(B)$

In general, the level of a formula B , written as $\text{lvl}(B)$, is the level of its λ -normal form.

A definition is *stratified* if we can assign levels to predicates in such a way that $\text{lvl}(B p \vec{x}) \leq \text{lvl}(p)$ for each clause $\forall \vec{x}. p \vec{x} \triangleq B p \vec{x}$ in that definition.

The $\text{def}\mathcal{L}$ and $\text{def}\mathcal{R}$ rules do not discriminate between any of the fixed points of a definition. We now allow the selection of least and greatest fixed points so as to support inductive and co-inductive definitions of predicates. Specifically, we denote an inductive clause by $\forall \vec{x}. p \vec{x} \stackrel{\mu}{\triangleq} B p \vec{x}$ and a co-inductive one by $\forall \vec{x}. p \vec{x} \stackrel{\nu}{\triangleq} B p \vec{x}$. As a refinement of the earlier restriction on definitions, a predicate may have at most one defining clause that is designated to be inductive, co-inductive or neither. The $\text{def}\mathcal{L}$ and $\text{def}\mathcal{R}$ rules may be

$$\frac{\vec{x} : B \ S \ \vec{x} \longrightarrow S \ \vec{x} \quad \Sigma : \Gamma, S \ \vec{t} \longrightarrow C}{\Sigma : \Gamma, p \ \vec{t} \longrightarrow C} \mathcal{IL}$$

provided p is defined as $\forall \vec{x}. p \ \vec{x} \stackrel{\mu}{=} B \ p \ \vec{x}$ and S is a term that has the same type as p

$$\frac{\Sigma : \Gamma \longrightarrow S \ \vec{t} \quad \vec{x} : S \ \vec{x} \longrightarrow B \ S \ \vec{x}}{\Sigma : \Gamma \longrightarrow p \ \vec{t}} \mathcal{CIR}$$

provided p is defined as $\forall \vec{x}. p \ \vec{x} \stackrel{\nu}{=} B \ p \ \vec{x}$ and S is a term that has the same type as p

Figure 3.5: The induction left and co-induction right rules

used with clauses in any one of these forms. Clauses that are inductive admit additionally the left rule \mathcal{IL} shown in Figure 3.5. This rule is based on the observation that the least fixed point of a monotone operator is the intersection of all its pre-fixed points; intuitively, anything that follows from any pre-fixed point should then also follow from the least fixed point. In a proof search setting, the term corresponding to the schema variable S in this rule functions like the induction hypothesis and is accordingly called the invariant of the induction. Clauses that are co-inductive, on the other hand, admit the right rule \mathcal{CIR} also presented in Figure 3.5. This rule reflects the fact that the greatest fixed point of a monotone operator is the union of all the post-fixed points; any member of such a post-fixed point must therefore also be a member of the greatest fixed point. The substitution that is used for S in this rule is called the co-invariant or the simulation of the co-induction. Just like the restriction on the body of clauses, in both \mathcal{IL} and \mathcal{CIR} , the (co-)invariant S must not contain any nominal constants.

As a simple illustration of the use of these rules, consider the clause $p \stackrel{\mu}{=} p$. The desired inductive reading of this clause implies that p must be false. In a proof-theoretic setting, we would therefore expect that the sequent $\cdot : p \longrightarrow \perp$ can be proved. This can, in fact, be done by using \mathcal{IL} with the invariant $S = \perp$. On the other hand, consider the clause $q \stackrel{\nu}{=} q$. The co-inductive reading intended here implies that q must be true. The logic \mathcal{G} satisfies this expectation: the sequent $\cdot : \cdot \longrightarrow q$ can be proved using \mathcal{CIR} with the co-invariant

$S = \top$.

The addition of inductive and co-inductive forms of clauses and the mixing of these forms in one setting might be expected to require stronger conditions than those described earlier in this section to guarantee consistency. One condition, in addition to the absence of nominal constants in the bodies of clauses and stratification based on levels, that suffices and that is also practically acceptable is the following that is taken from [TM09]: in a clause of any of the forms $\forall \vec{x}. p \vec{x} \triangleq B p \vec{x}$, $\forall \vec{x}. p \vec{x} \stackrel{\mu}{=} B p \vec{x}$ or $\forall \vec{x}. p \vec{x} \stackrel{\nu}{=} B p \vec{x}$, it must be that $\text{lvl}(B (\lambda \vec{x}. \top) \vec{x}) < \text{lvl}(p)$. This disallows any mutual recursion between clauses, a restriction which can easily be overcome by merging mutually recursive clauses into a single clause. We henceforth assume that all definitions satisfy all three conditions described for them in this section. Corollary 4.1.7 in Chapter 4 establishes the consistency of the logic under these restrictions.

3.4 A Pattern-Based Form for Definitions

When presenting a definition for a predicate, it is often convenient to write this as a collection of clauses whose applicability is also constrained by patterns appearing in the head. For example, in logics that support equality but not nominal abstraction, list membership may be defined by the two pattern based clauses shown below.

$$\text{member } X (X :: L) \triangleq \top \qquad \text{member } X (Y :: L) \triangleq \text{member } X L$$

These logics also include rules for directly treating definitions presented in this way. In understanding these rules, use may be made of the translation of the extended form of definitions to a version that does not use patterns in the head and in which there is at most one clause for each predicate. For example, the definition of the list membership predicate would be translated to the following form:

$$\text{member } X K \triangleq (\exists L. K = (X :: L)) \vee (\exists Y \exists L. K = (Y :: L) \wedge \text{member } X L)$$

The treatment of patterns and multiple clauses can now be understood in terms of the rules for definitions using a single clause and the rules for equality, disjunction, and existential

quantification.

In the logic \mathcal{G} , the notion of equality has been generalized to that of nominal abstraction. This allows us also to expand the pattern-based form of definitions to use nominal abstraction in determining the selection of clauses. By doing this, we would allow the head of a clausal definition to describe not only the term structure of the arguments, but also to place restrictions on the occurrences of nominal constants in these arguments. For example, suppose we want to describe the contexts in typing judgments by lists of the form $of\ c_1\ T_1 :: of\ c_2\ T_2 :: \dots :: nil$ with the further proviso that each c_i is a distinct nominal constant. We will allow this to be done by using the following pattern-based form of definition for the predicate ctx :

$$ctx\ nil \triangleq \top \qquad (\nabla x.ctx\ (of\ x\ T :: L)) \triangleq ctx\ L$$

Intuitively, the ∇ quantifier in the head of the second clause imposes the requirement that, to match it, the argument of ctx should have the form $of\ x\ T :: L$ where x is a nominal constant that does not occur in either T or L . To understand this interpretation, we could think of the earlier definition of ctx as corresponding to the following one that does not use patterns or multiple clauses:

$$ctx\ K \triangleq (K = nil) \vee (\exists T \exists L. (\lambda x.of\ x\ T :: L) \supseteq K \wedge ctx\ L)$$

Our objective in the rest of this section is to develop machinery for allowing the extended form of definitions to be used directly. We do this by presenting its syntax formally, by describing rules that allow us to work off of such definitions and, finally, by justifying the new rules by means of a translation of the kind indicated above.

Definition 3.4.1. *A pattern-based definition is a finite collection of clauses of the form*

$$\forall \vec{x}. (\nabla \vec{z}. p\ \vec{t}) \triangleq B\ p\ \vec{x}$$

where \vec{t} is a sequence of terms that do not have occurrences of nominal constants in them, p is a constant such that $p\ \vec{t}$ is of type o and B is a term devoid of occurrences of p , \vec{x} and

$$\frac{\Sigma : \Gamma \longrightarrow (B \ p \ \vec{x})[\theta]}{\Sigma : \Gamma \longrightarrow p \ \vec{s}} \text{ def}\mathcal{R}^p$$

for any clause $\forall \vec{x}.(\nabla \vec{z}.p \ \vec{t}) \triangleq B \ p \ \vec{x}$ in \mathcal{D} and any θ such that $\text{range}(\theta) \cap \Sigma = \emptyset$ and

$$(\lambda \vec{z}.p \ \vec{t})[\theta] \triangleright p \ \vec{s} \text{ holds}$$

$$\frac{\left\{ \begin{array}{l} \Sigma\theta : \Gamma[\theta], (B \ p \ \vec{x})[\theta] \longrightarrow C[\theta] \\ \forall \vec{x}.(\nabla \vec{z}.p \ \vec{t}) \triangleq B \ p \ \vec{x} \in \mathcal{D} \text{ and} \\ \theta \text{ is a solution to } ((\lambda \vec{z}.p \ \vec{t}) \triangleright p \ \vec{s}) \end{array} \right\}}{\Sigma : \Gamma, p \ \vec{s} \longrightarrow C} \text{ def}\mathcal{L}^p$$

Figure 3.6: Introduction rules for a pattern-based definition \mathcal{D}

nominal constants and such that $B \ p \ \vec{t}$ is of type o . Further, we expect such a collection of clauses to satisfy a stratification condition: there must exist an assignment of levels to predicate symbols such that for any clause $\forall \vec{x}.(\nabla \vec{z}.p \ \vec{t}) \triangleq B \ p \ \vec{x}$ occurring in the set, assuming p has arity n , it is the case that $\text{lvl}(B \ (\lambda \vec{x}.\top) \ \vec{x}) < \text{lvl}(p)$. Notice that we allow the collection to contain more than one clause for any given predicate symbol.

The logical rules for treating pattern-based definitions are presented in Figure 3.6. These rules encode the idea of matching an instance of a predicate with the head of a particular clause and then replacing the predicate with the corresponding clause body. The kind of matching involved is made precise through the construction of a nominal abstraction after replacing the ∇ quantifiers in the head of the clause by abstractions. The right rule embodies the fact that it is enough if an instance of any one clause can be used in this way to yield a successful proof. In this rule, the substitution θ that results from the matching must be applied in a nominal capture avoiding way to the body. However, since B does not contain nominal constants, the ordinary application of the substitution also suffices. To accord with the treatment in the right rule, the left rule must consider all possible ways in which an instance of an atomic assumption $p \ \vec{s}$ can be matched by a clause and must show that a proof can be constructed in each such case.

The soundness of these rules is the content of the following theorem whose proof also makes explicit the intended interpretation of the pattern-based form of definitions.

Theorem 3.4.2. *The pattern-based form of definitions and the associated proof rules do not add any new power to the logic. In particular, the $\text{def}\mathcal{L}^p$ and $\text{def}\mathcal{R}^p$ rules are admissible under the intended interpretation via translation of the pattern-based form of definitions.*

Proof. Let p be a predicate whose clauses in the definition being considered are given by the following set of clauses.

$$\{\forall \vec{x}_i. (\nabla \vec{z}_i. p \vec{t}_i) \triangleq B_i p \vec{x}_i\}_{i \in 1..n}$$

Let p' be a new constant symbol with the same argument types as p . Then the intended interpretation of the definition of p in a setting that does not allow the use of patterns in the head and that limits the number of clauses defining a predicate to one is given by the clause

$$\forall \vec{y}. p \vec{y} \triangleq \bigvee_{i \in 1..n} \exists \vec{x}_i. ((\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{y}) \wedge B_i p \vec{x}_i$$

in which the variables \vec{y} are chosen such that they do not appear in the terms \vec{t}_i for $1 \leq i \leq n$. Note also that we are using the term constructor p' here so as to be able to match the entire head of a clause at once, thus ensuring that the ∇ -bound variables in the head are assigned a consistent value for all arguments of the predicate.

Based on this translation, we can replace an instance of $\text{def}\mathcal{R}^p$,

$$\frac{\Gamma \longrightarrow (B_i p \vec{x}_i)[\theta]}{\Gamma \longrightarrow p \vec{s}} \text{def}\mathcal{R}^p$$

with the following sequence of rules, where a double inference line indicates that a rule is used multiple times.

$$\frac{\frac{\frac{\frac{\Gamma \longrightarrow (\lambda \vec{z}_i. p' \vec{t}_i)[\theta] \triangleright p' \vec{s}}{\Gamma \longrightarrow ((\lambda \vec{z}_i. p' \vec{t}_i)[\theta] \triangleright p' \vec{s}) \wedge (B_i p \vec{x}_i)[\theta]} \wedge \mathcal{R}}{\Gamma \longrightarrow \exists \vec{x}_i. ((\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{s}) \wedge B_i p \vec{x}_i} \exists \mathcal{R}}{\Gamma \longrightarrow \bigvee_{i \in 1..n} \exists \vec{x}_i. ((\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{s}) \wedge B_i p \vec{x}_i} \vee \mathcal{R}}{\Gamma \longrightarrow p' \vec{t}} \text{def}\mathcal{R}$$

Note that we have made use of the fact that θ instantiates only the variables x_i and thus has no effect on \vec{s} . Further, the side condition associated with the $def\mathcal{R}^p$ rule ensures that the $\supseteq\mathcal{R}$ rule that appears as a left leaf in this derivation is well applied.

Similarly, we can replace an instance of $def\mathcal{L}^p$,

$$\frac{\{\Sigma\theta : \Gamma[\theta], (B_i p \vec{x}_i)[\theta] \longrightarrow C[\theta] \mid \theta \text{ is a solution to } ((\lambda\vec{z}.p \vec{t}_i) \supseteq p \vec{s})\}_{i \in 1..n}}{\Sigma : \Gamma, p \vec{s} \longrightarrow C} def\mathcal{L}^p$$

with the following sequence of rules

$$\frac{\left\{ \frac{\left\{ \frac{\Gamma[\theta], (B_i p \vec{x}_i)[\theta] \longrightarrow C[\theta] \mid \theta \text{ is a solution to } ((\lambda\vec{z}.p' \vec{t}_i) \supseteq p' \vec{s})}{\Gamma, (\lambda\vec{z}_i.p' \vec{t}_i) \supseteq p' \vec{s}, B_i p \vec{x}_i \longrightarrow C} \supseteq\mathcal{L}}{\Gamma, ((\lambda\vec{z}_i.p' \vec{t}_i) \supseteq p' \vec{s}) \wedge B_i p \vec{x}_i \longrightarrow C} \wedge\mathcal{L}^*}{\Gamma, \exists\vec{x}_i.((\lambda\vec{z}_i.p' \vec{t}_i) \supseteq p' \vec{s}) \wedge B_i p \vec{x}_i \longrightarrow C} \exists\mathcal{L}} \right\}_{i \in 1..n}}{\Gamma, \bigvee_{i \in 1..n} \exists\vec{x}_i.((\lambda\vec{z}_i.p' \vec{t}_i) \supseteq p' \vec{s}) \wedge B_i p \vec{x}_i \longrightarrow C} \vee\mathcal{L}} \frac{\Gamma, \bigvee_{i \in 1..n} \exists\vec{x}_i.((\lambda\vec{z}_i.p' \vec{t}_i) \supseteq p' \vec{s}) \wedge B_i p \vec{x}_i \longrightarrow C}{\Gamma, p \vec{s} \longrightarrow C} def\mathcal{L}$$

Here $\wedge\mathcal{L}^*$ is an application of $c\mathcal{L}$ followed by $\wedge\mathcal{L}_1$ and $\wedge\mathcal{L}_2$ on the contracted formula. It is easy to see that the solutions to $(\lambda\vec{z}.p \vec{t}_i) \supseteq p \vec{s}$ and $(\lambda\vec{z}.p' \vec{t}_i) \supseteq p' \vec{s}$ are identical and hence the leaf sequents in this partial derivation are exactly the same as the upper sequents of the instance of the $def\mathcal{L}^p$ rule being considered. \square

A weak form of a converse to the above theorem also holds. Suppose that the predicate p is given by the following clauses

$$\{\forall\vec{x}_i. (\nabla\vec{z}_i.p \vec{t}_i) \triangleq B_i p \vec{x}_i\}_{i \in 1..n}$$

in a setting that uses pattern-based definitions and that has the $def\mathcal{L}^p$ and $def\mathcal{R}^p$ but not the $def\mathcal{L}$ and $def\mathcal{R}$ rules. In such a logic, it is easy to see that the following is provable:

$$\forall\vec{y}. \left[p \vec{y} \equiv \bigvee_{i \in 1..n} \exists\vec{x}_i.((\lambda\vec{z}_i.p' \vec{t}_i) \supseteq p' \vec{y}) \wedge B_i p \vec{x}_i \right]$$

Where $B \equiv C$ denotes $(B \supset C) \wedge (C \supset B)$. Thus, in the presence of cut , the $def\mathcal{L}$ and $def\mathcal{R}$ rules can be treated as derived ones relative to the translation interpretation of pattern-based definitions.

$$\frac{\{\vec{x}_i : B_i \ S \ \vec{x}_i \longrightarrow \nabla \vec{z}_i . S \ \vec{t}_i\}_{i \in 1..n} \quad \Sigma : \Gamma, S \ \vec{s} \longrightarrow C}{\Sigma : \Gamma, p \ \vec{s} \longrightarrow C} \mathcal{IL}^p$$

assuming p is defined by the set of clauses $\{\forall \vec{x}_i . (\nabla \vec{z}_i . p \ \vec{t}_i) \stackrel{\mu}{=} B_i \ p \ \vec{x}_i\}_{i \in 1..n}$

Figure 3.7: Induction rule for pattern-based definitions

We would like also to allow patterns to be used in the heads of clauses when writing definitions that are intended to pick out the least and greatest fixed points, respectively. Towards this end we admit in a definition also clauses of the form $\forall \vec{x} . (\nabla \vec{z} . p \ \vec{t}) \stackrel{\mu}{=} B \ p \ \vec{x}$ and $\forall \vec{x} . (\nabla \vec{z} . p \ \vec{t}) \stackrel{\nu}{=} B \ p \ \vec{x}$ with the earlier provisos on the form of B and \vec{t} and the types of B and p and with the additional requirement that all the clauses for any given predicate are un-annotated or annotated uniformly with either μ or ν . Further, a definition must satisfy stratification conditions as before. In reasoning about the least or greatest fixed point forms of definitions, we may use the translation into the earlier, non-pattern form together with the rules \mathcal{IL} and \mathcal{CTR} . It is possible to formulate an induction rule that works directly from pattern-based definitions using the idea that to show S to be an induction invariant for the predicate p , one must show that every clause of p preserves S . A rule that is based on this intuition is presented in Figure 3.7. The soundness of this rule is shown in the following theorem.

Theorem 3.4.3. *The \mathcal{IL}^p rule is admissible under the intended translation of pattern-based definitions.*

Proof. Let the clauses for p in the pattern-based definition be given by the set

$$\{\forall \vec{x}_i . (\nabla \vec{z}_i . p \ \vec{t}_i) \stackrel{\mu}{=} B_i \ p \ \vec{x}_i\}_{i \in 1..n}$$

in which case the translated form of the definition for p would be

$$\forall \vec{y} . p \ \vec{y} \stackrel{\mu}{=} \bigvee_{i \in 1..n} \exists \vec{x}_i . ((\lambda \vec{z}_i . p' \ \vec{t}_i) \triangleright p' \ \vec{y}) \wedge B_i \ p \ \vec{x}_i.$$

In this context, the rightmost upper sequents of the \mathcal{IL}^p and the \mathcal{IL} rules that are needed to derive a sequent of the form $\Sigma : \Gamma, p \vec{s} \longrightarrow C$ are identical. Thus, to show that \mathcal{IL}^p rule is admissible, it suffices to show that the left upper sequent in the \mathcal{IL} rule can be derived in the original calculus from all but the rightmost upper sequent in an \mathcal{IL}^p rule. Towards this end, we observe that we can construct the following derivation:

$$\frac{\left\{ \frac{\left\{ \frac{\{ (\vec{y}, \vec{x}_i)\theta : (B_i p \vec{x}_i)[\theta] \longrightarrow (S \vec{y})[\theta] \mid \theta \text{ is a solution to } ((\lambda \vec{z}. p' \vec{t}_i) \triangleright p' \vec{y}) \}}{\vec{y}, \vec{x}_i : (\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{y}, B_i S \vec{x}_i \longrightarrow S \vec{y}} \triangleright \mathcal{L} \right.}{\vec{y}, \vec{x}_i : ((\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{y}) \wedge B_i p \vec{x}_i \longrightarrow S \vec{y}} \wedge \mathcal{L}^*}{\vec{y} : \exists \vec{x}_i. ((\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{y}) \wedge B_i S \vec{x}_i \longrightarrow S \vec{y}} \exists \mathcal{L} \right.}{\vec{y} : \bigvee_{i \in 1..n} \exists \vec{x}_i. ((\lambda \vec{z}_i. p' \vec{t}_i) \triangleright p' \vec{y}) \wedge B_i S \vec{x}_i \longrightarrow S \vec{y}} \bigvee \mathcal{L} \right.}$$

Since the variables \vec{y} are distinct and do not occur in \vec{t}_i , the solutions to $(\lambda \vec{z}. p' \vec{t}_i) \triangleright p' \vec{y}$ have a simple form. In particular, let \vec{t}'_i be the result of replacing in \vec{t}_i the variables \vec{z} with distinct nominal constants. Then $\vec{y} = \vec{t}'_i$ will be a most general solution to the nominal abstraction. Thus the upper sequents of the invariant derivation above will be

$$\vec{x}_i : B_i p \vec{x}_i \longrightarrow S \vec{t}'_i$$

which are derivable if and only if the sequents

$$\vec{x}_i : B_i p \vec{x}_i \longrightarrow \nabla \vec{z}_i. S \vec{t}_i$$

are derivable. □

We do not introduce a co-induction rule for pattern-based definitions largely because it seems that there are few interesting co-inductive definitions that require patterns and multiple clauses.

3.5 Examples

We now provide some examples to illuminate the properties of nominal abstraction and its usefulness in both specification and reasoning tasks; while \mathcal{G} has many more features, their characteristics and applications have been exposed in other work (*e.g.*, see [MM02, MT03b,

Tiu04, TM08]). In the examples that are shown, use will be made of the pattern-based form of definitions described in Section 3.4. We will also use the convention that tokens given by capital letters denote variables that are implicitly universally quantified over the entire clause.

3.5.1 Properties of ∇ and Freshness

We can use nominal abstraction to gain a better insight into the behavior of the ∇ quantifier. Towards this end, let the *fresh* predicate be defined by the following clause.

$$(\nabla x. \text{fresh } x \ E) \triangleq \top$$

We have elided the type of *fresh* here; it will have to be defined at each type that it is needed in the examples we consider below. Alternatively, we can “inline” the definition by using nominal abstraction directly, *i.e.*, by replacing occurrences of *fresh* t_1 t_2 with $\exists E.(\lambda x. \langle x, E \rangle \triangleright \langle t_1, t_2 \rangle)$ for a suitably typed pairing construct $\langle \cdot, \cdot \rangle$.

Now let B be a formula whose free variables are among z, x_1, \dots, x_n , and let $\vec{x} = x_1 :: \dots :: x_n :: \text{nil}$ where $::$ and *nil* are constructors in the logic.² Then the following formulas logically imply one another in \mathcal{G} .

$$\nabla z. B \quad \exists z. (\text{fresh } z \ \vec{x} \wedge B) \quad \forall z. (\text{fresh } z \ \vec{x} \supset B)$$

Note that the type of z allows it to be an arbitrary term in the last two formulas, but its occurrence as the first argument of *fresh* will restrict it to being a nominal constant (even when $\vec{x} = \text{nil}$).

In the original presentation of the ∇ quantifier [MT03a], it was shown that one can move a ∇ quantifier inwards over universal and existential quantifiers by using raising to encode an explicit dependency. To illustrate this, let B be a formula with two variables

²We are, once again, finessing typing issues here in that the x_i variables may not all be of the same type. However, this problem can be solved by surrounding each of them with a constructor that yields a term with a uniform type.

abstracted out, and let $C \equiv D$ be shorthand for $(C \supset D) \wedge (D \supset C)$. The the following formulas are provable in the logic.

$$\nabla z.\forall x.(B z x) \equiv \forall h.\nabla z.(B z (h z)) \quad \nabla z.\exists x.(B z x) \equiv \exists h.\nabla z.(B z (h z))$$

In order to move a ∇ quantifier outwards over universal and existential quantifiers, one would need a way to make non-dependency (*i.e.*, freshness) explicit. This is now possible using nominal abstraction as shown by the following equivalences.

$$\forall x.\nabla z.(B z x) \equiv \nabla z.\forall x.(fresh z x \supset B z x) \quad \exists x.\nabla z.(B z x) \equiv \nabla z.\exists x.(fresh z x \wedge B z x)$$

Finally, we note that the two sets of equivalences for moving the ∇ quantifier interact nicely. Specifically, starting with a formula like $\nabla z.\forall x.(B z x)$ we can push the ∇ quantifier inwards and then outwards to obtain $\nabla z.\forall h.(fresh z (h z) \supset B z (h z))$. Here $fresh z (h z)$ will only be satisfied if h projects away its first argument, as expected.

3.5.2 Polymorphic Type Generalization

In addition to reasoning, nominal abstraction can also be useful in providing declarative specifications of computations. We consider the context of a type inference algorithm that is also discussed in [CU08] to illustrate such an application. In this setting, we might need a predicate *spec* that relates a polymorphic type σ , a list of distinct variables list of distinct variables $\vec{\alpha}$ (represented by nominal constants) and a monomorphic type τ just in the case that $\sigma = \forall \vec{\alpha}.\tau$. Using nominal abstraction, we can define this predicate as follows.

$$\begin{aligned} spec (monoTy T) nil T &\stackrel{\mu}{=} \top \\ (\nabla x.spec (polyTy P) (x :: L) (T x)) &\stackrel{\mu}{=} \nabla x.spec (P x) L (T x). \end{aligned}$$

Note that we use ∇ in the head of the second clause to associate the variable x at the head of the list L with its occurrences in the type $(T x)$. We then use ∇ in the body of this clause to allow for the recursive use of *spec*.

3.5.3 Arbitrarily Cascading Substitutions

Many reducibility arguments, such as Tait's proof of normalization for the simply typed λ -calculus [Tai67], are based on judgments over closed terms. During reasoning, however, one has often to work with open terms. To accommodate this requirement, the closed term judgment is extended to open terms by considering all possible closed instantiations of the open terms. When reasoning with \mathcal{G} , open terms are denoted by terms with nominal constants representing free variables. The general form of an open term is thus $M c_1 \cdots c_n$, and we want to consider all possible instantiations $M V_1 \cdots V_n$ where the V_i are closed terms. This type of arbitrary cascading substitutions is difficult to realize in reasoning systems where variables are given a simple type since M would have an arbitrary number of abstractions but the type of M would *a priori* fix that number of abstractions.

We can define arbitrary cascading substitutions in \mathcal{G} using nominal abstraction. In particular, we can define a predicate which holds on a list of pairs $\langle c_i, V_i \rangle$, a term with the form $M c_1 \cdots c_n$ and a term of the form $M V_1 \cdots V_n$. The idea is to iterate over the list of pairs and for each pair $\langle c, V \rangle$ use nominal abstraction to abstract c out of the first term and then substitute V before continuing. The following definition of the predicate *subst* is based on this idea.

$$\begin{aligned} \text{subst nil } T \ T &\stackrel{\mu}{=} \top \\ (\nabla x. \text{subst } (\langle x, V \rangle :: L) (T \ x) \ S) &\stackrel{\mu}{=} \text{subst } L (T \ V) \ S \end{aligned}$$

Given the definition of *subst* one may then show that arbitrary cascading substitutions have many of the same properties as normal higher-order substitutions. For instance, in the domain of the untyped λ -calculus, we can show that *subst* acts compositionally via the following lemmas.

$$\begin{aligned} \forall \ell, t, r, s. \text{subst } \ell (\text{app } t \ r) \ s \supset \exists u, v. (s = \text{app } u \ v \wedge \text{subst } \ell \ t \ u \wedge \text{subst } \ell \ r \ v) \\ \forall \ell, t, r. \text{subst } \ell (\text{abs } t) \ r \supset \exists s. (r = \text{abs } s \wedge \nabla z. \text{subst } \ell (t \ z) (s \ z)) \end{aligned}$$

Both of these lemmas have straightforward proofs by induction on *subst*.

We use this technique for describing arbitrary cascading substitutions again in Section 7.5 to formalize Girard's strong normalization argument for the simply-typed λ -calculus.

Some Properties of the Meta-logic

In this chapter we study some of the meta-theory of \mathcal{G} . There are two parts to our discussion. In the first part of the chapter, we prove various properties of the logic which show that the logic is well-designed and which are also useful when working within the logic. Most significantly, we prove the cut-elimination property for \mathcal{G} and then use this to establish the consistency of the logic. In the second part of the chapter we look at the question of how we can formally relate an object system to a potential encoding of it in \mathcal{G} . The naturalness of such a relationship is a strong recommendation for the meta-logic: it is ultimately this correspondence that allows us to use \mathcal{G} in establishing properties of an object system. Showing this type of relationship depends crucially on the earlier cut-elimination result which further justifies the emphasis we place on it.

4.1 Consistency of the Meta-logic

The logic \mathcal{G} , whose proof rules consist of the ones Figures 3.1, 3.2, 3.4, and 3.5, combines and extends the features in several logics such as $FO\lambda^{\Delta N}$ [MM00], $FO\lambda^{\Delta \nabla}$ [MT05], LG^ω [Tiu08] and Linc^- [TM09]. The relationship to Linc^- is of special interest to us below: \mathcal{G} is a conservative extension to this logic that is obtained by adding a treatment of the ∇ quantifier and the associated nominal constants and by generalizing the proof rules pertaining to equality to ones dealing with nominal abstraction. This correspondence will allow the proof of the critical meta-theoretic property of cut-elimination for Linc^- to be lifted to \mathcal{G} .

We shall actually establish three main properties of \mathcal{G} in this section. First, we shall show that the provability of a sequent is unaffected by the application of permutations of

nominal constants to formulas in the sequent. This property consolidates our understanding that nominal constants are quantified implicitly at the formula level; such quantification also renders irrelevant the particular names chosen for such constants. Second, we show that the application of substitution in a nominal capture-avoiding way preserves provability; by contrast, ordinary application of substitution does not have this property. Finally, we show that the *cut* rule can be dispensed with from the logic without changing the set of provable sequents. This implies that the left and right rules of the logic are balanced and moreover, that the logic is consistent. This is the main result of this section and its proof uses the earlier two results together with the argument for cut-elimination for Linc^- .

Several of our arguments will be based on induction on the heights of proofs. This measure is defined formally below. Notice that the height of a proof can be an infinite ordinal because the $\supseteq\mathcal{L}$ rule can have an infinite number of premises. Thus, we will be using a transfinite form of induction.

Definition 4.1.1. *The height of a derivation Π , denoted by $\text{ht}(\Pi)$, is 1 if Π has no premise derivations and is the least upper bound of $\{\text{ht}(\Pi_i) + 1\}_{i \in \mathcal{I}}$ if Π has the premise derivations $\{\Pi_i\}_{i \in \mathcal{I}}$ where \mathcal{I} is some index set.*

Many proof systems, such as Linc^- , include a weakening rule that allows formulas to be dropped (reading proofs bottom-up) from the left-hand sides of sequents. While \mathcal{G} does not include such a rule directly, its effect is captured in a strong sense as we show in the lemma below. Two proofs are to be understood here and elsewhere as having the same structure if they are isomorphic as trees, if the same rules appear at corresponding places within them and if these rules pertain to formulas that can be obtained one from the other via a renaming of eigenvariables and nominal constants.

Lemma 4.1.2. *Let Π be a proof of $\Sigma : \Gamma \longrightarrow B$ and let Δ be a multiset of formulas whose eigenvariables are contained in Σ . Then there exists a proof of $\Sigma : \Delta, \Gamma \longrightarrow B$ which has the same structure as Π . In particular $\text{ht}(\Pi) = \text{ht}(\Pi')$ and Π and Π' end with the same rule application.*

Proof. The lemma can be proved by an easy induction on $\text{ht}(\Pi)$. We omit the details. \square

The following lemma shows a strong form of the preservation of provability under permutations of nominal constants appearing in formulas, the first of our mentioned results.

Lemma 4.1.3. *Let Π be a proof of $\Sigma : B_1, \dots, B_n \longrightarrow B_0$ and let $B_i \approx B'_i$ for $i \in \{0, 1, \dots, n\}$. Then there exists a proof Π' of $\Sigma : B'_1, \dots, B'_n \longrightarrow B'_0$ which has the same structure as Π . In particular $\text{ht}(\Pi) = \text{ht}(\Pi')$ and Π and Π' end with the same rule application.*

Proof. The proof is by induction on $\text{ht}(\Pi)$ and proceeds specifically by considering the last rule used in Π . When this is a left rule, we shall assume without loss of generality that it operates on B_n .

The argument is easy to provide when the last rule in Π is one of $\perp\mathcal{L}$ or $\top\mathcal{R}$. If this rule is an *id*, i.e., if Π is of the form

$$\frac{B_j \approx B_0}{\Sigma : B_1, \dots, B_n \longrightarrow B_0} \textit{id}$$

then, since \approx is an equivalence relation, it must be the case that $B'_j \approx B'_0$. Thus, we can let Π' be the derivation

$$\frac{B'_j \approx B'_0}{\Sigma : B'_1, \dots, B'_n \longrightarrow B'_0} \textit{id}$$

If the last rule is a $\supseteq\mathcal{L}$ applied to a nominal abstraction $s \supseteq t$ that has no solutions, then, by Lemma 3.2.10, the sequent $\Sigma : B'_1, \dots, B'_n \longrightarrow B'_0$ also has a nominal abstraction with no solutions. Thus, Π' can be a derivation consisting of the single rule $\supseteq\mathcal{L}$. Lemma 3.2.10 similarly provides the key observation when the last rule in Π is an $\supseteq\mathcal{R}$.

All the remaining cases correspond to derivations of height greater than 1. We shall show that the last rule in Π in all these cases could also have $\Sigma : B'_1, \dots, B'_n \longrightarrow B'_0$ as a conclusion with the premises in this application of the rule being related via permutations in the way required by the lemma to the premises of the rule application in Π . The lemma then follows from the induction hypothesis.

In the case when the last rule in Π pertains to a binary connective—*i.e.*, when the rule is one of $\vee\mathcal{L}$, $\vee\mathcal{R}$, $\wedge\mathcal{L}$, $\wedge\mathcal{R}$, $\supset\mathcal{L}$ or $\supset\mathcal{R}$ —the desired conclusion follows naturally from the observation that permutations distribute over the connective. The proof can be similarly completed when a $\exists\mathcal{L}$, $\exists\mathcal{R}$, $\forall\mathcal{L}$ or $\forall\mathcal{R}$ rule ends the derivation, once we have noted that the application of permutations can be moved under the \exists and \forall quantifiers. For the *cut* and $c\mathcal{L}$ rules, we have to show that permutations can be extended to include the newly introduced formula in the upper sequent(s). This is easy: for the *cut* rule we use the identity permutation and for $c\mathcal{L}$ we replicate the permutation used to obtain B'_n from B_n .

The two remaining rules from the core logic are $\nabla\mathcal{L}$ and $\nabla\mathcal{R}$. The argument in these cases are similar and we consider only the later in detail. In this case, the last rule in Π is of the form

$$\frac{\Sigma : B_1, \dots, B_n \longrightarrow C[a/x]}{\Sigma : B_1, \dots, B_n \longrightarrow \nabla x.C} \nabla\mathcal{R}$$

where $a \notin \text{supp}(C)$. Obviously, $B'_0 = \nabla x.C'$ for some C' such that $C \approx C'$. Let d be a nominal constant such that $d \notin \text{supp}(C)$ and $d \notin \text{supp}(C')$. Such a constant must exist since both sets are finite. Then $C[a/x] \approx C[d/x] \approx C'[d/x]$. Thus the following

$$\frac{\Sigma : B'_1, \dots, B'_n \longrightarrow C'[d/x]}{\Sigma : B'_1, \dots, B'_n \longrightarrow \nabla x.C'} \nabla\mathcal{R}$$

is also an instance of the $\nabla\mathcal{R}$ rule and its upper sequent has the form desired.

The only case that remains to be treated when the last rule applies to a nominal abstraction is that of $\supseteq\mathcal{L}$ that has at least one upper sequent. In this case the rule has the structure

$$\frac{\{\Sigma\theta : B_1[[\theta]], \dots, B_{n-1}[[\theta]] \longrightarrow B_0[[\theta]] \mid \theta \text{ is a solution to } s \supseteq t\}}{\Sigma : B_1, \dots, s \supseteq t \longrightarrow B_0} \supseteq\mathcal{L}$$

Here we know that B'_n is a nominal abstraction $s' \supseteq t'$ that, by Lemma 3.2.10, has the same solutions as $s \supseteq t$. Further, by Lemma 3.2.3, $B_i[[\theta]] \approx B'_i[[\theta]]$ for any substitution θ . Thus

$$\frac{\{\Sigma\theta : B'_1[[\theta]], \dots, B'_{n-1}[[\theta]] \longrightarrow B'_0[[\theta]] \mid \theta \text{ is a solution to } s' \supseteq t'\}}{\Sigma : B'_1, \dots, s' \supseteq t' \longrightarrow B'_0} \supseteq\mathcal{L}$$

is also an instance of the $\supseteq\mathcal{L}$ rule and its upper sequents have the required property.

The arguments for the rules $def\mathcal{L}$ and $def\mathcal{R}$ are similar and we therefore only consider the case for the former rule in detail. Here, B_n must be of the form $p \vec{t}$ where p is a predicate symbol and the upper sequent must be identical to the lower one except for the fact that B_n is replaced by a formula of the form $B p \vec{t}$ where B contains no nominal constants. Further, B'_n is of the form $p \vec{s}$ where $p \vec{t} \approx p \vec{s}$. From this it follows that $B p \vec{t} \approx B p \vec{s}$ and hence that $\Sigma : B'_1, \dots, B'_n \longrightarrow B'_0$ can be the lower sequent of a rule whose upper sequent is related in the desired way via permutations to the upper sequent of the last rule in Π .

The only remaining rules to consider are \mathcal{IL} and \mathcal{CIR} . Once again, the arguments in these cases are similar and we therefore consider only the case for \mathcal{IL} in detail. Here, Π ends with a rule of the form

$$\frac{\vec{x} : B S \vec{x} \longrightarrow S \vec{x} \quad \Sigma : B_1, \dots, S \vec{t} \longrightarrow B_0}{\Sigma : B_1, \dots, p \vec{t} \longrightarrow B_0} \mathcal{IL}$$

where p is a predicate symbol defined by a clause of the form $\forall \vec{x}. p \vec{x} \stackrel{\mu}{=} B p \vec{x}$ and S contains no nominal constants. Now, B'_n must be of the form $p \vec{r}$ where $p \vec{t} \approx p \vec{r}$. Noting the proviso on S , it follows that $S \vec{t} \approx S \vec{r}$. But then the following

$$\frac{\vec{x} : B S \vec{x} \longrightarrow S \vec{x} \quad \Sigma : B'_1, \dots, S \vec{r} \longrightarrow B'_0}{\Sigma : B'_1, \dots, p \vec{r} \longrightarrow B'_0} \mathcal{IL}$$

is also an instance of the \mathcal{IL} rule and its upper sequents are related in the manner needed to those of the \mathcal{IL} rule used in Π . \square

Several rules in \mathcal{G} require the selection of new eigenvariables and nominal constants. Lemma 4.1.3 shows that we obtain what is essentially the same proof regardless of how we choose nominal constants in such rules so long as the local non-occurrence conditions are satisfied. A similar observation with regard to the choice of eigenvariables is also easily verified. We shall therefore identify below proofs that differ only in the choices of eigenvariables and nominal constants.

We now turn to the second of our desired results, the preservation of provability under substitutions.

Lemma 4.1.4. *Let Π be a proof of $\Sigma : \Gamma \longrightarrow C$ and let θ be a substitution. Then there is a proof Π' of $\Sigma\theta : \Gamma[\theta] \longrightarrow C[\theta]$ such that $\text{ht}(\Pi') \leq \text{ht}(\Pi)$.*

Proof. We show how to transform the proof Π into a proof Π' for the modified sequent. The transformation is by recursion on $\text{ht}(\Pi)$, the critical part of it being a consideration of the last rule in Π . The transformation is, in fact, straightforward in all cases other than when this rule is $\supseteq\mathcal{L}$, $\forall\mathcal{R}$, $\exists\mathcal{L}$, $\exists\mathcal{R}$, $\forall\mathcal{L}$, $\mathcal{I}\mathcal{L}$ and $\mathcal{C}\mathcal{I}\mathcal{R}$. In these cases, we simply apply the substitution in a nominal capture avoiding way to the lower and any possible upper sequents of the rule. It is easy to see that the resulting structure is still an instance of the same rule and its upper sequents are guaranteed to have proofs (of suitable heights) by induction.

Suppose that the last rule in Π is an $\supseteq\mathcal{L}$, *i.e.*, it is of the form

$$\frac{\{\Sigma\rho : \Gamma[\rho] \longrightarrow C[\rho] \mid \rho \text{ is a solution to } s \supseteq t\}}{\Sigma : \Gamma, s \supseteq t \longrightarrow C} \supseteq\mathcal{L}$$

Then the following

$$\frac{\{\Sigma(\theta \bullet \rho') : \Gamma[\theta \bullet \rho'] \longrightarrow C[\theta \bullet \rho'] \mid \rho' \text{ is a solution to } (s \supseteq t)[\theta]\}}{\Sigma\theta : \Gamma[\theta], (s \supseteq t)[\theta] \longrightarrow C[\theta]} \supseteq\mathcal{L}$$

is also an $\supseteq\mathcal{L}$ rule. Noting that if ρ' is a solution to $(s \supseteq t)[\theta]$, then $\theta \bullet \rho'$ is a solution to $s \supseteq t$, we see that the upper sequents of this rule are contained in the upper sequents of the rule in Π . It follows that we can construct a proof of the lower sequent whose height is less than or equal to that of Π .

The argument is similar in the cases when the last rule in Π is a $\forall\mathcal{R}$ or a $\exists\mathcal{L}$ so we consider only the former in detail. In this case the rule has the form

$$\frac{\Sigma, h : \Gamma \longrightarrow B[h \bar{c}/x]}{\Sigma : \Gamma \longrightarrow \forall x.B} \forall\mathcal{R}$$

where $\{\bar{c}\} = \text{supp}(\forall x.B)$. Let $\{\bar{a}\} = \text{supp}((\forall x.B)[\theta])$. Further, let h' be a new variable name. We assume without loss of generality that neither h nor h' appear in the domain or range of θ . Letting $\rho = \theta \cup \{\lambda\bar{c}.h' \bar{a}/h\}$, consider the structure

$$\frac{(\Sigma, h)\rho : \Gamma[\rho] \longrightarrow B[h \bar{c}/x][\rho]}{\Sigma\theta : \Gamma[\theta] \longrightarrow (\forall x.B)[\theta]}$$

The upper sequent here is equivalent under λ -conversion to $\Sigma\theta, h' : \Gamma[\theta] \longrightarrow (B[\theta])[h' \vec{a}/x]$ so this structure is, in fact, also an instance of the $\forall\mathcal{R}$ rule. Moreover, its upper sequent is obtained via substitution from the upper sequent of the rule in Π . The lemma then follows by induction.

The arguments for the cases when the last rule is an $\exists\mathcal{R}$ or an $\forall\mathcal{L}$ are similar and so we provide it explicitly only for the former. In this case, we have the rule

$$\frac{\Sigma, \mathcal{K}, \mathcal{C} \vdash t : \tau \quad \Sigma : \Gamma \longrightarrow B[t/x]}{\Sigma : \Gamma \longrightarrow \exists_{\tau} x.B} \exists\mathcal{R}$$

ending Π . Assuming that the substitution $(\exists_{\tau} x.B)[\theta]$ uses the permutation π to avoid the capture of nominal constants, consider the structure

$$\frac{\Sigma, \mathcal{K}, \mathcal{C} \vdash \pi.t : \tau \quad \Sigma\theta : \Gamma[\theta] \longrightarrow B[\theta][\pi.t/x]}{\Sigma\theta : \Gamma[\theta] \longrightarrow (\exists_{\tau} x.B)[\theta]}$$

This is also obviously an instance of the $\exists\mathcal{R}$ rule and its right upper sequent is related via substitution to that of the rule in Π . The lemma follows from these observations by induction.

The only remaining cases for the last rule are \mathcal{IL} and \mathcal{CIR} . The arguments in these cases are, yet again, similar and it suffices to make only the former explicit. In this case, the end of Π has the form

$$\frac{\vec{x} : B \quad S \vec{x} \longrightarrow S \vec{x} \quad \Sigma : \Gamma, S \vec{t} \longrightarrow C}{\Sigma : \Gamma, p \vec{t} \longrightarrow C} \mathcal{IL}$$

But then the following

$$\frac{\vec{x} : B \quad S \vec{x} \longrightarrow S \vec{x} \quad \Sigma\theta : \Gamma[\theta], (S \vec{t})[\theta] \longrightarrow C[\theta]}{\Sigma\theta : \Gamma[\theta], (p \vec{t})[\theta] \longrightarrow C[\theta]}$$

is also an instance of the \mathcal{IL} rule. Moreover, the same proof as in Π can be used for the left upper sequent and the right upper sequent has the requisite form for using the induction hypothesis. \square

The proof of Lemma 4.1.4 effectively defines a transformation of a derivation Π based on a substitution θ . We shall use the notation $\Pi[\theta]$ to denote the transformed derivation. Note that $\text{ht}(\Pi[\theta])$ can be less than $\text{ht}(\Pi)$. This may happen because the transformed version of a $\supseteq\mathcal{L}$ rule can have fewer upper sequents.

Corollary 4.1.5. *The following rules are admissible.*

$$\frac{\Sigma, h : \Gamma \longrightarrow B[h \vec{a}/x]}{\Sigma : \Gamma \longrightarrow \forall x.B} \forall\mathcal{R}^* \qquad \frac{\Sigma, h : \Gamma, B[h \vec{a}/x] \longrightarrow C}{\Sigma : \Gamma, \exists x.B \longrightarrow C} \exists\mathcal{L}^*$$

where $h \notin \Sigma$ and \vec{a} is any listing of distinct nominal constants which contains $\text{supp}(B)$.

Proof. Let Π be a derivation for $\Gamma \longrightarrow B[h \vec{a}/x]$, let h' be a variable that does not appear in Π , and let $\{\vec{c}\} = \text{supp}(B)$. By Lemma 4.1.4, $\Pi[\lambda\vec{a}.h' \vec{c}/h]$ is a valid derivation. Since \vec{a} contains \vec{c} , no nominal constants appear in the substitution $\{\lambda\vec{a}.h' \vec{c}/h\}$. It can now be seen that the last sequent in $\Pi[\lambda\vec{a}.h' \vec{c}/h]$ has the form $\Sigma, h' : \Gamma' \longrightarrow B'$ where $B' \approx B[h' \vec{c}/h]$ and Γ' results from replacing some of the formulas in Γ by ones that they are equivalent to under \approx . But then, by Lemma 4.1.3, there must be a derivation for $\Sigma, h' : \Gamma \longrightarrow B[h' \vec{c}/h]$. Using a $\forall\mathcal{R}$ rule below this we get a derivation for $\Sigma : \Gamma \longrightarrow \forall x.B$, verifying the admissibility of $\forall\mathcal{R}^*$. The argument for $\exists\mathcal{L}^*$ is analogous. \square

We now turn to the main result of this section, the redundancy from a provability perspective of the *cut* rule in \mathcal{G} . The usual approach to proving such a property is to define a set of transformations called cut reductions on derivations that leave the end sequent unchanged but that have the effect of pushing occurrences of *cut* up the proof tree to the leaves where they can be immediately eliminated. The difficult part of such a proof is showing that these cut reductions always terminate. In simpler sequent calculi such as the one for first-order logic, this argument can be based on an uncomplicated measure such as the size of the cut formula. However, the presence of definitions in a logic like \mathcal{G} renders this measure inadequate. For example, the following is a natural way to define a cut reduction

between a $def\mathcal{L}$ and a $def\mathcal{R}$ rule that work on the cut formula:

$$\frac{\frac{\Sigma : \Gamma \longrightarrow B \ p \ \vec{t}}{\Sigma : \Gamma \longrightarrow p \ \vec{t}} \text{ def}\mathcal{R} \quad \frac{\frac{\Pi'}{\Sigma : \Gamma \longrightarrow B \ p \ \vec{t}} \text{ def}\mathcal{L} \quad \frac{\Pi''}{\Sigma : B \ p \ \vec{t}, \Delta \longrightarrow C} \text{ def}\mathcal{L}}{\Sigma : p \ \vec{t}, \Delta \longrightarrow C} \text{ def}\mathcal{L}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{ cut}$$

↓

$$\frac{\frac{\Pi'}{\Sigma : \Gamma \longrightarrow B \ p \ \vec{t}} \quad \frac{\Pi''}{\Sigma : B \ p \ \vec{t}, \Delta \longrightarrow C}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{ cut}$$

Notice that $B \ p \ \vec{t}$, the cut formula in the new cut introduced by this transformation, could be more complex than $p \ \vec{t}$, the old cut formula. To overcome this difficulty, a more complicated argument based on the idea of reducibility in the style of Tait [Tai67] is often used. Tiu and Momigliano [TM09] in fact formulate a notion of parametric reducibility for derivations that is based on the Girard's proof of strong normalizability for System F [GTL89] and that works in the presence of the induction and co-induction rules for definitions. Our proof makes extensive use of this notion and the associated argument structure.

Theorem 4.1.6. *The cut rule can be eliminated from \mathcal{G} without affecting the provability relation.*

Proof. The relationship between \mathcal{G} and the logic Linc^- treated by Tiu and Momigliano can be understood as follows: Linc^- does not treat the ∇ quantifier and therefore has no rules for it. Consequently, it does not have nominal constants, it does not use raising over nominal constants in the rules $\forall\mathcal{R}$ and $\exists\mathcal{L}$, it has no need to consider permutations in the *id* (or initial) rule and has equality rules in place of nominal abstraction rules. The rules in \mathcal{G} other than the ones for ∇ , including the ones for definitions, induction, and co-induction, are essentially identical to the ones in Linc^- except for the additional attention to nominal constants.

Tiu and Momigliano's proof can be extended to \mathcal{G} in a fairly direct way since the addition of nominal constants and their treatment in the rules is quite modular and does not create any new complexities for the reduction rules. The main issues in realizing this extension

is building in the idea of identity under permutations of nominal constants and lifting the Linc^- notion of substitution on terms, sequents, and derivations to a form that avoids capture of nominal constants. The machinery for doing this has already been developed in Lemmas 4.1.3 and 4.1.4. In the rest of this proof we assume a familiarity with the argument for cut-elimination for Linc^- and discuss only the changes to the cut reductions of Linc^- to accommodate the differences.

The *id* rule in \mathcal{G} identifies formulas which are equivalent under \approx which is more permissive than equality under λ -convertibility that is used in the Linc^- initial rule. Correspondingly, we have to be a bit more careful about the cut reductions associated with the *id* (initial) rule. For example, consider the following reduction:

$$\frac{\frac{B \approx B'}{\Sigma : \Gamma, B \longrightarrow B'} \textit{id} \quad \frac{\Sigma : B', \Delta \longrightarrow C}{\Sigma : B', \Delta \longrightarrow C} \textit{cut}}{B, \Gamma, \Delta \longrightarrow C} \textit{cut} \quad \Longrightarrow \quad \frac{\Sigma : B', \Delta \longrightarrow C}{\Sigma : B', \Delta \longrightarrow C} \textit{cut}$$

This reduction has not preserved the end sequent. However, we know $B \approx B'$ and so we can now use Lemma 4.1.3 to replace Π' with a derivation of $\Sigma : B, \Delta \longrightarrow C$. Then we can use Lemma 4.1.2 to produce a derivation of $\Sigma : B, \Gamma, \Delta \longrightarrow C$ as desired. The changes to the cut reduction when *id* applies to the right upper sequent of the *cut* rule are similar.

The $\forall\mathcal{R}$ and $\exists\mathcal{L}$ rules of \mathcal{G} extend the corresponding rules of Linc^- by raising over nominal constants in the support of the quantified formula. The $\forall\mathcal{L}$ and $\exists\mathcal{R}$ rules of \mathcal{G} also extend the corresponding rules in Linc^- by allowing instantiations which contain nominal constants. Despite these changes, the cut reductions involving these quantifier rules remain unchanged for \mathcal{G} except for the treatment of essential cuts that involve an interaction between $\forall\mathcal{R}$ and $\forall\mathcal{L}$ and, similarly, between $\exists\mathcal{R}$ and $\exists\mathcal{L}$. The first of these is

treated as follows:

$$\frac{\frac{\frac{\Pi'}{\Sigma, h : \Gamma \longrightarrow B[h \vec{c}/x]}{\Sigma : \Gamma \longrightarrow \forall x.B} \forall \mathcal{R} \quad \frac{\frac{\Pi''}{\Sigma : \Delta, B[t/x] \longrightarrow C}}{\Sigma : \Delta, \forall x.B \longrightarrow C} \forall \mathcal{L}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{cut}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{cut}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{cut}$$

$$\Downarrow$$

$$\frac{\frac{\frac{\Pi'[\lambda \vec{c}.t/h]}{\Sigma : \Gamma \longrightarrow B[t/x]} \quad \frac{\Pi''}{\Sigma : \Delta, B[t/x] \longrightarrow C}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{cut}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{cut}}$$

The existence of the derivation $\Pi'[\lambda \vec{c}.t/h]$ (with height at most that of Π') is guaranteed by Lemma 4.1.4. The end sequent of this derivation is $\Sigma : \Gamma[\lambda \vec{c}.t/h] \longrightarrow B[h \vec{c}/x][\lambda \vec{c}.t/h]$. However, $\Gamma[\lambda \vec{c}.t/h] \approx \Gamma$ because h is new to Γ and $B[h \vec{c}/x][\lambda \vec{c}.t/h] \approx B[t/x]$ because $\{\vec{c}\} = \text{supp}(B)$ and so $\lambda \vec{c}.t$ has no nominal constants in common with $\text{supp}(B)$. Thus, by Lemma 4.1.3 and by an abuse of notation, we may consider $\Pi'[\lambda \vec{c}.t/h]$ to also be a derivation of $\Sigma : \Gamma \longrightarrow B[t/x]$. The reduction for a cut involving an interaction between an $\exists \mathcal{R}$ and an $\exists \mathcal{L}$ rule is analogous.

The logic \mathcal{G} extends the equality rules in Linc^- to treat the more general case of nominal abstraction. Our notion of nominal capture-avoiding substitution correspondingly generalizes the Linc^- notion of substitution, and we have shown in Lemma 4.1.4 that this preserves provability. Thus the reductions for nominal abstraction are the same as for equality, except that we use nominal capture-avoiding substitution in place of regular substitution. For example, the essential cut involving an interaction between an $\supseteq \mathcal{R}$ and an $\supseteq \mathcal{L}$ rule is treated as follows:

$$\frac{\frac{\Sigma : \Gamma \longrightarrow s \supseteq t}{\Sigma : \Gamma, \Delta \longrightarrow C} \supseteq \mathcal{R} \quad \left\{ \frac{\frac{\Pi_\theta}{\Sigma \theta : \Delta[\theta] \longrightarrow C[\theta]}}{\Sigma : \Delta, s \supseteq t \longrightarrow C} \supseteq \mathcal{L} \right\}}{\Sigma : \Gamma, \Delta \longrightarrow C} \text{cut}}{\Sigma : \Delta \longrightarrow C} \text{cut} \implies \frac{\Pi_\epsilon}{\Sigma : \Delta \longrightarrow C}$$

Here we know $s \supseteq t$ holds and thus ϵ , the identity substitution, is a solution to this nominal abstraction. Therefore we have the derivation Π_ϵ as needed. We can then apply Lemma 4.1.2 to weaken this derivation to one for $\Sigma : \Gamma, \Delta \longrightarrow C$. For the other cuts

involving nominal abstraction, we make use of the fact proved in Lemma 4.1.4 that nominal capturing avoiding substitution preserves provability. This allows us to commute other rules with $\triangleright\mathcal{L}$. For example, consider the following reduction of a cut where the upper right derivation uses an $\triangleright\mathcal{L}$ on a formula different from the cut formula:

$$\frac{\Sigma : \Gamma \xrightarrow{\Pi'} B \quad \left\{ \frac{\Sigma\theta : B[\theta], \Delta[\theta] \xrightarrow{\Pi_\theta} C[\theta]}{\Sigma : B, \Delta, s \triangleright t \rightarrow C} \right\}}{\Sigma : \Gamma, \Delta, s \triangleright t \rightarrow C} \text{ cut} \triangleright\mathcal{L}$$

$$\Downarrow$$

$$\left\{ \frac{\frac{\Sigma\theta : \Gamma[\theta] \xrightarrow{\Pi'[\theta]} B[\theta] \quad \Sigma\theta : B[\theta], \Delta[\theta] \xrightarrow{\Pi_\theta} C[\theta]}{\Sigma\theta : \Gamma[\theta], \Delta[\theta] \rightarrow C[\theta]} \text{ cut}}{\Sigma : \Gamma, \Delta, s \triangleright t \rightarrow C} \right\} \triangleright\mathcal{L}$$

Finally, \mathcal{G} has new rules for treating the ∇ -quantifier. The only reduction rule which deals specifically with either the $\nabla\mathcal{L}$ or $\nabla\mathcal{R}$ rule is the essential cut between both rules which is treated as follows:

$$\frac{\frac{\Sigma : \Gamma \xrightarrow{\Pi'} B[a/x] \quad \nabla\mathcal{R} \quad \frac{\Sigma : B[a/x], \Delta \xrightarrow{\Pi''} C}{\Sigma : \nabla x.B, \Delta \rightarrow C} \nabla\mathcal{L}}{\Sigma : \Gamma, \Delta \rightarrow C} \text{ cut}}{\Sigma : \Gamma, \Delta \rightarrow C} \text{ cut}$$

$$\Downarrow$$

$$\frac{\Sigma : \Gamma \xrightarrow{\Pi'} B[a/x] \quad \Sigma : B[a/x], \Delta \xrightarrow{\Pi''} C}{\Sigma : \Gamma, \Delta \rightarrow C} \text{ cut}.$$

With these changes, the cut-elimination argument for Linc^- extends to \mathcal{G} , *i.e.*, \mathcal{G} admits cut-elimination. □

The consistency of \mathcal{G} is an easy consequence of Theorem 4.1.6.

Corollary 4.1.7. *The logic \mathcal{G} is consistent, *i.e.*, not all sequents are provable in it.*

Proof. The sequent $\rightarrow \perp$ has no cut-free proof and, hence, no proof in \mathcal{G} . □

4.2 Adequacy of Encodings and Theorems in the Meta-logic

The logic \mathcal{G} provides various features such as λ -terms, definitions, and ∇ -quantification which form a convenient vehicle for encoding computational systems. With all these features, one might rightfully ask if our encodings in \mathcal{G} are faithful representations of the computational systems they describe. This kind of property for encodings, which is formally known as *adequacy*, is similar to the one that we have already encountered with respect to the specification logic. A proof of adequacy establishes a relationship between terms and judgments in an object system and their encoding in \mathcal{G} in such a way that we can relate reasoning results proven about the encoding to results about the original system. In this section we discuss adequacy in more detail, we describe the general approach to proving adequacy, and we present an example which illustrates some of the nuances which may arise for particular encodings.

At a philosophical level, adequacy is the method by which we assign meaning to our logic. Without adequacy, the logic has only behavior. Thus, one may naively ask a question such as, “what does the ∇ -quantifier mean?” To which a valid answer is that the ∇ -quantifier has no meaning in itself. It has the behavior of introducing a fresh nominal constant into a formula, but it is only through adequacy that we can interpret this behavior and provide it with some meaning. For instance, we might establish a correspondence between nominal constants in a \mathcal{G} formula and free variables in a typing judgment for an object system. In this setting, the meaning of ∇ -quantification can be interpreted as quantifying over fresh free variables.

A proof of adequacy for an encoding of an object system in \mathcal{G} consists of two parts:

1. the description of a bijection between the terms of the object system and their encoding in \mathcal{G} , and
2. a proof, based on this bijection, that a judgment in the object system holds if and only if its encoding in \mathcal{G} is provable.

For the second point, the cut-elimination result from Section 4.1 is of critical importance

since it allows us to restrict the sort of proofs we must consider. Without an independent proof of the cut-elimination property, proving adequacy would require establishing something like a cut-elimination theorem relative to each encoding that we wish to prove adequate.

Our ultimate objective is, of course, to prove theorems about the original system. However, this follows naturally from the proof of a relevant theorem in \mathcal{G} and the adequacy of encodings in the following way: 1) using adequacy, object level judgments are translated into \mathcal{G} formulas, 2) the relevant theorem proven in \mathcal{G} is used as a lemma on these formulas, and 3) using adequacy, the result of that lemma application is then translated back into an object level judgment. The end result is that the theorem is proven for the object system while most of the reasoning takes place within \mathcal{G} . The *cut* rule plays an essential role here as it allows us to use theorems proven in \mathcal{G} as lemmas which is very useful in reasoning and absolutely vital in the adequacy argument outlined above. It is for this reason that we cannot simply exclude the *cut* rule from our logic and hope to avoid the work involved in showing cut-elimination.

It is important to remember that adequacy is only an interface issue, *i.e.*, it is only a question about the “inputs” and “outputs” of \mathcal{G} . We show that an encoding of an object system (the “input”) is adequate and we use this to relate reasoning results in \mathcal{G} (the “output”) to results about the original system. Any auxiliary notions that we use within the logic in order to establish the results of interest do not matter for the purposes of adequacy. This is not to say that we do not care what goes on in between. Certainly we have designed the logic \mathcal{G} so that the intermediate reasoning can closely mimic the informal reasoning that is typically done. But in the end, the correctness of the reasoning that is performed depends only on the adequacy results and the cut-elimination property for \mathcal{G} .

As an example, let us now consider the adequacy of a proof of determinacy for an evaluation relation on untyped λ -terms. The evaluation relation of interest is presented in Figure 4.1. This example will be sufficient to illustrate the key issues involved in showing adequacy for an encoding in \mathcal{G} , while a more thorough example is presented later in

$$\frac{}{(\lambda x.r) \Downarrow (\lambda x.r)} \qquad \frac{m \Downarrow (\lambda x.r) \quad r[x := n] \Downarrow v}{(m \ n) \Downarrow v}$$

Figure 4.1: An evaluation relation for untyped λ -terms

$$\begin{aligned} \text{eval } (\text{abs } R) \text{ } (\text{abs } R) &\stackrel{\mu}{=} \top \\ \text{eval } (\text{app } M \ N) \ V &\stackrel{\mu}{=} \exists R. \text{eval } M \ (\text{abs } R) \wedge \text{eval } (R \ N) \ V \end{aligned}$$

Figure 4.2: An encoding of the evaluation relation in Figure 4.1

Section 6.5.1.

To represent untyped λ -terms in \mathcal{G} , we introduce the type tm along with the constructors $\text{app} : tm \rightarrow tm \rightarrow tm$ and $\text{abs} : (tm \rightarrow tm) \rightarrow tm$. Then we encode the evaluation relation as a definition for a predicate $\text{eval} : tm \rightarrow tm \rightarrow o$ as shown in Figure 4.2. Given this definition, we can prove the following determinacy result in \mathcal{G} :

$$\forall t, v_1, v_2. (\text{eval } t \ v_1 \wedge \text{eval } t \ v_2) \supset v_1 = v_2.$$

What we want to do is use this result to obtain a similar determinacy result for evaluation in the original system. We will develop the bijections and the associated adequacy lemmas below to be able to obtain such a translation.

We begin by defining a mapping $\ulcorner \cdot \urcorner$ from untyped λ -terms to their representation in \mathcal{G} :

$$\ulcorner x \urcorner = x \qquad \ulcorner t_1 \ t_2 \urcorner = \text{app } \ulcorner t_1 \urcorner \ulcorner t_2 \urcorner \qquad \ulcorner (\lambda x.t) \urcorner = \text{abs } (\lambda x. \ulcorner t \urcorner)$$

Note that we conflate the names of variables in untyped λ -terms with the corresponding names in \mathcal{G} . In truth, the bound variables of untyped λ -terms will be mapped to bound variables of type tm in \mathcal{G} , while the free variables of untyped λ -terms will be mapped to nominal constants of type tm in \mathcal{G} . Assuming a one-to-one correspondence between such terms, the above mapping is obviously bijective. Moreover, closed untyped λ -terms will

map to terms in \mathcal{G} without nominal constants and vice-versa. Thus our representation of untyped λ -terms is adequate.¹

Since we use the substitution mechanism of \mathcal{G} in the definition of *eval* to encode substitution on untyped λ -terms, we will later need to know that these two substitution relations are related via $\ulcorner \cdot \urcorner$ in the following sense.

Lemma 4.2.1. *Let t_1 and t_2 be untyped λ -terms. Then $\ulcorner t_1[x := t_2] \urcorner = \ulcorner t_1 \urcorner[\ulcorner t_2 \urcorner/x]$ where the substitution on the left takes place in the context of untyped λ -terms and the substitution on the right takes place in \mathcal{G} .*

Proof. The proof is by a straightforward induction on the structure of t_1 . □

Next we want to show an if-and-only-if relationship between the original evaluation judgment and its encoding in \mathcal{G} . This is formalized as follows.

Lemma 4.2.2. *$t \Downarrow v$ has a derivation if and only if $\longrightarrow \text{eval} \ulcorner t \urcorner \ulcorner v \urcorner$ is provable in \mathcal{G} .*

Proof. The proof in the forward direction is by straightforward induction on the derivation of $t \Downarrow v$.

For the backward direction we first note that $\longrightarrow \text{eval} \ulcorner t \urcorner \ulcorner v \urcorner$ must have a cut-free derivation by Theorem 4.1.6. The proof will be by induction on the height of this cut-free derivation. The cut-free derivation must end with *defR* though for ease of presentation we may suppose that it ends with *defR^p*.² The interesting case is when considering the second clause for *eval*, i.e., when $t = (m \ n)$ and the derivation ends as follows.

$$\frac{\frac{\frac{\longrightarrow \text{eval} \ulcorner m \urcorner (\text{abs } R) \quad \longrightarrow \text{eval} (R \ulcorner n \urcorner) \ulcorner v \urcorner}{\longrightarrow \text{eval} \ulcorner m \urcorner (\text{abs } R) \wedge \text{eval} (R \ulcorner n \urcorner) \ulcorner v \urcorner} \wedge \mathcal{R}}{\longrightarrow \exists r. \text{eval} \ulcorner m \urcorner (\text{abs } r) \wedge \text{eval} (r \ulcorner n \urcorner) \ulcorner v \urcorner} \exists \mathcal{R}}{\longrightarrow \text{eval} (\text{app} \ulcorner m \urcorner \ulcorner n \urcorner) \ulcorner v \urcorner} \text{defR}^p}$$

¹ A subtle but important point: we do not permit ∇ -quantification at type $tm \rightarrow tm$. Allowing this would mean that we will have terms in \mathcal{G} such as *abs c* for a nominal constant *c*. Since such a term cannot be the image of any untyped λ -term, the representation would then not be adequate.

² Note that cut-elimination was shown for the logic containing *defL* and *defR*, whereas *defL^p* and *defR^p* are only admissible additions to the logic.

Here R is a term of type $tm \rightarrow tm$. By the bijectivity of $\ulcorner \cdot \urcorner$, we know that $(abs R)$ is the representation of an untyped λ -term and thus we can apply the inductive hypothesis to the upper left sequent. Similarly, we can apply the inductive hypothesis to the upper right sequent after using Lemma 4.2.1 to convert $(R \ulcorner n \urcorner)$ to the representation of a substitution over untyped λ -terms. \square

It was essential to applying the inductive hypothesis in the proof of the lemma above that our mapping $\ulcorner \cdot \urcorner$ was a bijection. This property would not hold, for instance, if we restricted attention to only closed untyped λ -terms in the object language and we still allowed ∇ -quantification at type tm and, hence, admitted nominal constants of this type; specifically, we would have terms of type tm in \mathcal{G} that do not correspond to any closed untyped λ -terms. We would then not have been able to apply the inductive hypothesis in the proof of Lemma 4.2.2 because we would have to consider the possibility that particular occurrences of the $\exists\mathcal{R}$ rule generalize on terms of type tm that contain one or more nominal constants. However, it is still possible to use a proof in \mathcal{G} to establish a property about the original system even in this case. To do this, we would have to introduce a definition in \mathcal{G} for the class of terms of type tm that *do not* contain nominal constants and we would have to relativize the theorem we prove in \mathcal{G} to the class of terms satisfying this definition. From this perspective, adequacy is not always just a matter of mapping terms in the object system to terms in \mathcal{G} : we may need to map terms in the object system to terms satisfying a particular predicate in \mathcal{G} .

We now return to showing how a theorem in \mathcal{G} about the determinacy of the evaluation relation can be combined with the adequacy property for the encoding of untyped λ -terms to yield a theorem about the determinacy of the evaluation relation in the original calculus.

Theorem 4.2.3. *If $t \Downarrow v_1$ and $t \Downarrow v_2$ then v_1 equals v_2 .*

Proof. Suppose $t \Downarrow v_1$ and $t \Downarrow v_2$ both have derivations. By Lemma 4.2.2, that means we have proofs of $\longrightarrow eval \ulcorner t \urcorner \ulcorner v_1 \urcorner$ and $\longrightarrow eval \ulcorner t \urcorner \ulcorner v_2 \urcorner$. We also know from before that the

following has a derivation in \mathcal{G} :

$$\longrightarrow \forall t, v_1, v_2. (\text{eval } t \ v_1 \wedge \text{eval } t \ v_2) \supset v_1 = v_2.$$

Then using the rules $\forall\mathcal{L}$, $\supset\mathcal{L}$, $\wedge\mathcal{R}$, *id*, and *cut*, we can construct a derivation of $\longrightarrow \ulcorner v_1 \urcorner = \ulcorner v_2 \urcorner$. By Theorem 4.1.6 we know that $\longrightarrow \ulcorner v_1 \urcorner = \ulcorner v_2 \urcorner$ must have a cut-free derivation. This derivation must end with $\supseteq\mathcal{R}$ which applies only if $\ulcorner v_1 \urcorner$ is equal to $\ulcorner v_2 \urcorner$. Since $\ulcorner \cdot \urcorner$ is a bijection, this means that v_1 equals v_2 . \square

The discussion of adequacy in this section is reminiscent of an earlier discussion relative to the specification logic and hence raises the question of what, if anything, is different. The main observation here is that the logic \mathcal{G} is significantly richer than the hH^2 logic. In particular, when proving properties about an hH^2 specification, reasoning is conducted using general mathematical techniques, while for proving properties about an encoding in \mathcal{G} , the reasoning is conducted within \mathcal{G} itself. Thus, when working with \mathcal{G} , we use adequacy to connect results proven in \mathcal{G} with corresponding results about the original system. One may informally think of this as establishing adequacy for the theorems in \mathcal{G} relative to their counterparts about the original system.

An Interactive Theorem Prover for the Meta-logic

As part of this thesis, we have developed an interactive theorem prover called Abella for the logic \mathcal{G} [Gac08, Gac09]. Abella is implemented in OCaml and currently comprises approximately 4,000 lines of code. This system has been available to the public as open source software since March 2008 and has, in fact, been downloaded by several researchers. One of the key components of a theorem prover for \mathcal{G} is the treatment of nominal abstraction problems. We have discussed in Section 3.2.4 how the task of finding a solution to particular instances of the nominal abstraction predicate can be reduced to solving higher-order unification problems. Abella makes use of this reduction. Moreover, it assumes that the resulting unification problems lie within a restricted class known as the *higher-order pattern unification* class [Mil91, Nip93]. To solve such problems, it uses an algorithm developed by Nadathur and Linnell [NL05] that was initially implemented in Standard ML and that has subsequently been adapted to OCaml.

In this chapter, we briefly describe the architecture of Abella; this discussion serves the auxiliary purpose of building up ideas and terminology that we need for presenting applications of \mathcal{G} in Chapter 7. Abella requires proofs to be constructed through an interaction with a user. At any time, the state of a proof is represented as a collection of subgoals, all of which need to be proved for the overall proof to succeed. The user applies a *tactic* to a subgoal in order to make progress towards a completed proof. If we think of the proof as a derivation constructed in \mathcal{G} , then the subgoals in Abella correspond to sequents in the derivation which do not themselves have derivations as yet. Tactics then correspond to schemes for applying the rules of \mathcal{G} to such sequents in order to (incrementally) fill out their derivations.

There are two guiding principles for designing tactics in Abella:

1. they should correspond to some combination of rules from \mathcal{G} , and
2. they should correspond to natural reasoning steps.

For the most part, the rules of \mathcal{G} themselves resemble natural reasoning steps. The role of many tactics therefore, is simply to chain these together into larger steps. For example, given a goal of the form

$$\Sigma : \Gamma \longrightarrow \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$$

we may want to transition in one step into a goal of the following form:

$$\Sigma, \vec{x} : \Gamma, H_1, \dots, H_n \longrightarrow C.$$

Tactics are also used to group together many alternative rules. For example, a “case analysis” tactic may actually perform $\forall\mathcal{L}$, $\wedge\mathcal{L}$, $\perp\mathcal{L}$, $def\mathcal{L}$, $\exists\mathcal{L}$, or $\nabla\mathcal{L}$ based on the structure of the formula to which it is applied.

In the rest of this chapter, we describe two areas in which tactics greatly massage the rules of \mathcal{G} into a convenient form. The first concerns how hypotheses or lemmas of a particular form can be applied to other hypotheses. The second concerns a treatment of induction and co-induction which can naturally accommodate even sophisticated inductive and co-inductive arguments.

5.1 A Framework for Using Lemmas

Suppose we have a hypothesis of the form

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$$

and further hypotheses H'_1, \dots, H'_n which match H_1, \dots, H_n under proper instantiations of the \vec{x} . Then we would like a tactic to apply the first hypothesis to H'_1, \dots, H'_n , *i.e.*, a tactic which finds the proper instantiations for \vec{x} and chains together the rules of \mathcal{G} to generate a

new hypothesis C' that is the corresponding instantiation of C . To be more specific, let Γ contain H'_1, \dots, H'_n . Then we want a tactic which constructs the derivation

$$\frac{\frac{\frac{\Pi_n}{\Gamma \longrightarrow H_n[\vec{t}/\vec{x}]} \quad \frac{\Pi}{\Gamma, C[\vec{t}/\vec{x}] \longrightarrow B}}{\Gamma, H_n[\vec{t}/\vec{x}] \supset C[\vec{t}/\vec{x}] \longrightarrow B} \supset \mathcal{L}}{\frac{\frac{\Pi_1}{\Gamma \longrightarrow H_1[\vec{t}/\vec{x}]} \quad \vdots \quad \frac{\Pi_n}{\Gamma, H_n[\vec{t}/\vec{x}] \supset C[\vec{t}/\vec{x}] \longrightarrow B}}{\Gamma, H_1[\vec{t}/\vec{x}] \supset \dots \supset H_n[\vec{t}/\vec{x}] \supset C[\vec{t}/\vec{x}] \longrightarrow B} \supset \mathcal{L}}{\frac{\Gamma, \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C \longrightarrow B}{\Gamma, \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C \longrightarrow B} \forall \mathcal{L}}$$

where each Π_i is just the identity rule. In an actual implementation, this construction may be accomplished by replacing the variables \vec{x} with instantiatable meta-variables \vec{v} and using unification between $H_i[\vec{v}/\vec{x}]$ and H'_i to determine specific values for the \vec{v} .

Using the above construction, we can think of more sophisticated ways in which H'_i will match $H_i[\vec{t}/\vec{x}]$. All that we effectively require is that a derivation of $H'_i \longrightarrow H_i[\vec{t}/\vec{x}]$ can be constructed automatically. One useful case arises when $H_i[\vec{t}/\vec{x}]$ has the form $\nabla \vec{z}. H''_i$ for some formula H''_i , and where H'_i will match $H''_i[\vec{a}/\vec{z}]$ for some distinct listing of nominal constants \vec{a} which are not in the support of H''_i . If such a case holds, then a derivation of $H'_i \longrightarrow \nabla \vec{z}. H''_i$ can be constructed by repeated use of $\nabla \mathcal{R}$ followed by the initial rule. As before, in an actual implementation, we might be working with $H_i[\vec{v}/\vec{x}] = \nabla \vec{z}. H'''_i$ where \vec{v} are instantiatable meta-variables. In such a case, we can determine proper instantiations for the \vec{v} by solving the nominal abstraction $\lambda \vec{z}. H'''_i \supseteq H'_i$.

Typically, lemmas also have the form

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset C.$$

If we have independently proven such a lemma, then we can use *cut* to bring it in as a hypothesis at any time. Then we can use this lemma together with other hypotheses as described above so as to derive a suitable instance of C .

By supporting an easy and direct use of lemmas, the system encourages large proofs to be broken down into separate lemmas which build towards a final result. In practice, these

intermediate lemmas and the points at which they are used are often the most important pieces in the development of a proof. In fact, the structure of most arguments is the following: use the induction rule, then perform case analysis and finally use particular lemmas and the induction hypothesis to obtain the goal. Thus in actual presentation of proofs, the detailed proof steps are hidden by default, and instead the focus is on the series of lemmas that lead to the desired conclusions [Gac09].

A final point worth mentioning is that we deliberately consider formulas of the form

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$$

even though the following form is equivalent and perhaps more easy to read for humans:

$$\forall \vec{x}. H_1 \wedge \dots \wedge H_n \supset C.$$

The reason we prefer the first form is two-fold: 1) it has a recursive structure which is easier to work with in an implementation, and 2) in the degenerate case the when $n = 0$, then first form is $\forall \vec{x}. C$ while the second is the more obtuse $\forall \vec{x}. \top \supset C$. In the future, we shall always work with formulas in the first form.

5.2 An Annotation Based Scheme for Induction

The rule for induction in \mathcal{G} can be somewhat awkward to use from a traditional reasoning perspective: it requires one to formulate an invariant S , prove that S is truly an invariant, and then use S in place of the predicate that was given by the inductive definition under consideration. In traditional reasoning, these steps are often merged into a single idea which is called simply “reasoning by induction.” In this section we present a treatment of induction based on annotating formulas which aims to capture this simplified approach to induction. Further, we justify this treatment by translating the tactic that underlies it into a particular application of the logical rules of \mathcal{G} .

Let us consider a very simple inductive argument to introduce the annotation based

treatment of induction. Suppose we define *even* and *odd* on natural numbers as follows.

$$\begin{array}{ll} \text{even } z \stackrel{\mu}{=} \top & \text{odd } (s z) \stackrel{\mu}{=} \top \\ \text{even } (s (s N)) \stackrel{\mu}{=} \text{even } N & \text{odd } (s (s N)) \stackrel{\mu}{=} \text{odd } N \end{array}$$

Suppose we want to prove that if N is even then $s N$ is odd:

$$\forall N. \text{even } N \supset \text{odd } (s N).$$

The proof is by induction on the *even* hypothesis. The annotation based treatment of this induction proceeds by creating a new hypothesis (called the inductive hypothesis) of the form

$$\forall N. (\text{even } N)^* \supset \text{odd } (s N)$$

and changing the goal to

$$\forall N. (\text{even } N)^{\textcircled{a}} \supset \text{odd } (s N).$$

The $*$ annotation indicates that the inductive hypothesis can only be applied to an argument which has that same annotation. The \textcircled{a} annotation indicates that when this atomic formula is subjected to case analysis, any recursive calls to *even* will be annotated with $*$. In all other respects, the annotations are to be ignored, and besides the induction tactic there is no way to introduce these annotations. In this way, Abella allows the inductive hypothesis to be applied only when the distinguished inductive argument has been subjected to case analysis.

Coming back to the proof, let us abbreviate the inductive hypothesis by *IH*. Then we can eventually do case analysis on the *even* hypothesis which leads to the following sequents.

$$IH \longrightarrow \text{odd } (s z) \qquad IH, (\text{even } N')^* \longrightarrow \text{odd } (s (s (s N')))$$

The first of these is easily provable. In the second we apply the inductive hypothesis which is allowed based on the annotations, and this produces a hypothesis of *odd* ($s N'$). The rest of the proof is straightforward.

We will now show how this annotation based treatment of induction is sound by translating it to rules from \mathcal{G} . Suppose we want to prove the following.

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$$

Further, assume that we want to do this by induction on $H_i = p \vec{t}$ where p is defined by $\forall \vec{y}. p \vec{y} \stackrel{\mu}{=} B p \vec{y}$. Then we define the invariant S as

$$S = \lambda \vec{y}. \forall \vec{x}. \vec{y} = \vec{t} \supset H_1 \supset \dots \supset H_n \supset C$$

where $\vec{y} = \vec{t}$ denotes an equality between appropriately typed tuples involving the indicated terms. Using this invariant, we can construct the following derivation in \mathcal{G} .

$$\frac{\frac{\frac{\frac{\frac{\Pi_S}{\vec{y} : B S \vec{y} \longrightarrow S \vec{y}}}{\vec{x} : p \vec{t}, H_1, \dots, H_n \longrightarrow C} \text{c}\mathcal{L}}{\vec{x} : H_1, \dots, H_n \longrightarrow C} \supset \mathcal{R}}{\vec{x} : \cdot \longrightarrow H_1 \supset \dots \supset H_n \supset C} \supset \mathcal{R}}{\cdot : \cdot \longrightarrow \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C} \forall \mathcal{R}}{\frac{\frac{\Pi}{\vec{x} : S \vec{t}, H_1, \dots, H_n \longrightarrow C}}{\vec{x} : S \vec{t}, H_1, \dots, H_n \longrightarrow C} \text{I}\mathcal{L}}{\cdot : \cdot \longrightarrow \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C} \forall \mathcal{R}} \text{I}\mathcal{L}$$

Now, the missing derivation Π is trivial to construct using $\forall \mathcal{L}$, $\supset \mathcal{L}$, $\supseteq \mathcal{R}$ and id . We fill in the other missing derivation, Π_S , as follows:

$$\frac{\frac{\frac{\frac{\Pi'_S}{\vec{x} : B S \vec{t}, H_1, \dots, H_n \longrightarrow C}}{\vec{x}, \vec{y} : B S \vec{y}, \vec{y} = \vec{t}, H_1, \dots, H_n \longrightarrow C} \supseteq \mathcal{L}_{CSNAS}}{\vec{x}, \vec{y} : B S \vec{y} \longrightarrow \vec{y} = \vec{t} \supset H_1 \supset \dots \supset H_n \supset C} \supset \mathcal{R}}{\vec{y} : B S \vec{y} \longrightarrow \forall \vec{x}. \vec{y} = \vec{t} \supset H_1 \supset \dots \supset H_n \supset C} \forall \mathcal{R}} \supset \mathcal{R}$$

Then we fill in Π'_S based on the content of the inductive argument carried out within the annotation based scheme.

To complete this picture, let us consider how uses of the induction hypothesis in the annotation based treatment of induction correspond to making use of the hypothesis $B S \vec{t}$ in constructing the derivation Π'_S . Within the annotation based treatment, the induction hypothesis has the following form:

$$\forall \vec{x}. H_1 \supset \dots \supset (p \vec{t})^* \supset \dots \supset H_n \supset C.$$

Given the restrictions on annotations, this hypothesis can only be used if instantiations are found for the \vec{x} such that $(p \vec{t})^*$ is equal to one of the $(p \vec{s})^*$ which occurs as a result of case analysis on the original hypothesis of $(p \vec{t})^\circledast$. By understanding case analysis as $\text{def}\mathcal{L}$ in \mathcal{G} , we see that these occurrences of $(p \vec{s})^*$ for which the induction hypothesis is applicable are exactly those occurrences of p in $B p \vec{t}$. In turn, the induction invariant is available for those same occurrences of p when constructing the derivation Π'_S , which is precisely what is realized via the hypothesis $B S \vec{t}$. Thus the annotation based treatment of induction can be translated to a proper derivation in \mathcal{G} , and therefore the treatment is sound.

5.3 Extensions to the Basic Scheme for Induction

The treatment of induction that we have just described can be extended in a few different ways. Each of these brings some additional complications to the construction of a corresponding derivation in \mathcal{G} . For clarity of presentation, we shall consider each extension in isolation, but we note that they could all be combined.

5.3.1 Induction on a Predicate in the Scope of Generic Quantifiers

We can extend the annotation based treatment of induction to work with predicates which occur underneath ∇ -quantifiers. Suppose again we want to prove

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$$

where, this time, we want to induct on $H_i = \nabla \vec{z}. p \vec{t}$ where p is defined by $\forall \vec{y}. p \vec{y} \stackrel{\mu}{=} B p \vec{y}$. Within the annotation based treatment, nothing needs to be changed to cater to this situation: $(p \vec{t})$ is annotated with $*$ in the inductive hypothesis and with \circledast in the goal and the rules for applying an inductive hypothesis with ∇ s over the inductive argument are the same as those described in Section 5.1.

We justify this treatment by defining the invariant S as follows.

$$S = \lambda \vec{y}. \forall \vec{x}. (\lambda \vec{z}. \vec{t} \succeq \vec{y}) \supset H_1 \supset \dots \supset H_n \supset C$$

We can follow the original construction with this invariant, and the only wrinkle is in the construction of Π_S , a derivation of $\vec{y} : B S \vec{y} \longrightarrow S \vec{y}$. We construct this as follows.

$$\frac{\frac{\frac{\Pi'_S}{\vec{x} : B S \vec{t}, H_1, \dots, H_n \longrightarrow C}}{\vec{x}, \vec{y} : B S \vec{y}, (\lambda \vec{z}. \vec{t} \triangleright \vec{y}), H_1, \dots, H_n \longrightarrow C} \triangleright \mathcal{L}_{CSNAS}}{\vec{x}, \vec{y} : B S \vec{y} \longrightarrow (\lambda \vec{z}. \vec{t} \triangleright \vec{y}) \supset H_1 \supset \dots \supset H_n \supset C} \supset \mathcal{R}}{\vec{y} : B S \vec{y} \longrightarrow \forall \vec{x}. (\lambda \vec{z}. \vec{t} \triangleright \vec{y}) \supset H_1 \supset \dots \supset H_n \supset C} \forall \mathcal{R}$$

Here and in the future, we simplify the presentation by treating the free variables \vec{z} in \vec{t} as nominal constants. Now we fill in Π'_S based on the content of the inductive argument carried out within the annotation based scheme. After using $\nabla \mathcal{L}$ and case analysis on $H_i = \nabla \vec{z}. p \vec{t}$ we will have $B p \vec{t}$ and also $B S \vec{t}$. Thus we have the inductive hypothesis available for the recursive calls to p . The restrictions enforced by the nominal abstraction in S are the same as those enforced when applying hypotheses which have embedded occurrences of ∇ , as per the discussion in Section 5.1. Thus this treatment is sound.

5.3.2 Induction in the Presence of Additional Premises

We extend the annotation based treatment of induction by allowing induction in the context of other hypotheses. That is, instead of proving $\cdot : \cdot \longrightarrow \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$, we prove

$$\Sigma : \Gamma \longrightarrow \forall \vec{x}. H_1 \supset \dots \supset H_n \supset C$$

Within the annotation based treatment of induction, there is nothing that needs to be changed to handle this case: we annotate the goal and generate an annotated induction hypothesis which is added to the other hypotheses.

To verify the soundness of this extension, we reconstruct the original soundness argument using the invariant $S' = \lambda \vec{y}. \forall \Sigma. \bigwedge \Gamma \supset S \vec{y}$ where S is the invariant prescribed in the original construction and $\bigwedge \Gamma$ denotes the conjunction of all formulas in Γ . Then the only significant change in the construction is that Π_S needs to be a derivation of $\vec{y} : B S' \vec{y} \longrightarrow S' \vec{y}$. Using $\forall \mathcal{R}$, $\supset \mathcal{R}$, and $\wedge \mathcal{L}$ this becomes $\Sigma, \vec{y} : \Gamma, B S' \vec{y} \longrightarrow S' \vec{y}$. Finally,

we know $\forall \Sigma. \forall \vec{y}. \bigwedge \Gamma \supset S' \vec{y} \supset S \vec{y}$ by the definition of S' , and since B does not use its first argument negatively (due to stratification), we know $\forall \Sigma. \forall \vec{y}. \bigwedge \Gamma \supset B S' \vec{y} \supset B S \vec{y}$. By using this, all we have left to show is $\Sigma, \vec{y} : \Gamma, B S \vec{y} \longrightarrow S \vec{y}$ which we can unfold as in the original construction and what is left matches the work done in the annotation based treatment.

5.3.3 Delayed Applications of the Induction Hypothesis

Another extension we can make is to allow the inductive hypothesis to be applied not just for immediate recursive calls, but for finitely nested ones as well. This is supported in the annotation based treatment by saying that case analysis on a hypothesis with a $*$ annotation results in recursive calls which also have the $*$ annotation. For example, taking *even* and *odd* as before, suppose we want to prove every natural number is either even or odd:

$$\forall N. \text{nat } N \supset \text{even } N \vee \text{odd } N.$$

The proof is by induction on $\text{nat } N$. Thus we have the inductive hypothesis *IH* as follows:

$$\forall N. (\text{nat } N)^* \supset \text{even } N \vee \text{odd } N.$$

When we perform case analysis on the hypothesis $(\text{nat } N)^{\textcircled{a}}$ in the goal it leads to the following sequents.

$$IH \longrightarrow \text{even } z \vee \text{odd } z \qquad IH, (\text{nat } N')^* \longrightarrow \text{even } (s N') \vee \text{odd } (s N')$$

The first sequent is trivial to prove, and we can apply case analysis to $(\text{nat } N')^*$ in the second to get the following two sequents.

$$IH \longrightarrow \text{even } (s z) \vee \text{odd } (s z) \qquad IH, (\text{nat } N'')^* \longrightarrow \text{even } (s (s N'')) \vee \text{odd } (s (s N''))$$

Again the first sequent is trivial. In the second sequent we can apply the inductive hypothesis to get the sequent

$$\dots, \text{even } N'' \vee \text{odd } N'' \longrightarrow \text{even } (s (s N'')) \vee \text{odd } (s (s N'')).$$

Now we can apply $\forall\mathcal{L}$ and the rest of the proof is trivial to construct.

The justification for this extension in \mathcal{G} is to use the invariant $S' = \lambda\vec{y}.S \vec{y} \wedge B S \vec{y}$ in the original construction where S is the original invariant. Then only significant change in the construction is that we are required to fill out the following derivation

$$\frac{\frac{\Pi_1}{\vec{y} : B S' \vec{y} \longrightarrow S \vec{y}} \quad \frac{\Pi_2}{\vec{y} : B S' \vec{y} \longrightarrow B S \vec{y}}}{\vec{y} : B S' \vec{y} \longrightarrow S' \vec{y}} \wedge\mathcal{R}}$$

Now note that $\forall\vec{x}. S' \vec{x} \supset S \vec{x}$ and $\forall\vec{x}. S' \vec{x} \supset B S \vec{x}$ are both trivially provable after expanding the definition of S' . Since B does not allow its first argument to occur negatively (due to stratification) this means we can inductively construct derivations of $\forall\vec{x}. B S' \vec{x} \supset B S \vec{x}$ and $\forall\vec{x}. B S' \vec{x} \supset B (B S) \vec{x}$. The construction of the derivation Π_2 follows directly from the first of these. The derivation Π_1 contains the real content of the inductive proof. If case analysis is eventually used on $H_i = p \vec{t}$ in this derivation then the \vec{y} will have been instantiated with \vec{t} so that we have the hypothesis $B S' \vec{t}$. Thus we will have $B S \vec{t}$ which is the regular inductive hypothesis and also $B (B S) \vec{t}$ which is the inductive hypothesis applied to recursive calls nested at depth two. This depth can be extended to any finite number by repeating the above construction with the appropriate S' .

5.3.4 Nested Inductions

The use of annotations can be extended to allow nested inductions. For example, suppose we define the following predicate *ack* for computing the Ackermann function.

$$\begin{aligned} \text{ack } z N (s N) &\stackrel{\mu}{=} \top \\ \text{ack } (s M) z R &\stackrel{\mu}{=} \text{ack } M (s z) R \\ \text{ack } (s M) (s N) R &\stackrel{\mu}{=} \exists R'. \text{ack } (s M) N R' \wedge \text{ack } M R' R \end{aligned}$$

And suppose we want to prove that this function is total in its first two arguments:

$$\forall M, N. \text{nat } M \supset \text{nat } N \supset \exists R. \text{nat } R \wedge \text{ack } M N R$$

The proof requires an outer induction on $\text{nat } M$ and an inner induction on $\text{nat } N$. In the annotation based treatment of induction, this is realized as follows. Applying induction to $\text{nat } M$ produces the outer inductive hypothesis

$$\forall M, N. (\text{nat } M)^* \supset \text{nat } N \supset \exists R. \text{nat } R \wedge \text{ack } M N R$$

and the goal

$$\forall M, N. (\text{nat } M)^{\textcircled{a}} \supset \text{nat } N \supset \exists R. \text{nat } R \wedge \text{ack } M N R.$$

Then applying induction to $\text{nat } N$ in this goal produces the inner inductive hypothesis

$$\forall M, N. (\text{nat } M)^{\textcircled{a}} \supset (\text{nat } N)^{**} \supset \exists R. \text{nat } R \wedge \text{ack } M N R$$

and the goal

$$\forall M, N. (\text{nat } M)^{\textcircled{a}} \supset (\text{nat } N)^{\textcircled{a}\textcircled{a}} \supset \exists R. \text{nat } R \wedge \text{ack } M N R.$$

The treatment of annotations is the same as described before. The annotations $*$ and $**$ as well as \textcircled{a} and $\textcircled{a}\textcircled{a}$ are considered distinct and unrelated. Thus the outer inductive hypothesis applies as before, while the inner inductive hypothesis can only be applied to $(\text{nat } M)^{\textcircled{a}}$ from the goal and something with the $**$ annotation which can only come from case analysis on $(\text{nat } N)^{\textcircled{a}\textcircled{a}}$.

We will use this treatment to finish the proof of totality for the Ackermann function. Let IH and IH' be the outer and inner induction hypotheses, respectively. Then the interesting part of the proof comes after we have done case analysis on both $(\text{nat } M)^{\textcircled{a}}$ and $(\text{nat } N)^{\textcircled{a}\textcircled{a}}$. In particular, in the case where $M = s M'$ and $N = s N'$ we need to prove the following sequent.

$$IH, IH', (\text{nat } (s M'))^{\textcircled{a}}, (\text{nat } M')^*, (\text{nat } N')^* \longrightarrow \exists R. \text{nat } R \wedge \text{ack } (s M') (s N') R$$

Note that we must have performed contraction on $(\text{nat } M)^{\textcircled{a}}$ prior to case analysis in order to keep a copy of it. Then we can apply the inner induction hypothesis to $(\text{nat } (s M'))^{\textcircled{a}}$ and $(\text{nat } N')^*$ to get the hypotheses $\text{nat } R'$ and $\text{ack } (s M') N R'$ for some new variable R' .

Applying the outer inductive hypothesis to $(\text{nat } M')^*$ and $\text{nat } R'$ produces the hypotheses $\text{nat } R''$ and $\text{ack } M' R' R''$. Then we can apply $\exists\mathcal{R}$ with $R = R''$, and the rest of the proof is trivial.

We now justify the annotation based treatment of nested induction. As in the original construction, suppose we want to prove

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset C.$$

And suppose the proof is by an outer induction on $H_i = p \vec{t}$ where p is defined by $\forall \vec{y}. p \vec{y} \stackrel{\mu}{=} B p \vec{y}$ and an inner induction on $H_j = q \vec{s}$ where q is defined by $\forall \vec{z}. q \vec{z} \stackrel{\mu}{=} B' q \vec{z}$. We proceed with the original construction using the original invariant S for the outer induction. This leaves us with a need to prove the following.

$$\vec{x} : B S \vec{t}, H_1, \dots, H_n \longrightarrow C$$

Now we apply contraction on $H_j = q \vec{s}$ and induct on one of the copies using the following invariant.

$$S' = \lambda \vec{z}. \forall \vec{x}. B S \vec{t} \supset \vec{z} = \vec{s} \supset H_1 \supset \dots \supset H_n \supset C$$

The only non-trivial sequent to prove will be $\vec{z} : B' S' \vec{z} \longrightarrow S' \vec{z}$. Applying $\forall\mathcal{R}$, $\supset\mathcal{R}$, and $\supseteq\mathcal{L}_{CSNAS}$, this reduces to showing

$$\vec{x} : B' S' \vec{s}, B S \vec{t}, H_1, \dots, H_n \longrightarrow C$$

Now from $B S \vec{t}$ we have the outer induction invariant available for the recursive calls to p which arise from case analysis on $H_i = p \vec{t}$. From $B' S' \vec{s}$ we have the inner induction invariant available for the recursive calls to q which arise from case analysis on $H_j = q \vec{s}$. The caveat is that the inner induction invariant S' requires a proof of $B S \vec{t}$. This constrains the variables \vec{x} in the inner induction variant based on their occurrences in \vec{t} . In the annotation based treatment, the requirement of a hypothesis with a @ annotation enforces exactly this condition for the inner inductive hypothesis.



Figure 5.1: Transition diagrams for two different processes

5.4 An Annotation Based Scheme for Co-induction

We can also use annotations to treat co-induction. To illustrate how this works, we will take an example from the domain of process calculi. Let us consider the two processes depicted in Figure 5.4. Here the circles represent states and the arrows represent possible transitions between those states. We say that a P is *simulated by* a state Q if for every transition that P can make to a state P' there exists a state Q' to which Q can transition and such that P' is simulated by Q' . We consider the notion of simulation as co-inductive so a state can be simulated by another state even if both have infinite (possibly cyclic) chains of transitions from them. Suppose then, that we want to show that the state p_0 is simulated by the state q_0 . We can see that this is true by considering all possible transitions from these states and recognizing that p_1 is simulated by the state q_1 .

Let us now think of conducting this example in \mathcal{G} . We start by encoding the two processes using the following definition of *step*.

$$\begin{array}{lll} \text{step } p_0 \ p_1 \triangleq \top & \text{step } p_1 \ p_0 \triangleq \top & \\ \text{step } q_0 \ q_1 \triangleq \top & \text{step } q_1 \ q_0 \triangleq \top & \text{step } q_1 \ q_2 \triangleq \top \end{array}$$

Then we define simulation as a co-inductive predicate $\text{sim } P \ Q$ which holds when the process P is simulated by the process Q . The precise definition is as follows.

$$\text{sim } P \ Q \stackrel{\nu}{=} \forall P'. \text{step } P \ P' \supset \exists Q'. \text{step } Q \ Q' \wedge \text{sim } P' \ Q'$$

Our goal is then to prove $\text{sim } p_0 \ q_0$ which we generalize based on the argument sketched above into the following formula to prove:

$$\forall P, Q. (P = p_0 \wedge Q = q_0) \vee (P = p_1 \wedge Q = q_1) \supset \text{sim } P \ Q.$$

If we apply annotation based co-induction to this goal we get the co-inductive hypothesis

$$\forall P, Q. (P = p_0 \wedge Q = q_0) \vee (P = p_1 \wedge Q = q_1) \supset (\text{sim } P \ Q)^+$$

and the new goal

$$\forall P, Q. (P = p_0 \wedge Q = q_0) \vee (P = p_1 \wedge Q = q_1) \supset (\text{sim } P \ Q)^{\#}.$$

Note that the annotations for co-induction apply to the consequent of an implication rather than one of the hypotheses. The rules for these new annotations are as follows. If we unfold (*i.e.*, use $\text{def}\mathcal{R}$ on) a co-inductive definition with a $\#$ annotation then all of its recursive calls have the $+$ annotation. Hypotheses with a $+$ annotation are obtained from the co-inductive hypothesis and can *only* be used to match a goal with the $+$ annotation. For all other purposes, the annotations can be ignored. The proof of the above simulation eventually reduces to the following two sequents where CH is the co-inductive hypothesis.

$$CH \longrightarrow (\text{sim } p_0 \ q_0)^{\#} \qquad CH \longrightarrow (\text{sim } p_1 \ q_1)^{\#}$$

The proofs of these two sequents are similar, so we will consider only the first one. Here if we apply $\text{def}\mathcal{R}$ we will eventually end up with the sequent

$$CH \longrightarrow (\text{sim } p_1 \ q_1)^+.$$

At this point we can apply the co-inductive hypothesis to get a hypothesis which will match the goal.

We can justify the annotation based treatment of co-induction by translating it into appropriate rules from \mathcal{G} . Suppose we want to prove the following where p is defined by $\forall \vec{y}. p \ \vec{y} \stackrel{\vee}{=} B \ p \ \vec{y}$.

$$\forall \vec{x}. H_1 \supset \dots \supset H_n \supset p \ \vec{t}$$

We proceed as in the construction for induction to get the sequent

$$\vec{x} : H_1, \dots, H_n \longrightarrow p \ \vec{t}.$$

We then apply co-induction with the invariant S as follows.

$$S = \lambda \vec{y}. \exists \vec{x}. \vec{y} = \vec{t} \wedge H_1 \wedge \dots \wedge H_n$$

The CTR rule applied to the earlier sequent requires us to show $\vec{x} : H_1, \dots, H_n \longrightarrow S \vec{t}$ which is trivial and $\vec{y} : S \vec{y} \longrightarrow B S \vec{y}$ which contains the real content of the co-inductive proof. A derivation of this later sequent can be constructed as follows.

$$\frac{\frac{\frac{\vec{x} : H_1, \dots, H_n \longrightarrow B S \vec{t}}{\vec{y}, \vec{x} : \vec{y} = \vec{t}, H_1, \dots, H_n \longrightarrow B S \vec{y}} \triangleright \mathcal{L}_{CSNAS}}{\vec{y}, \vec{x} : \vec{y} = \vec{t} \wedge H_1 \wedge \dots \wedge H_n \longrightarrow B S \vec{y}} \wedge \mathcal{L}}{\vec{y} : \exists \vec{x}. \vec{y} = \vec{t} \wedge H_1 \wedge \dots \wedge H_n \longrightarrow B S \vec{y}} \exists \mathcal{L}}$$

The derivation for the upper-most sequent here can be constructed based on the argument carried out in the the annotation based treatment. Within that argument, when the goal $(p \vec{t})^\#$ is unfolded, the recursive calls will be annotated with $+$ and will be provable using the co-inductive hypothesis. This is what is given in the formal derivation by the goal $B S \vec{t}$.

This annotation based treatment of co-induction can be extended in ways similar to the inductive treatment. For example, we can allow co-induction within a context of other hypotheses, or we can allow the goal to be unfolded multiple times before applying the co-inductive hypotheses. The soundness arguments for these extensions are similar to the inductive case.

A Two-level Logic Approach to Reasoning

One approach to reasoning about object systems is to encode their descriptions directly into definitions in \mathcal{G} and to then use the inference rules of \mathcal{G} with these definitions. In this chapter we explore an alternative approach. In particular, we show how the meta-logic \mathcal{G} can be used to encode the specification logic hH^2 and to then reason about hH^2 specifications through this encoding. This is the two-level logic approach to reasoning that was enunciated by McDowell and Miller earlier in the context of the meta-logic $FO\lambda^{\Delta\mathbb{N}}$ [MM02].

An important part of assessing the value of the two-level logic approach to reasoning is understanding both its benefits and its costs. One benefit is that the specification logic carves out a useful subset of the specifications that are possible in the meta-logic while at the same time possessing a complete proof search procedure which make it possible to execute the specifications. A second benefit is that by encoding an entire specification logic in the meta-logic, we can formalize properties of the specification logic and make them available during reasoning. An auxiliary observation in this context is that because of the way the specification logic can be used to encode object systems, the properties of this logic that are used in meta-logic reasoning often turn out to be based on intuitions about the properties of the object systems themselves. From a cost perspective, one issue with the two-level logic approach to reasoning is that there is an additional overhead to reasoning about specifications through the encoded semantics of the specification logic rather than directly. Another cost to be considered is that because the specification logic is only a subset of the full range of specifications allowed by the meta-logic, this approach in some ways limits what we are able to say within a specification.

After all aspects are taken into account, we believe that the combination of the hH^2 specification logic and the meta-logic \mathcal{G} seems to provide a nice balance between the benefits and costs of the two-level logic approach to reasoning. The specification logic hH^2 elegantly encodes many systems of interest, and there are efficient implementations of this specification logic. Moreover, as we saw in Section 2.3, the properties of hH^2 provide useful results during reasoning. Finally, as we shall see in this chapter, the encoding of hH^2 into \mathcal{G} is lightweight and therefore imposes little overhead on the reasoning process.

The rest of this chapter is laid out as follows. Section 6.1 describes the encoding of hH^2 into \mathcal{G} . Section 6.2 formalizes some properties of hH^2 as theorems in \mathcal{G} ; these theorems can then be used as lemmas to simplify subsequent reasoning. Section 6.3 illustrates our specific realization of the two-level logic approach to reasoning and demonstrates its power by using it to formalize the informal proof that we have presented in Chapter 1 of the fact that types are preserved by evaluation in the simply-typed λ -calculus. Finally, Section 6.5 discusses the issue of adequacy relative to the two-level logic approach to reasoning.

6.1 Encoding the Specification Logic

There are two components to our encoding of the specification logic hH^2 into the meta-logic \mathcal{G} . First, we encode the syntax by defining a mapping ψ from specification logic types and terms to meta-logic types and terms. Since both logics are constructed from Church's simple theory of types and hence contain subsets of expressions that are isomorphic, this encoding can be very shallow. Second, we encode the semantics of hH^2 (*i.e.*, the provability relation) via the definition of a suitably chosen atomic judgment in \mathcal{G} . This encoding is lightweight which makes later reasoning fairly transparent. To aid in that reasoning we observe some formulas that can be proved in \mathcal{G} involving the judgment that encodes specification logic provability. These theorems of \mathcal{G} can be used as lemmas to shorten other proofs that we would want to construct in \mathcal{G} .

6.1.1 Encoding the Syntax of the Specification Logic

The types of our specification logic are mapped to isomorphic types in the meta-logic. We define the mapping ψ on types as follows.

$$\psi(\tau) = \tau \quad \text{if } \tau \text{ is a base type} \qquad \psi(\tau_1 \rightarrow \tau_2) = \psi(\tau_1) \rightarrow \psi(\tau_2)$$

For each specification type, we assume a bijective mapping between eigenvariables of that type (in the specification logic) and nominal constants of that type (in the meta-logic). We denote this mapping using subscripts: the eigenvariable h maps to the nominal constant a_h and the nominal constant a maps to the eigenvariable h_a . Using this, we define the encoding of specification terms as follows.

$$\begin{aligned} \psi(c) = c \quad \text{if } c \text{ is a constant} & \qquad \psi(h) = a_h \quad \text{if } h \text{ is an eigenvariable} \\ \psi(x) = x \quad \text{if } x \text{ is a variable} & \qquad \psi(\lambda x.t) = \lambda x.\psi(t) \qquad \psi(t_1 \ t_2) = \psi(t_1) \ \psi(t_2) \end{aligned}$$

Now for clarity and correctness of the encoding, we make two adjustments to this mapping. First, the specification logic type o for formulas is mapped to a distinguished type frm to avoid conflicting with the type o for meta-logic formulas. Second, we introduce a distinguished type atm for atomic specification logic formulas and a constructor $\langle \cdot \rangle : atm \rightarrow frm$ to inject such atoms into formulas. We then modify the type of the specification logic \supset connective to $atm \rightarrow frm \rightarrow frm$ to enforce the restriction that the left-hand side of an implication is atomic.

Note that we map specification logic constants to constants of the same name in the meta-logic. This means, for example, that the meta-logic will have two constants called \wedge . One will be the logical connective of \mathcal{G} with type $o \rightarrow o \rightarrow o$, and the other will be a term constructor for representations of specification logic formulas with type $frm \rightarrow frm \rightarrow frm$. We will always be able to distinguish between such constants based on the context in which they are used.

Our encoding is clearly bijective. Furthermore, typing judgments are preserved by the bijection in the following sense. Let \mathcal{K} denote the set of meta-logic constants which represent

the constants of the specification logic, then $\Sigma \vdash t : \tau$ is a valid specification logic typing if and only if $\psi(\Sigma), \mathcal{K} \vdash \psi(t) : \psi(\tau)$ is a valid meta-logic typing where $\psi(\Sigma) = \{\psi(h) \mid h \in \Sigma\}$. Since our mapping ψ is bijective we will use the mapping ψ^{-1} freely.

6.1.2 Encoding the Semantics of the Specification Logic

In the encoding of the semantics of our specification logic, we shall use two auxiliary notions. First, we introduce a type nt for natural numbers with the constructors $z : nt$ and $s : nt \rightarrow nt$ and the predicate $nat : nt \rightarrow o$ defined by

$$nat\ z \stackrel{\mu}{=} \top \qquad nat\ (s\ N) \stackrel{\mu}{=} nat\ N$$

As we see below, these numbers will be used to capture the idea of the height of a derivation in our encoding of the provability relation of the specification logic. Second, we introduce a type $atmlist$ with constructors $nil : atmlist$ and the infix $:: : atm \rightarrow atmlist \rightarrow atmlist$ and the predicate $member : atm \rightarrow atmlist \rightarrow o$ defined by

$$member\ A\ (A :: L) \stackrel{\mu}{=} \top \qquad member\ A\ (B :: L) \stackrel{\mu}{=} member\ B\ L$$

We shall use lists of this kind and the corresponding membership predicate to encode the addition to premise sets when trying to prove implicational formulas in hH^2 .

We encode hH^2 provability in \mathcal{G} through the predicate $seq : nt \rightarrow atmlist \rightarrow frm \rightarrow o$ that is defined by the clauses in Figure 6.1. This encoding of hH^2 provability derives from McDowell and Miller [MM02]. As described in Chapter 2, proofs in hH^2 contain sequents of the form $\Sigma : \Delta, \mathcal{L} \vdash G$ where Δ is a fixed set of closed D -formulas and \mathcal{L} is a varying set of atomic formulas. The eigenvariables in Σ are encoded as nominal constants in \mathcal{G} . The meta-logic predicate $prog : atm \rightarrow frm \rightarrow o$ is used to represent the D -formulas in Δ : the D formula $\forall \vec{x}. [G_1 \supset \cdots \supset G_n \supset A]$ is encoded as the clause $\forall \vec{x}. prog\ A\ (G_1 \wedge \cdots \wedge G_n) \triangleq \top$ and $\forall \vec{x}. A$ is encoded by the clause $\forall \vec{x}. prog\ A\ \top \triangleq \top$. We denote these $prog$ clauses by $\Psi(\Delta)$, and we note that such clauses do not contain any nominal constants since the formulas of Δ are closed. Finally, the hH^2 sequent is encoded as $seq_N\ \psi(\mathcal{L})\ \psi(G)$ where we define ψ

$$\begin{aligned}
& seq_{(s \ N)} L \top \stackrel{\mu}{=} \top \\
& seq_{(s \ N)} L (B \vee C) \stackrel{\mu}{=} seq_N L B \\
& seq_{(s \ N)} L (B \vee C) \stackrel{\mu}{=} seq_N L C \\
& seq_{(s \ N)} L (B \wedge C) \stackrel{\mu}{=} seq_N L B \wedge seq_N L C \\
& seq_{(s \ N)} L (A \supset B) \stackrel{\mu}{=} seq_N (A :: L) B \\
& seq_{(s \ N)} L (\forall B) \stackrel{\mu}{=} \nabla x. seq_N L (B \ x) \\
& seq_{(s \ N)} L (\exists B) \stackrel{\mu}{=} \exists x. seq_N L (B \ x) \\
& seq_{(s \ N)} L \langle A \rangle \stackrel{\mu}{=} member \ A \ L \\
& seq_{(s \ N)} L \langle A \rangle \stackrel{\mu}{=} \exists b. prog \ A \ b \wedge seq_N L \ b
\end{aligned}$$

Figure 6.1: Second-order hereditary Harrop logic in \mathcal{G}

on lists of atomic formulas as $\psi(A_n, \dots, A_1) = A_1 :: \dots :: A_n :: nil$. The argument N , written as a subscript, roughly corresponds to the height of the proof tree and is used in inductive arguments. To simplify notation, we write $L \Vdash_n G$ for $seq_n L G$ and $L \Vdash G$ for $\exists n. nat \ n \wedge seq_n L G$. When L is nil we write simply $\Vdash_n G$ or $\Vdash G$.

Proofs of universally quantified G formulas in hH^2 are generic in nature. A natural encoding of this (object-level) quantifier in the definition of seq uses a (meta-level) ∇ -quantifier. In the case of proving an implication, the atomic assumption is maintained in a list (the second argument of seq). The last clause for seq implements backchaining over a fixed hH^2 specification (stored as $prog$ atomic formulas). The matching of atomic judgments to heads of clauses is handled by the treatment of definitions in the logic \mathcal{G} , thus the last rule for seq simply performs this matching and makes a recursive call on the corresponding clause body.

Note that for each specification type τ we have the constants $\forall_\tau : (\tau \rightarrow frm) \rightarrow frm$ and $\exists_\tau : (\tau \rightarrow frm) \rightarrow frm$, thus we should have seq clauses for each of these. However,

here and going forward, we present only general rules for \forall and \exists , knowing that the actual rules are easily derived from these.

With this kind of an encoding, we can now formulate and prove in \mathcal{G} statements about what is or is not provable in hH^2 . In constructing such proofs, we shall sometimes need induction over the height of derivations. Such arguments can be realized via induction on the predicate $nat\ n$ in a formula of the form $\exists n. nat\ n \wedge seq_n\ L\ G$ occurring on the left of a sequent. We may sometimes also want to use strong induction in our arguments. Towards this end, we introduce the auxiliary predicate $It : nt \rightarrow nt \rightarrow o$ defined as follows.

$$\begin{aligned} It\ z\ (s\ N) &\stackrel{\mu}{=} \top \\ It\ (s\ M)\ (s\ N) &\stackrel{\mu}{=} It\ M\ N \end{aligned}$$

Now, a formula such as $\forall n.(nat\ n) \supset P$ can be proven using strong induction by proving $\forall n, m.(nat\ n \wedge It\ n\ m \wedge nat\ m) \supset P$ and using induction on $nat\ m$. Section 6.3 contains an example that uses this approach. Finally, the $def\mathcal{L}$ rule can be used to realize case analysis based reasoning in the derivation of an atomic goal. Using this rule leading eventually to a consideration of the different ways in which an atomic judgment may have been inferred in the specification logic.

In the rest of this chapter, we shall conduct all of our reasoning by constructing derivations in \mathcal{G} , with the exception of adequacy arguments where we will need to reason over \mathcal{G} derivations. Thus, when we say that “a formula F is provable” or that “a formula F is provable in \mathcal{G} ”, we shall mean that the sequent $\longrightarrow F$ is provable in \mathcal{G} . Moreover, when we talk about the “proof of a formula F ” we shall mean the derivation in \mathcal{G} of the sequent $\longrightarrow F$. When we say that such proofs are constructed “by induction” we shall mean that we use the \mathcal{IL} rule of \mathcal{G} with an induction invariant derived from the entire sequent being considered. We shall also talk about proving a formula by induction on one of its hypotheses (*i.e.*, one of its subformulas to the left of a \supset) by which we mean following the constructions for induction described in Chapter 5. The construction of the derivations in \mathcal{G} is often straightforward, with only a few sequents which may be interesting, and so we

shall frequently skip directly to such sequents. Finally, we shall often use running text to describe the construction of a derivation in \mathcal{G} ; this is possible since the rules of \mathcal{G} often mimic traditional mathematical reasoning, but it must be remembered that the proof is still being carried out within \mathcal{G} .

Several of the results that we present below concern the provability of formulas in \mathcal{G} . While our proofs of these results here involve arguing about derivations in \mathcal{G} , it is important to note that these arguments sketch a scheme for actually carrying out the proof *within* a system such as Abella. Thus, the justification for using such formulas in subsequent arguments is completely formalized through actual mechanical proofs and the lemma mechanism of Abella; in particular, the resulting style of (mechanized) argument does not rely on the informal proofs we present to justify the approach.

6.1.3 Some Provable Properties of the Specification Logic

It is often convenient to reason directly with formulas of the form $L \Vdash G$ rather than expanding them into $\exists n.nat\ n \wedge seq_n\ L\ G$. In this section, we show that certain schematic formulas corresponding to \Vdash judgments are provable in \mathcal{G} . Using these as lemmas allows us to encode certain direct forms of reasoning about \Vdash in \mathcal{G} proofs. The particular formulas that we show to be provable in \mathcal{G} closely mirror the clauses which define the *seq* predicate.

Lemma 6.1.1. *The following formulas are provable in \mathcal{G} .*

1. $\forall \ell. (\ell \Vdash \top)$
2. $\forall \ell, g_1, g_2. (\ell \Vdash g_1) \supset (\ell \Vdash g_1 \vee g_2)$
3. $\forall \ell, g_1, g_2. (\ell \Vdash g_2) \supset (\ell \Vdash g_1 \vee g_2)$
4. $\forall \ell, g_1, g_2. (\ell \Vdash g_1) \wedge (\ell \Vdash g_2) \supset (\ell \Vdash g_1 \wedge g_2)$
5. $\forall \ell, a, g. (a :: \ell \Vdash g) \supset (\ell \Vdash a \supset g)$
6. $\forall \ell, g. (\nabla x. (\ell \Vdash (g\ x))) \supset (\ell \Vdash \forall g)$

$$7. \forall \ell, g, t. (\ell \Vdash (g \ t)) \supset (\ell \Vdash \exists g)$$

Proof. It is easy to see that the formulas 1, 2, 3, 5, and 7 are provable in \mathcal{G} by unfolding (*i.e.*, using $\text{def}\mathcal{R}$ on) the goal formulas.

In the straightforward construction of a proof of formula 4, we shall need to construct a proof of the following sequent.

$$\text{nat } n, \text{seq}_n \ell \ g_1, \text{nat } m, \text{seq}_m \ell \ g_2 \longrightarrow \exists p. \text{nat } p \wedge \text{seq}_p \ell \ (g_1 \wedge g_2).$$

To prove this we must reconcile the measures n and m . Towards this end, we might first show that the following formula that relates n and m is provable in \mathcal{G} :

$$\forall m, n. (\text{nat } m) \wedge (\text{nat } n) \supset (\text{It } m \ n) \vee (m = n) \vee (\text{It } n \ m).$$

This can be proved by induction on one of the nat hypotheses. Then we can also prove the following formula which allows us to increase the measure of a derivation:

$$\forall m, n, \ell, g. (\text{It } m \ n) \wedge (\ell \Vdash_m g) \supset (\ell \Vdash_n g).$$

This is proved by induction on $\text{It } m \ n$. Using these two lemmas the rest of the proof is straightforward.

In constructing a proof of Formula 6 we will find it necessary to construct a proof of the sequent

$$\exists n. \text{nat } n \wedge \text{seq}_n \ell \ (g \ a) \longrightarrow \exists m. \text{nat } m \wedge \text{seq}_m \ell \ (\forall g).$$

where a is a nominal constant. Now when we apply $\exists\mathcal{L}$, we have the sequent

$$\text{nat } (n' \ a) \wedge \text{seq}_{(n' \ a)} \ell \ (g \ a) \longrightarrow \exists m. \text{nat } m \wedge \text{seq}_m \ell \ (\forall g).$$

The raising of n' over a here prevents this proof from going through immediately, thus we need the following lemma.

$$\forall n. (\nabla x. \text{nat } (n \ x)) \supset \exists p. n = \lambda y. p$$

This is proved by induction on nat . Once we apply this lemma we have $n' = \lambda y. p$ for some p and rest of the proof is straightforward. \square

6.2 Formalizing Meta-Theoretic Properties of the Specification Logic

In Section 2.2 we observed certain meta-theoretic properties of hH^2 which are useful in reasoning about hH^2 specifications. Since we have encoded the entire specification logic into \mathcal{G} , we can formalize such properties of the specification logic within \mathcal{G} . In particular, we can consider particular formulas in \mathcal{G} that encode these properties and then we can show that these formulas are provable in \mathcal{G} . Doing this will allow us to later bring these properties to bear on particular reasoning tasks that are carried out using \mathcal{G} . The particular properties of hH^2 that we consider in this way in this section are monotonicity, instantiation, and cut admissibility. With one exception, the proofs of these properties never use a *prog* formula except in the initial rule and thus the proofs are independent of any particular specification encoded in *prog*. The one exception is specifically noted, and even here the proof is independent of the specification.

Monotonicity The statement of monotonicity for hH^2 , expressed as a formula of \mathcal{G} , is

$$\forall n, \ell_1, \ell_2, g. (\ell_1 \Vdash_n g) \wedge (\forall e. \text{member } e \ell_1 \supset \text{member } e \ell_2) \supset (\ell_2 \Vdash_n g).$$

The proof is by straightforward induction on the hypothesis *nat* n in $\ell_1 \Vdash_n g$.

Instantiation The instantiation property recovers the notion of universal quantification from our representation of the specification logic \forall using ∇ . This property is expressed in \mathcal{G} through the formula

$$\forall \ell, g. (\nabla x. (\ell x) \Vdash_n (g x)) \supset \forall t. (\ell t) \Vdash_n (g t).$$

Stated another way, although ∇ quantification cannot be replaced by \forall quantification in general, it can be replaced in this way when dealing with specification judgments. The proof of this formula is by induction on the hypothesis *nat* n in $(\ell x) \Vdash_n (g x)$, and the following two auxiliary results are useful in constructing this proof.

$$\forall \ell, a. (\nabla x. \text{member } (a x) (\ell x)) \supset \forall t. \text{member } (a t) (\ell t)$$

$$\forall a, b. (\nabla x. \text{prog} (a x) (b x)) \supset \forall t. \text{prog} (a t) (b t)$$

The first is proved by induction on the *member* hypothesis. The second depends on the particular specification encoded in *prog*, but the core of the proof is always applying *defL* to *prog* (a x) (b x) followed by *defR* on *prog* (a t) (b t). This will succeed for any specification since *prog* only performs pattern matching and contains no “logic.”

Cut admissibility The cut admissibility property of hH^2 is expressed in \mathcal{G} through the formula

$$\forall \ell, a, g. (\ell \Vdash \langle a \rangle) \wedge (a :: \ell \Vdash g) \supset (\ell \Vdash g).$$

The proof is by induction on the *nat* n assumption in $\exists n. \text{nat } n \wedge \text{seq}_n (a :: \ell) g$. If $n = z$ then the *seq* judgment is impossible, thus we know $n = s m$ for some m . The proof proceeds by case analysis on the *seq* judgment.

1. One case is when $g = \langle a' \rangle$ and *member* a' (a :: ℓ). Applying *defL* to this *member* hypothesis results in two additional cases: either $a = a'$ so that $\ell \Vdash \langle a \rangle$ holds by assumption, or we know *member* $a' \ell$ and thus $\ell \Vdash \langle a' \rangle$ holds by applying *defR^p* and *init*.
2. Another case is when $g = a' \supset g'$ so that we have $a' :: a :: \ell \Vdash_m g'$. We then apply the monotonicity property once to get $a :: a' :: \ell \Vdash_m g'$ and another time to get $a' :: \ell \Vdash \langle a \rangle$. Then we can apply the inductive hypothesis to get $a' :: \ell \Vdash g'$ and therefore $\ell \Vdash a' \supset g'$.
3. The remaining cases follow directly from the inductive hypothesis and the results in Lemma 6.1.1.

6.3 An Example of the Two-level Logic Reasoning Approach

Within this framework of the two-level logic approach to reasoning, we come back to the example of evaluation and typing for the simply-typed λ -calculus. We use the hH^2 specification of these notions given in Section 2.3 which yields the *prog* clauses shown in Figure 6.2. We can now formalize the type preservation theorem completely in the meta-logic:

$$\begin{aligned}
& \text{prog } (\text{eval } (\text{abs } A \ R) \ (\text{abs } A \ R)) \top \triangleq \top \\
& \text{prog } (\text{eval } (\text{app } M \ N) \ V) \ (\langle \text{eval } M \ (\text{abs } A \ R) \rangle \wedge \langle \text{eval } (R \ N) \ V \rangle) \triangleq \top \\
& \text{prog } (\text{of } (\text{app } M \ N) \ B) \ (\langle \text{of } M \ (\text{arr } A \ B) \rangle \wedge \langle \text{of } N \ A \rangle) \triangleq \top \\
& \text{prog } (\text{of } (\text{abs } A \ R) \ (\text{arr } A \ B)) \ (\forall x. \text{of } x \ A \supset \langle \text{of } (R \ x) \ B \rangle) \triangleq \top
\end{aligned}$$

Figure 6.2: *prog* clauses for simply-typed λ -calculus

Theorem 6.3.1. *The following formula is derivable in \mathcal{G} .*

$$\forall e, t, v. (\Vdash \langle \text{eval } e \ v \rangle) \wedge (\Vdash \langle \text{of } e \ t \rangle) \supset (\Vdash \langle \text{of } v \ t \rangle)$$

Proof. The informal argument for the proof of type preservation presented in Section 2.3 is based on strong induction over the height of hH^2 derivations. We will now show how we can mimic that same style of induction in \mathcal{G} . We first generalize the formula we want to prove to the following.

$$\forall e, t, v, i, j. (\text{nat } j) \wedge (\text{It } i \ j) \wedge (\text{seq}_i \ \text{nil } \langle \text{eval } e \ v \rangle) \wedge (\Vdash \langle \text{of } e \ t \rangle) \supset (\Vdash \langle \text{of } v \ t \rangle)$$

If we prove this generalization, then we can use the *cut* rule to bring it in as a hypothesis in a proof of the original formula. The resulting sequent will then be easily provable. To prove the generalization, we use induction on *nat j*. In the case where $j = z$, the proof is trivial since *It i z* is unsatisfiable. In the other case we have $j = s \ j'$ and we know the result holds for any i such that *It i j'*. In this way, we can completely handle the strong induction within our logic.

The rest of proof of the generalization closely follows the informal argument with only the following points worthy of note.

Case analysis on specification judgments in the informal argument is realized in the construction of a derivation in \mathcal{G} by using *def \mathcal{L}* twice. Specifically, if we want to do case analysis on a derivation such as $\text{seq}_i \ \text{nil } \langle \text{eval } e \ v \rangle$ then we apply *def \mathcal{L}* which results in two

cases. The first is that $member (eval e v) nil$ holds which is impossible. The second is that $\exists b.prog (eval e v) b \wedge seq_{i'} nil b$ holds for some i' such that $i = s i'$. Then we can apply $def\mathcal{L}$ on $prog (eval e v) b$ which gives us the two cases corresponding to the clauses for forming $eval$ judgments.

The instantiation and cut admissibility properties of our specification logic which are used the informal argument are now formal lemmas which are applied in this proof. Thus the entire proof is formally constructed within \mathcal{G} while still using meta-theoretic properties of hH^2 . \square

6.4 Architecture of a Two-level Logic Based Theorem Prover

The architecture of the Abella theorem prover for \mathcal{G} presented in Chapter 5 can be naturally extended to support the two-level logic approach to reasoning that is the topic of discussion in this current chapter. In fact, the Abella system already incorporates such an extension [Gac09]. In this section we briefly describe the architectural changes which facilitate this support. Most of these changes can be motivated from the type preservation example shown in the previous section which we will refer to as simply “the example.”

The first step in the two-level logic approach to reasoning is encoding a specification into the proper $prog$ statement. Abella facilitates this by reading specifications written in the subset of λ Prolog which corresponds to hH^2 . In this way, the specifications used by Abella are directly executable by λ Prolog implementations such as Teyjus without the potentially error-prone need to translate between different input languages.

To reduce syntactic overhead associated with the two-level logic approach to reasoning, Abella has specialized syntax for representing judgments of the form $\ell \Vdash g$. Direct reasoning on these judgments is enabled by incorporating the derived rules of inference from Section 6.1.3. Case analysis on judgments of the form $\ell \Vdash g$ in Abella corresponds to applying $def\mathcal{L}$ to underlying the seq judgment followed by applying $def\mathcal{L}$ to the resulting $prog$ judgment. Trivial cases such as $member E nil$ are handled automatically. Thus much of the overhead which is shown in the example is hidden when working with Abella.

The monotonicity, instantiation, and cut-admissibility properties of the specification logic (Section 6.2) are incorporated into Abella in the form of tactics. Moreover, the monotonicity property is incorporated into some other existing tactics since it seems to be used most often. For example, when determining if $\ell \Vdash g$ implies $\ell' \Vdash g$ the system checks if ℓ is an obvious subset of ℓ' . Such checks arise often, for example, when applying a lemma to hypotheses.

Abella simulates strong induction on hH^2 derivations using the technique shown in the example. In general, the induction tactic applied to a judgment of the form $\ell \Vdash g$ is treated as strong induction on the underlying measure. This is approximated using the annotation based treatment of induction from Section 5.2 applied directly to specification judgments. This has the benefit of removing much of the tedious reasoning about natural numbers which would otherwise clutter a proof. As an example of this annotation based treatment, suppose we want to prove a formula of the form

$$\forall \vec{x}. (\ell \Vdash g) \supset F.$$

Then the induction scheme creates the following inductive hypothesis and goal, respectively:

$$\forall \vec{x}. (\ell \Vdash g)^* \supset F \qquad \forall \vec{x}. (\ell \Vdash g)^{\textcircled{a}} \supset F.$$

Eventual case analysis on $(\ell \Vdash g)^{\textcircled{a}}$ results in recursive judgments of the form $(\ell' \Vdash g')^*$ which are subject to the inductive hypothesis. The monotonicity and instantiation properties of the specification logic preserve the height of hH^2 derivations, and thus tactics which implement them preserve induction annotations as well (since induction is being carried out on the underlying height measure). Finally, suppose we want to deal with mutual induction on specification judgments. For example, suppose we have a goal of the form

$$(\forall \vec{x}_1. (\ell_1 \Vdash g_1) \supset F_1) \wedge (\forall \vec{x}_2. (\ell_2 \Vdash g_2) \supset F_2).$$

We can perform induction on both of the specification judgments simultaneously by instead considering the following goal

$$\forall n. \text{nat } n \supset (\forall \vec{x}_1. (\ell_1 \Vdash_n g_1) \supset F_1) \wedge (\forall \vec{x}_2. (\ell_2 \Vdash_n g_2) \supset F_2),$$

and performing induction on *nat* n . Once this new goal is proven, the original is an easy consequence. We extend the annotation based treatment of induction to treat this kind of mutual induction directly. Specifically, it creates the following two inductive hypotheses

$$(\forall \vec{x}_1. (\ell_1 \Vdash g_1)^* \supset F_1) \qquad (\forall \vec{x}_2. (\ell_2 \Vdash g_2)^* \supset F_2),$$

and the goal becomes

$$(\forall \vec{x}_1. (\ell_1 \Vdash g_1)^{\textcircled{a}} \supset F_1) \wedge (\forall \vec{x}_2. (\ell_2 \Vdash g_2)^{\textcircled{a}} \supset F_2).$$

The proof then proceeds as normal. When case analysis is performed on a judgment with a \textcircled{a} annotation, the recursive calls will have the $*$ annotation and thus be candidates for either of the inductive hypotheses.

6.5 Adequacy for the Two-level Logic Approach to Reasoning

Adequacy within the framework based on the two-level logic approach to reasoning has three components:

1. Our encoding of the object system into hH^2 must be adequate.
2. Our encoding of hH^2 into \mathcal{G} must be adequate.
3. We must show that information about object system properties can be extract from theorems in \mathcal{G} via the two encodings.

The first component is particular to the object system of interest. For example, adequacy for the hH^2 encoding of evaluation and typing for the simply-typed λ -calculus was shown in Section 2.4. In the current section we are primarily concerned with latter two components which deal with adequacy relative to \mathcal{G} . The second component is a general result about hH^2 and its encoding in the predicate *seq* (we shall often call this simply “the adequacy of *seq*”). The proof of this result is carried out in the next subsection, and it never needs to be changed since hH^2 and *seq* are fixed. The last component of adequacy is particular

to the theorems of interest, and in Section 6.5.2 we show this adequacy for the example of type preservation for the simply-typed λ -calculus.

There is some difficulty in establishing adequacy relative to \mathcal{G} . When we represent objects in \mathcal{G} we usually denote bound variables using λ -terms and free variables using nominal constants. Then, when we quantify over such objects, we are usually interested only in objects whose free variables are restricted to a particular set (*e.g.*, we may care only about closed objects). The \forall and \exists quantifiers of \mathcal{G} , however, allow nominal constants to appear freely in the terms that instantiate them. There are two ways to address this mismatch (without modifying the logic \mathcal{G}). The first is to define an explicit typing of objects (*e.g.*, through a predicate *typeof* $L T A$ where L is a context of nominal constants), and to attach this typing judgment to all quantified variables. This is a very heavy approach and requires explicitly maintaining a context of which nominal constants are allowed to appear in objects. An alternative approach, and the one we use to establish the adequacy of *seq* in the next subsection, is to restrict the use of nominal constants in such a way that adequacy can still be established. How exactly this is done depends on the particular system of interest and how nominal constants are treated by it. In the case of *seq* we know that nominal constants can always be instantiated, thus the only restriction we need is that nominal constants are allowed only at inhabited types.

6.5.1 Adequacy of Encoding of the Specification Logic

We now show that our encoding of the specification logic hH^2 in the definition of *seq* and *prog* is adequate. The critical aspect of this result is showing that theoremhood in the two systems is preserved under an appropriate mapping.

Theorem 6.5.1. *Let Δ be a list of closed D -formulas, \mathcal{L} a list of atoms, G a G -formula, and Σ a set of eigenvariables containing at least the free variables of Δ , \mathcal{L} , and \mathcal{G} . Suppose that all non-logical specification logic constants and types are represented by equivalent constants and types in \mathcal{G} . Suppose also that specification logic \forall -quantification (eigenvariables) and meta-logic ∇ -quantification (nominal constants) are allowed only at inhabited types. Then*

$\Sigma : \Delta, \mathcal{L} \vdash G$ has a derivation in hH^2 if and only if $\psi(\mathcal{L}) \Vdash \psi(G)$ is provable in \mathcal{G} with the clauses for *nat*, *member*, and *seq* as stated before and the clauses for *prog* as given by $\Psi(\Delta)$.

Proof. Note that in this proof we will desugar the representation of quantification and substitution in the specification logic.

Forward direction. Given a derivation of $\Sigma : \Delta, \mathcal{L} \vdash G$ in hH^2 , we will construct a proof of $\psi(\mathcal{L}) \Vdash \psi(G)$ in \mathcal{G} . The construction uses structural induction on the hH^2 derivation and proceeds by cases on the last rule used in the derivation.

1. Suppose the derivation ends with OR_1 :

$$\frac{\Sigma : \Delta, \mathcal{L} \vdash G_1}{\Sigma : \Delta, \mathcal{L} \vdash G_1 \vee G_2} \text{OR}_1$$

By the inductive hypothesis we know $\psi(\mathcal{L}) \Vdash \psi(G_1)$ is provable in \mathcal{G} . Then we know $\psi(\mathcal{L}) \Vdash \psi(G_1 \vee G_2)$ using the appropriate formula from Lemma 6.1.1.

2. Suppose the derivation ends with TRUE , OR_2 , AND , or AUGMENT : these cases are similar to the previous one.

3. Suppose the derivation ends with GENERIC :

$$\frac{\Sigma, c : \Delta, \mathcal{L} \vdash G' c}{\Sigma : \Delta, \mathcal{L} \vdash \forall G'} \text{GENERIC}$$

By the inductive hypothesis we know $\psi(\mathcal{L}) \Vdash \psi(G' c)$ is provable in \mathcal{G} . We also know $\psi(G' c) = \psi(G') a_c$ where a_c is a nominal constant not in $\psi(\Sigma)$ (and therefore not occurring in $\psi(\mathcal{L})$ or $\psi(G')$). Thus we know there is a proof of $\nabla x.(\psi(\mathcal{L}) \Vdash (\psi(G') x))$. Using the appropriate formula from Lemma 6.1.1, there must be a proof of $\psi(\mathcal{L}) \Vdash \forall \psi(G')$.

4. Suppose the derivation ends with INSTANCE :

$$\frac{\Sigma : \Delta, \mathcal{L} \vdash G' t}{\Sigma : \Delta, \mathcal{L} \vdash \exists_r G'} \text{INSTANCE}$$

By the inductive hypothesis we know $\psi(\mathcal{L}) \Vdash \psi(G' t)$ is provable in \mathcal{G} . We also know $\psi(G' t) = \psi(G') \psi(t)$. Using the appropriate formula from Lemma 6.1.1, there must be a proof of $\psi(\mathcal{L}) \Vdash \exists \psi(G')$.

5. Suppose the derivation ends with BACKCHAIN:

$$\frac{\Sigma : \Delta, \mathcal{L} \vdash G_1 \vec{t} \quad \dots \quad \Sigma : \Delta, \mathcal{L} \vdash G_m \vec{t}}{\Sigma : \Delta, \mathcal{L} \vdash A} \text{BACKCHAIN}$$

where $\forall \vec{x}. (G_1 \vec{x} \supset \dots \supset G_m \vec{x} \supset A' \vec{x}) \in \Delta, \mathcal{L}$ and $A' \vec{t} = A$. We distinguish two cases based on whether the formula is in Δ or in \mathcal{L} .

(a) Suppose $\forall \vec{x}. (G_1 \vec{x} \supset \dots \supset G_m \vec{x} \supset A' \vec{x}) \in \Delta$. Then we must have the following clause.

$$\forall \vec{x}. \text{prog } (\psi(A') \vec{x}) (\psi(G_1) \vec{x} \wedge \dots \wedge \psi(G_m) \vec{x}) \triangleq \top$$

By the inductive hypothesis we have a proof of $\psi(\mathcal{L}) \Vdash \psi(G_i \vec{t})$ for each $i \in \{1, \dots, m\}$. By repeatedly using the appropriate formula from Lemma 6.1.1 we can construct a proof of $\psi(\mathcal{L}) \Vdash (\psi(G_1 \vec{t}) \wedge \dots \wedge \psi(G_m \vec{t}))$, which we can write as $\psi(\mathcal{L}) \Vdash (\psi(G_1) \overrightarrow{\psi(t)}) \wedge \dots \wedge \psi(G_m) \overrightarrow{\psi(t)}$. Finally we know $\psi(A) = \psi(A' \vec{t}) = \psi(A') \overrightarrow{\psi(t)}$. Thus we know $\exists b. \text{prog } \psi(A) b \wedge (\psi(\mathcal{L}) \Vdash b)$ and we can construct a proof of $\psi(\mathcal{L}) \Vdash \langle \psi(A) \rangle$.

(b) Suppose $\forall \vec{x}. (G_1 \vec{x} \supset \dots \supset G_m \vec{x} \supset A' \vec{x}) \in \mathcal{L}$. Since \mathcal{L} contains only atoms we must have $A = A'$ and thus $A \in \mathcal{L}$. Then *member* $\psi(A) \psi(\mathcal{L})$ is provable and thus so is $\psi(\mathcal{L}) \Vdash \langle \psi(A) \rangle$.

Backward direction. It suffices to show if $\text{nat } (s \ n)$ and $\text{seq}_{(s \ n)} \psi(\mathcal{L}) \psi(G)$ have cut-free proofs in \mathcal{G} , then we can construct a derivation of $\Sigma : \Delta, \mathcal{L} \vdash G$ in hH^2 for any Σ which contains at least the eigenvariables of \mathcal{L} and G . The proof is by induction on the natural number denoted by $(s \ n)$ (which we know is a natural number since $\text{nat } (s \ n)$ has a proof). This proof will always end with *defR^p* (or can be seen to) and we will consider cases based on the definitional clause used in this rule.

1. The cases for the first five clauses of seq are all similar and thus we will consider just one instance. Suppose the cut-free proof ends with,

$$\frac{\longrightarrow seq_n \psi(\mathcal{L}) \psi(G_1)}{\longrightarrow seq_{(s\ n)} \psi(\mathcal{L}) (\psi(G_1) \vee \psi(G_2))} def\mathcal{R}^p$$

By the inductive hypothesis we know there is a derivation of $\Sigma : \Delta, \mathcal{L} \vdash G_1$ and we can construct the following.

$$\frac{\Sigma : \Delta, \mathcal{L} \vdash G_1}{\Sigma : \Delta, \mathcal{L} \vdash G_1 \vee G_2} OR_1$$

2. Suppose the cut-free proof ends with,

$$\frac{\frac{\longrightarrow seq_n \psi(\mathcal{L}) (\psi(G') a)}{\longrightarrow \nabla x.seq_n \psi(\mathcal{L}) (\psi(G') x)} \nabla\mathcal{R}}{\longrightarrow seq_{(s\ n)} \psi(\mathcal{L}) (\forall\psi(G'))} def\mathcal{R}^p$$

Since $\psi(G') a = \psi(G' h_a)$ we know from the inductive hypothesis that there is a derivation of $\Sigma, h_a : \Delta, \mathcal{L} \vdash G' h_a$. Thus we can construct the following.

$$\frac{\Sigma, h_a : \Delta, \mathcal{L} \vdash G' h_a}{\Sigma : \Delta, \mathcal{L} \vdash \forall G'} GENERIC$$

3. Suppose the cut-free proof ends with,

$$\frac{\mathcal{C}, \mathcal{K} \vdash t : \tau \longrightarrow seq_n \psi(\mathcal{L}) (\psi(G') t)}{\frac{\longrightarrow \exists_\tau x.seq_n \psi(\mathcal{L}) (\psi(G') x)}{\longrightarrow seq_{(s\ n)} \psi(\mathcal{L}) (\exists_\tau \psi(G'))} \exists\mathcal{R}} \exists\mathcal{R}$$

Now t may contain any nominal constants and therefore $t' = \psi^{-1}(t)$ may contain eigenvariables not in Σ . Thus when we apply the inductive hypothesis to $seq_n \psi(\mathcal{L}) \psi(G' t')$ we get a derivation of $\Sigma' : \Delta, \mathcal{L} \vdash G' t'$ where Σ' may contain additional eigenvariables. To reconcile this, we use the restriction that eigenvariables are allowed only at inhabited types. For each eigenvariable in t' and not in Σ , we select an inhabitant of the corresponding type and substitute it for the eigenvariable using the instantiation property of hH^2 . Since these eigenvariables do not occur in Σ , they also do not occur

in \mathcal{L} or G and therefore the instantiations affect only t' . Thus the result of all these instantiations is a derivation of $\Sigma : \Delta, \mathcal{L} \vdash G' t''$ for some t'' . Then we can construct the following.

$$\frac{\Sigma : \Delta, \mathcal{L} \vdash G' t''}{\Sigma : \Delta, \mathcal{L} \vdash \exists G'} \text{INSTANCE}$$

4. Suppose the cut-free proof ends with,

$$\frac{\longrightarrow \text{member } \psi(A) \psi(\mathcal{L})}{\longrightarrow \text{seq}_{(s \ n)} \psi(\mathcal{L}) \langle \psi(A) \rangle} \text{def}\mathcal{R}^p$$

Then it must be that $A \in \mathcal{L}$, and so we can construct the following.

$$\overline{\Sigma : \Delta, \mathcal{L} \vdash A} \text{BACKCHAIN}$$

5. Suppose the cut-free proofs ends with,

$$\frac{\frac{\frac{\longrightarrow \text{prog } \psi(A) b}{\longrightarrow \text{prog } \psi(A) b \wedge \text{seq}_n \psi(\mathcal{L}) b} \wedge \mathcal{R}}{\longrightarrow \exists b. \text{prog } \psi(A) b \wedge \text{seq}_n \psi(\mathcal{L}) b} \exists \mathcal{R}}{\longrightarrow \text{seq}_{(s \ n)} \psi(\mathcal{L}) \langle \psi(A) \rangle} \text{def}\mathcal{R}^p$$

for some instantiation of b . Suppose also that $\text{prog } \psi(A) b$ holds by matching with some clause,

$$\forall \vec{x}. \text{prog } (\psi(A') \vec{x}) (\psi(G_1) \vec{x} \wedge \dots \wedge \psi(G_m) \vec{x}) \triangleq \top.$$

Then we know $\forall \vec{x}. (G_1 \vec{x} \supset \dots \supset G_m \vec{x} \supset A' \vec{x}) \in \Delta$. From matching with the *prog* clause we know there exists \vec{t} such that $\psi(A) = \psi(A') \vec{t}$, so let $\vec{s} = \psi^{-1}(\vec{t})$. Then b is $\psi(G_1 \vec{s}) \wedge \dots \wedge \psi(G_m \vec{s})$ and we have proofs of $\text{seq}_n \psi(\mathcal{L}) \psi(G_i \vec{s})$ for each $i \in \{1, \dots, m\}$. By the inductive hypothesis we have derivations of $\Sigma' : \Delta, \mathcal{L} \vdash G_i \vec{s}$ where Σ' contains the eigenvariables of $\mathcal{L}, G_1, \dots, G_m$, and \vec{s} . Note that as was the case for the *seq* rule governing the existential quantifier, Σ' may contain some eigenvariables from \vec{s} which do not occur in Σ . As with that case, we can use the restriction on specification logic eigenvariables to instantiate all such eigenvariables with inhabitants therefore yielding derivations $\Sigma : \Delta, \mathcal{L} \vdash G_i \vec{r}$ where \vec{r} is the result of the instantiations

on \vec{s} . Finally, we know $A = A' \vec{s}$ but we need to know $A = A' \vec{r}$. Note that A' contains no eigenvariables and the eigenvariables of A are a subset of Σ , thus the eigenvariables in \vec{s} but not in Σ play no role in the equality $A = A' \vec{s}$. Therefore instantiating those eigenvariables does not change the equality and we have $A = A' \vec{r}$. Thus we can construct the following.

$$\frac{\Sigma : \Delta, \mathcal{L} \vdash G_1[\vec{r}/\vec{x}] \quad \cdots \quad \Sigma : \Delta, \mathcal{L} \vdash G_m[\vec{r}/\vec{x}]}{\Sigma : \Delta, \mathcal{L} \vdash A} \text{BACKCHAIN} \quad \square$$

Note that this theorem restricts the definitions of the predicates *nat*, *member*, *seq*, and *prog*, but makes no explicit reference to other predicates. Indeed, the definitions of other predicates have no affect on the adequacy of the encoding of the specification logic. Additionally, \mathcal{G} may make use of additional constants and types which are unconnected to the constants and types used to represent the specification logic without affecting the adequacy of the encoding.

Another point of interest is the following condition of the previous theorem: specification logic \forall -quantification and meta-logic ∇ -quantification are allowed only at inhabited types. This condition arises because we have chosen to do a shallow encoding of the typing judgment of the specification logic. That is, rather than encode an explicit typing judgment for specification logic terms, we have instead relied on the typing judgment of \mathcal{G} to enforce the well-formedness of terms. Due to the lack of restrictions on the occurrences of nominal constants, the typing judgment in \mathcal{G} is more permissive than the specification logic typing. As the previous theorem shows, however, this difference only manifests itself for uninhabited types. A deeper encoding involving an explicit typing judgment would avoid this condition, but would also impose some overhead additional costs in terms of reasoning about and through the encoding. We find the shallow encoding to be a good balance in practice.

6.5.2 Adequacy of Type Preservation Example

We can now use our adequacy results to extract a proof of type preservation for the simply-typed λ -calculus from the proof of its encoding in \mathcal{G} .

Theorem 6.5.2. *If $t \Downarrow v$ and $\vdash t : a$ then $\vdash v : a$.*

Proof. Suppose $t \Downarrow v$ and $\vdash t : a$, then by the adequacy results in Section 2.4, we know that $\Delta \vdash \text{eval } \phi(t) \phi(v)$ and $\Delta \vdash \text{of } \phi(t) \phi(a)$ have derivations in hH^2 where ϕ is the bijection between the object language and its specification logic representation and Δ is the specification of *eval* and *of*. By Theorem 6.5.1, we know $\Vdash \langle \text{eval } \psi(\phi(t)) \psi(\phi(v)) \rangle$ and $\Vdash \langle \text{of } \psi(\phi(t)) \psi(\phi(v)) \rangle$ have proofs in \mathcal{G} . Using these proofs and the proof of the formula in Theorem 6.3.1 together with various rules of \mathcal{G} (notably the *cut* rule), we can construct a proof of $\Vdash \langle \text{of } \psi(\phi(v)) \psi(\phi(a)) \rangle$ in \mathcal{G} . Then using the backwards direction of Theorem 6.5.1 we know $\Delta \vdash \text{of } \phi(v) \phi(a)$ has a derivation in hH^2 , and using adequacy results from Section 2.4 we find that $\vdash v : a$ must hold. \square

Chapter 7

Applications of The Framework

In this chapter we consider various applications of the proposed framework, focusing mainly on the reasoning component. The purpose of these applications is illustrate both the strengths and the weaknesses of the framework. From this perspective, we are interested in the *quality* of the encodings and associated reasoning, *e.g.*, properties such as naturalness, expressiveness, complexity, and overhead. We will try to expose and highlight these traits in this chapter.

We begin in Section 7.1 with a proof of type uniqueness for the simply-typed λ -calculus which provides a simple example of how judgment contexts and the related variable freshness information is handled in the framework. In Section 7.2 we present a solution to part of the POPLmark challenge [ABF⁺05] which demonstrates the more sophisticated inductive reasoning that is possible within \mathcal{G} . Section 7.3 contains an example of proving the equivalence of λ -terms based on the set of paths they contain, and shows how easily the framework handles formulas with a more sophisticated quantification structure. In Section 7.4 we describe a translation between higher-order abstract syntax and de Bruijn notation for λ -terms, and we show that this translation is deterministic in both directions. This example highlights a more expressive use of definitions to describe the structure of judgment contexts. Finally, in Section 7.5 we show how Girard’s proof of strong normalization for the simply-typed λ -calculus can be encoded. This is by far the largest application in this chapter, and it uses many of the features highlighted by previous examples as well as introducing new ones such as a way of dealing with an arbitrary number of substitutions applied to a term.

There have been many other applications of the reasoning component of our framework that we do not discuss explicitly in this thesis. These include the following.

- Properties of big and small step evaluation and typing in the simply-typed λ -calculus
- Translation among combinatory logic, natural deduction, and sequent calculus
- Soundness and completeness for a focused sequent calculus
- Cut-admissibility for LJ
- Takahashi's proof of the Church-Rosser theorem
- Properties of bi-simulation in CCS and the π -calculus
- Tait's argument for weak normalization of the simply-typed λ -calculus [GMN08b].
- The substitution theorem for Canonical LF.

All of the applications mentioned above and the ones presented in this chapter are available on the Abella website [Gac09]. We note that some of these examples have been developed by other researchers. Randy Pollack contributed the formalization of the Church-Rosser result. The formalization of the substitution theorem for Canonical LF was contributed by Todd Wilson and is the largest development done in Abella to date. This development includes two sophisticated results: one which uses a triply nested induction where the innermost induction is an eight-way mutual induction and another which uses a doubly nested induction with an outer strong induction and an inner three-way mutual induction. The richness and elegance of this development serves as a powerful example of the expressivity of Abella.

Finally, before we proceed to the examples we establish a few common items and conventions which simplify the presentation. First, in specification formulas we elide the outermost universal quantifiers and assume that tokens given by capital letters denote variables that are implicitly universally quantified over the entire formula. Second, for judgments of the form $(L \Vdash \langle A \rangle)$ we write simply $(L \Vdash A)$ since we will only ever display this with atomic formulas on the right of the judgment. We assume the following definition of *name* (with

appropriate type based on the application):

$$(\nabla x.name\ x) \triangleq \top.$$

We will use the following result about the (non)occurrences of nominal constants in lists:

$$\forall L, E. \nabla x. member\ (E\ x)\ L \supset \exists E'. (E = \lambda y. E').$$

This says that if an element of a list depends on a nominal constant and the list itself does not, then the element's dependency must be vacuous. The proof is by induction on the *member* hypothesis. We will leave out the details of most proofs except to note the uses of induction or the particularly interesting cases. Also, we will freely and implicitly make use of the properties of the specification logic.

7.1 Type-uniqueness for the Simply-typed λ -calculus

The type of a λ -term in the simply-typed λ -calculus is unique. Proving this type uniqueness property requires reasoning inductively about typing judgments which, in turn, requires generalizing the context in which typing judgments are made. We can encode such arguments directly in our framework so long as we can describe the structure of the judgment contexts. Such descriptions can be naturally expressed using nominal abstraction and, in fact, this is the most common use of nominal abstraction. Thus, we use the present example to demonstrate how nominal abstraction can be used in this way and to point out the related lemmas that often go along with such descriptions.

We will use the specification of the simply-typed λ -calculus developed thus far in the thesis (Section 2.3). Relative to this, we can formally state type uniqueness as

$$\forall E, T_1, T_2. (\Vdash\ of\ E\ T_1) \supset (\Vdash\ of\ E\ T_2) \supset (T_1 = T_2).$$

Suppose we try to prove this directly by induction on one of the typing judgments. Then, when we consider the case where E is an abstraction, the typing context will grow which

$$\begin{aligned}
\text{ctx } nil &\stackrel{\mu}{=} \top \\
\text{ctx } (\text{of } X \ A :: L) &\stackrel{\mu}{=} (\forall M, N. X = \text{app } M \ N \supset \perp) \wedge \\
&\quad (\forall R, B. X = \text{abs } B \ R \supset \perp) \wedge \\
&\quad (\forall B. \text{member } (\text{of } X \ B) \ L \supset \perp) \wedge \\
&\quad \text{ctx } L
\end{aligned}$$

Figure 7.1: Potential *ctx* definition without nominal abstraction

means the inductive hypothesis will not be able to apply. Instead, we need to generalize the statement of type uniqueness to the following.

$$\forall L, E, T_1, T_2. \text{ctx } L \supset (L \Vdash \text{of } E \ T_1) \supset (L \Vdash \text{of } E \ T_2) \supset (T_1 = T_2).$$

Where *ctx* is a definition which restricts L so that the formula is provable. In particular, $\text{ctx } L$ should enforce that L has the structure $(x_1, A_1) :: \dots :: (x_n, A_n) :: nil$ where each x_i is atomic and unique. In the logics which preceded \mathcal{G} , these atomicity and uniqueness properties could not be directly described and instead one needed to encode them by explicitly excluding the other possibilities as shown in Figure 7.1. However, using nominal abstraction we define *ctx* as

$$\text{ctx } nil \stackrel{\mu}{=} \top \qquad (\nabla x. \text{ctx } (\text{of } x \ A :: L)) \stackrel{\mu}{=} \text{ctx } L.$$

Note that in $(\text{of } x \ A :: L)$, the atomicity of x is enforced by it being ∇ quantified while the uniqueness is enforced by L being quantified outside the scope of x . Had we wanted to allow x to occur later in the context we could have written $(L \ x)$ in place of L .

The definition of *ctx* enforces atomicity and uniqueness properties for the first element of the context and then calls itself recursively on the remaining portion of the context. Thus, to know that an arbitrary element of the context has the atomicity and uniqueness properties requires inductive reasoning. We state these properties in the following two

lemmas.

$$\forall L, X, A. \text{ ctx } L \supset \text{ member } (\text{of } X \ A) \ L \supset \text{ name } X$$

$$\forall L, X, A_1, A_2. \text{ ctx } L \supset \text{ member } (\text{of } X \ A_1) \ L \supset \text{ member } (\text{of } X \ A_2) \supset (A_1 = A_2)$$

Both of these lemmas have direct proofs using induction on one of the *member* hypotheses.

With the above lemmas in place, the rest of the type uniqueness proof is straightforward. There is an interesting point to be noted here, though, concerning the treatment of abstractions, *i.e.*, when considering the typing in the context L of a λ -term of the form $\text{abs } A \ R$. The use of a universal quantifier in the specification of typing in this case and the interpretation in the meta-logic of such universal quantifiers via ∇ -quantifiers ensures that the typing of R will be done in a context given by $\text{of } x \ A :: L$ where x is a nominal constant not appearing in L . In the type uniqueness proof, we will need to show that this extended typing context is well-formed. This is done by showing that $\text{ctx } (\text{of } x \ A :: L)$ follows from $\text{ctx } L$ which is clear based on the definition of ctx and the way x was introduced in the typing process. If a definition such as in Figure 7.1 were used, this argument would be more complicated.

7.2 The POPLmark Challenge

The POPLmark challenge is a call to researchers to develop tools and methodologies for animating and for reasoning about systems with binding [ABF⁺05]. The particular challenge proposed focuses on System F_{\leq} , a polymorphic λ -calculus with subtyping [CMMS94, CG94]. This challenge is of interest to us primarily because it provides a common benchmark on which various frameworks may be compared. In addition, some of the reasoning required for this problem illustrates the sophistication and naturalness of the reasoning tools available in our framework.

The POPLmark challenge consists of three challenge problems which focus on 1) the type system, 2) evaluation, type preservation, and progress, and 3) animation. In this section we explain the solution to the first challenge problem which requires sophisticated induc-

$$\begin{array}{c}
\Gamma \vdash S \prec: \text{Top} \qquad\qquad\qquad (\text{SA-Top}) \\
\\
\Gamma \vdash X \prec: X \qquad\qquad\qquad (\text{SA-Ref-TVar}) \\
\\
\frac{X \prec: U \in \Gamma \quad \Gamma \vdash U \prec: T}{\Gamma \vdash X \prec: T} \qquad\qquad\qquad (\text{SA-Trans-TVar}) \\
\\
\frac{\Gamma \vdash T_1 \prec: S_1 \quad \Gamma \vdash S_2 \prec: T_2}{\Gamma \vdash S_1 \rightarrow S_2 \prec: T_1 \rightarrow T_2} \qquad\qquad\qquad (\text{SA-Arrow}) \\
\\
\frac{\Gamma \vdash T_1 \prec: S_1 \quad \Gamma, X \prec: T_1 \vdash S_2 \prec: T_2}{\Gamma \vdash (\forall X \prec: S_1. S_2) \prec: (\forall X \prec: T_1. T_2)} \qquad\qquad\qquad (\text{SA-All})
\end{array}$$

Figure 7.2: Algorithmic subtyping rules for System F_{\prec} :

tion schemes and some reasoning about binding structure. The second challenge problem requires a significant amount of reasoning about binding structure, but since we take binding as fundamental in our framework, this challenge problem is straightforward and fairly mundane in our framework (the development is available on the Abella website). Finally, the last challenge problem could be addressed through an animation system for λ Prolog, but we do not explore this in this section. The first and second challenge problems also have an additional component that asks for proofs to be repeated for System F_{\prec} : extended with records and patterns. This extension requires a significant amount of additional work without providing much additional insight in the framework, and thus we do not pursue this extension.

The first POPLmark challenge problem focuses on the type system of System F_{\prec} . In particular, given an algorithmic presentation of the subtyping rules for System F_{\prec} , the challenge asks one to show that the subtyping relation is reflexive and transitive, the key results needed to show equivalence between the algorithmic and declarative descriptions of subtyping. Reflexivity turns out to be straightforward, while transitivity requires sophisticated inductive reasoning. In the rest of this section we focus on the proof of transitivity.

Types and typing contexts in System F_{\prec} are described by the following grammars.

$$\begin{aligned} T &::= X \mid \text{Top} \mid T \rightarrow T \mid \forall X \prec: T. T \\ \Gamma &::= \emptyset \mid \Gamma, X \prec: T \end{aligned}$$

Here X denotes a variable occurrence, and $\forall X \prec: T_1. T_2$ denotes that the variable X is bound within the scope of T_2 (but not in the scope of T_1). In $\Gamma, X \prec: T$ it is assumed that X does not occur in Γ . The algorithmic subtyping relation of System F_{\prec} is denoted by $\Gamma \vdash S \prec: T$, and is defined by the rules in Figure 7.2.

The challenge problem is to prove that the subtyping relation is transitive: if $\Gamma \vdash S \prec: Q$ and $\Gamma \vdash Q \prec: T$ then $\Gamma \vdash S \prec: T$. The proof of this property requires another result called narrowing to be proved simultaneously: if $\Gamma, X \prec: Q, \Delta \vdash M \prec: N$ and $\Gamma \vdash P \prec: Q$ then $\Gamma, X \prec: P, \Delta \vdash M \prec: N$. The proof of these two properties requires a mutual induction on the structure of the type Q . Within this induction the transitivity property is proved by induction on the structure of $\Gamma \vdash S \prec: Q$ and it uses the narrowing property for structurally smaller types Q . The narrowing property is proved by an inner induction on the structure of $\Gamma, X \prec: Q, \Delta \vdash M \prec: N$ and uses the transitivity property for the type Q . With the proper induction schemes as described, the details of the proof are straightforward.

To formalize System F_{\prec} types we introduce the type ty and the following constants.

$$\text{top} : ty \qquad \text{arrow} : ty \rightarrow ty \rightarrow ty \qquad \text{all} : ty \rightarrow (ty \rightarrow ty) \rightarrow ty$$

Typing contexts will be represented using the context of specification logic judgments. We introduce the constant $\text{bound} : ty \rightarrow ty \rightarrow o$ for representing individual type bindings within that context.

We encode subtyping rules of System F_{\prec} as specification logic formulas concerning the constant $\text{sub} : ty \rightarrow ty \rightarrow o$ as presented in Figure 7.3. Note that we do not explicitly represent the typing context, but instead make assumptions of the form $\text{bound } X \ T$ to denote a typing assumption of $X \prec: T$. Also, in the formal rules SA-REFL-TVAR and SA-TRANS-TVAR the variable X represents only type variables while our translation of these

$$\begin{aligned}
& \text{sub } S \text{ top} \\
& \text{bound } X \ U \supset \text{sub } X \ X \\
& \text{bound } X \ U \supset \text{sub } U \ T \supset \text{sub } X \ T \\
& \text{sub } T_1 \ S_1 \supset \text{sub } S_2 \ T_2 \supset \text{sub } (\text{arrow } S_1 \ S_2) \ (\text{arrow } T_1 \ T_2) \\
& \text{sub } T_1 \ S_1 \supset (\forall x. \text{bound } x \ T_1 \supset \text{sub } (S_2 \ x) \ (T_2 \ x)) \supset \text{sub } (\text{all } S_1 \ S_2) \ (\text{all } T_1 \ T_2)
\end{aligned}$$
Figure 7.3: Specification of algorithmic subtyping for System F_{\leq} :

rules do not directly enforce this constraint. Instead, our translations require that any such X satisfy a $\text{bound } X \ U$ judgment for some U . Since we only make such judgments for X which denotes a type variable, our encoding remains adequate.

To reason about subtyping we first formalize the notion that a typing context is well-formed. Strictly speaking, a context is well-formed if it is either \emptyset or $\Gamma, X \prec T$ where X is a variable which does not occur in Γ . For reasons we discuss later, we deliberately weaken this notion and require only that X is a variable. We recognized such well-formed contexts with the following definition.

$$\text{ctx } nil \stackrel{\mu}{=} \top \qquad \text{ctx } (\text{bound } X \ U :: L) \stackrel{\mu}{=} \text{name } X \wedge \text{ctx } L$$

We also prove the following associated lemma.

$$\forall E, L. \text{ctx } L \supset \text{member } E \ L \supset \exists X, U. (E = \text{bound } X \ U) \wedge \text{name } X$$

This is proved by a simple induction on the *member* hypothesis.

The logic \mathcal{G} allows for induction only on definitions and not on terms. Thus to induct on the structure of a System F_{\leq} type we must create a definition which recognizes such

types. We define a predicate $wfty : ty \rightarrow o$ as follows.

$$\begin{aligned} wfty \text{ top} &\stackrel{\mu}{=} \top \\ (\nabla x. wfty x) &\stackrel{\mu}{=} \top \\ wfty (\text{arrow } T_1 T_2) &\stackrel{\mu}{=} wfty T_1 \wedge wfty T_2 \\ wfty (\text{all } T_1 T_2) &\stackrel{\mu}{=} wfty T_1 \wedge \nabla x. wfty (T_2 x) \end{aligned}$$

Induction on $wfty Q$ will correspond to structural induction on the type Q as needed. Note that we could impose additional well-formedness constraints which restrict variable occurrences relative to some context of type variables, but such restrictions are unnecessary for the proof at hand.

We can state the combined transitivity and narrowing property as follows.


$$\begin{aligned} \forall Q. wfty Q \supset & \\ (\forall L, S, T. \text{ctx } L \supset (L \Vdash \text{sub } S Q) \supset (L \Vdash \text{sub } Q T) \supset (L \Vdash \text{sub } S T)) \wedge & \\ (\forall L, P, X, M, N. \text{ctx } (\text{bound } X Q :: L) \supset (L \Vdash \text{sub } P Q) \supset & \\ (\text{bound } X Q :: L \Vdash \text{sub } M N) \supset (\text{bound } X P :: L \Vdash \text{sub } M N)) & \end{aligned}$$

The proof is by an outer induction on $wfty Q$. To prove the inner conjunction we use the following derived rule of \mathcal{G} .

$$\frac{\Gamma \longrightarrow B \quad \Gamma, B \longrightarrow C}{\Gamma \longrightarrow B \wedge C} \wedge \mathcal{R}^*$$

This rule is clearly admissible using *cut* and $\wedge \mathcal{R}$. We use this rule with B as the transitivity result for the type Q and C as the narrowing result for the type Q . Thus this rule allows us to use the transitivity result for the type Q while proving the corresponding narrowing result. Once this is applied we can prove transitivity using a further induction on $(L \Vdash \text{sub } S Q)$ and narrowing using a further induction on $(\text{bound } X Q :: L \Vdash \text{sub } M N)$. The reasoning which remains is straightforward.

Notice that in the original statement of narrowing, the distinguished typing assumption $X \triangleleft Q$ is taken from the middle of the typing context, while in our formalized statement we consider the assumption $\text{bound } X Q$ only at the front. By formalizing narrowing in

Figure 7.4: Tree form of $\lambda x.(x(\lambda y.y))$

this way, we greatly simplify the associated reasoning (*e.g.*, we do not need to talk about appending contexts as we would with a direct statement). The cost is that when we add other elements to the context, we must show that the distinguished binding can always be moved to the front. This is possible since we have weakened the *ctx* judgment to not contain any freshness information, and therefore no ordering information. Since freshness information is not relevant to the transitivity and narrowing results, there is no cost to leaving this information out. To establish adequacy, we can use a more precise description of typing contexts and still make use of these results proved for the looser description.

7.3 Path Equivalence for λ -terms

We can characterize λ -terms by means of their paths, where a path formalizes the idea of descending through the abstract syntax tree of a term. For example, the tree for the λ -term $\lambda x.(x(\lambda y.y))$ is shown in Figure 7.3 has two paths:

1. descend through the binder for x , go left at the application, stop at x , and
2. descend through the binder for x , go right at the application, descend through the binder for y , stop at y

$$\begin{aligned}
& \text{term } M \supset \text{term } N \supset \text{term } (\text{app } M \ N) \\
& (\forall x. \text{term } x \supset \text{term } (R \ x)) \supset \text{term } (\text{abs } R) \\
& \text{path } M \ P \supset \text{path } (\text{app } M \ N) \ (\text{left } P) \\
& \text{path } N \ P \supset \text{path } (\text{app } M \ N) \ (\text{right } P) \\
& (\forall x. \forall p. \text{path } x \ p \supset \text{path } (R \ x) \ (S \ p)) \supset \text{path } (\text{abs } R) \ (\text{bnd } S)
\end{aligned}$$

Figure 7.5: Specification of paths through λ -terms

Our goal in this section is to show that if two λ -terms share all the same paths, then the terms must be equal. We call this the *path equivalence* property.

We are interested in the path equivalence property since it expresses a model checking-like property over terms with binding structure. This type of property is difficult or impossible to formalize in competing frameworks like Twelf [PS99] since expressing the hypothetical property that two λ -terms have all the same paths requires a sufficiently rich logic. However, in our framework, we find that this property can be stated and reasoned about directly. Also, this application illustrates how we can use definitions to describe the structure of multiple judgment contexts which have related structure. Finally, a complication in this application demonstrates the need for occasional vacuity properties to be established regarding the occurrences of nominal constants in terms.

We introduce a type tm for untyped λ -terms and pt for paths together with the following constructors.

$$\begin{aligned}
& \text{app} : tm \rightarrow tm \rightarrow tm & \text{abs} : (tm \rightarrow tm) \rightarrow tm \\
& \text{left} : pt \rightarrow pt & \text{right} : pt \rightarrow pt & \text{bnd} : (pt \rightarrow pt) \rightarrow pt
\end{aligned}$$

We then introduce the predicates $\text{term} : tm \rightarrow o$ and $\text{path} : tm \rightarrow pt \rightarrow o$ defined by the specification logic formulas in Figure 7.5.

Given this description of paths through λ -terms we can state the path equivalence property as follows.

$$\forall M, N. (\Vdash \text{term } M) \supset (\forall P. (\Vdash \text{path } M P) \supset (\Vdash \text{path } N P)) \supset (M = N)$$

Note that we have added the explicit assumption $(\Vdash \text{term } M)$ so that we can induct on the structure of M . Also, we have stated only that the paths in M are also in N , but not vice-versa. It turns out that this weaker property is sufficient to prove the result.

Before we can proceed with the proof of the above statement, we need to strengthen it. In particular, when M is an abstraction we need to consider how the contexts for the *term* and *path* judgments will grow. This is done with the following definition of *ctxs* which describes not only how each context grows, but how the two contexts are related.

$$\text{ctxs } \text{nil } \text{nil} \stackrel{\mu}{=} \top \quad (\nabla x. \nabla p. \text{ctxs } (\text{term } x :: L) (\text{path } x p :: K)) \stackrel{\mu}{=} \text{ctxs } L K$$

Along with this definition, we need the following lemmas which allow us to extract information about a term based on its membership in one of the contexts described by *ctxs*.

$$\forall X, L, K. \text{ctxs } L K \supset \text{member } (\text{term } X) L \supset$$

$$\text{name } X \wedge \exists P. \text{member } (\text{path } X P) K$$

$$\forall X, P, L, K. \text{ctxs } L K \supset \text{member } (\text{path } X P) K \supset \text{name } X \wedge \text{name } P$$

The proofs of both lemma are by straightforward induction on the *member* hypotheses.

We can state the strengthened equivalence property as follows.

$$\forall L, K, M, N. \text{ctxs } L K \supset (L \Vdash \text{term } M) \supset$$

$$(\forall P. (K \Vdash \text{path } M P) \supset (K \Vdash \text{path } N P)) \supset (M = N)$$

The proof of this statement is by induction on $(L \Vdash \text{term } M)$. In the base case we need the following lemma which is proved by induction one of the *member* hypotheses.

$$\forall L, K, X_1, X_2, P. \text{ctxs } L K \supset$$

$$\text{member } (\text{path } X_1 P) K \supset \text{member } (\text{path } X_2 P) K \supset (X_1 = X_2)$$

In the other cases of the proof, we need to show that the top-level constructor of M is also the top-level constructor for N . We do by finding a path through M and using the hypothesis that M and N share the same paths to find the same path in N . The top-level constructor of that path will determine the top-level constructors of M and N . However, this requires that we can always find a path through a term which we formalize this as the following lemma.

$$\forall L, K, M, P. \text{ctxs } L \ K \supset (L \Vdash \text{term } M) \supset \exists P. (K \Vdash \text{path } M \ P)$$

The proof of this lemma is by induction on $(L \Vdash \text{term } M)$.

There is one last complication in the proof of path equivalence which comes from the inductive case concerning abstractions. Suppose $M = \text{abs } R$ and $N = \text{abs } R'$. Here we know

$$\forall P. (K \Vdash \text{path } (\text{abs } R) \ P) \supset (K \Vdash \text{path } (\text{abs } R') \ P)$$

but in order to use the inductive hypothesis we must show

$$\forall P. (\text{path } x \ p :: K \Vdash \text{path } (R \ x) \ P) \supset (\text{path } x \ p :: K \Vdash \text{path } (R' \ x) \ P)$$

where x and p are nominal constants. Now the problem is that when we go to prove this latter formula, the $\forall\mathcal{R}$ rule says that we must replace P by $P' \ x \ p$ for some new eigenvariable P' . Note that P' is raised over both x and p even though the dependency on x must be vacuous. We must prove this vacuity to finish this case of the proof, and thus we need the following lemma.

$$\forall K, M, P. \nabla x, p. (\text{path } x \ p :: K \Vdash \text{path } (M \ x) \ (P \ x \ p)) \supset \exists P'. (P = \lambda z. P')$$

This is proved by induction on the *path* judgment. With this issue resolved, the rest of the path equivalence proof is straightforward.

As we have seen, the path equivalence property is expressed naturally in our framework through the use of a formula with a nested universal quantifier and implication. We briefly discuss the adequacy considerations regarding such a formula. The goal is to use the

path equivalence property proven in \mathcal{G} in order to prove the path equivalence property for the object system. To do this, we need to show that the hypotheses we have about the object system imply that there are proofs in \mathcal{G} of the corresponding hypotheses for the formalization of the path equivalence problem; if we can show this, then we will obtain the desired result by using the bijectivity of the mappings for terms. Looking more carefully at the hypothesis, we see that the main concern is showing that if every path in a λ -term m is a path in another λ -term n then the following is provable in \mathcal{G} :

$$\forall P. (\Vdash path \psi(\phi(m)) P) \supset (\Vdash path \psi(\phi(n)) P) \quad (7.1)$$

Here ϕ is the bijection between object terms and their specification logic representations, and ψ is the bijection between specification logic terms and their meta-logic representations.

To complete this discussion, we provide a sketch of how a proof of (7.1) might be constructed. We start with the knowledge that every path in m is a path in n . Then, assuming that the specification of *path* is adequate, we know that whenever $\Delta \vdash path \phi(m) \phi(p)$ has an hH^2 derivation, it must be that $\Delta \vdash path \phi(n) \phi(p)$ also has an hH^2 derivation where Δ is the specification of *path* and *term*. By the adequacy of *seq* established in Theorem 6.5.1, we know that whenever $\Vdash path \psi(\phi(m)) \psi(\phi(p))$ is provable in \mathcal{G} , it must be that $\Vdash path \psi(\phi(n)) \psi(\phi(p))$ is also provable in \mathcal{G} . We will use this knowledge shortly. Now to prove (7.1) in \mathcal{G} we start by applying the $\forall\mathcal{R}$ and $\supset\mathcal{R}$ rules. Then we repeatedly apply appropriate left rules starting with the assumption $\Vdash path \psi(\phi(m)) P$. Since $\psi(\phi(m))$ has no eigenvariables and *path* always deconstructs its first argument, this repeated application of left rules can be made to result only in sequents with no formulas on the left and where P is instantiated with a term such that $\Vdash path \psi(\phi(m)) P$ is provable in \mathcal{G} . Now using our knowledge from before and the assumption that ϕ and ψ are bijections, it must be that $\Vdash path \psi(\phi(n)) P$ is provable in \mathcal{G} . This is exactly the form of the right side of each of the sequents which results from the repeated application of left rules. Thus each such sequent must be provable, and therefore (7.1) must also be provable in \mathcal{G} .

$$\begin{aligned}
& \text{add } z \ C \ C. \\
& \text{add } A \ B \ C \supset \text{add } (s \ A) \ B \ (s \ C) \\
& \text{ho2db } M \ D \ M' \supset \text{ho2db } N \ D \ N' \supset \text{ho2db } (\text{app } M \ N) \ D \ (\text{dapp } M' \ N') \\
& \text{depth } X \ D_X \supset \text{add } D_X \ X' \ D \supset \text{ho2db } X \ D \ (\text{dvar } X') \\
& (\forall x. \text{depth } x \ D \supset \text{ho2db } (R \ x) \ (s \ D) \ R') \supset \text{ho2db } (\text{abs } R) \ D \ (\text{abs } R')
\end{aligned}$$

Figure 7.6: Specification of translation between HOAS and de Bruijn notation

7.4 Conversion between HOAS and de Bruijn Notation

De Bruijn notation is a first-order representation of binding which uses numeric indices to associate variable occurrences with their binders. More precisely, the index denoting a variable occurrence corresponds the number of abstractions between the occurrence and its binder. In this section we describe a translation between higher-order abstract syntax representation and de Bruijn notation for untyped λ -terms, and we prove that this translation is deterministic in both directions. This example highlights the use of a definition for describing a context which carries more than just variable freshness information.

We start by introducing the type tm for the higher-order abstract syntax representation of untyped λ -terms with the constructors $\text{app} : tm \rightarrow tm \rightarrow tm$ and $\text{abs} : (tm \rightarrow tm) \rightarrow tm$. For natural numbers we use the type nt with constructors $z : nt$ and $s : nt \rightarrow nt$. Finally, for de Bruijn notation terms we introduce the type db with the following constructors.

$$\text{dabs} : db \rightarrow db \qquad \text{dapp} : db \rightarrow db \rightarrow db \qquad \text{dvar} : nt \rightarrow db$$

We translate from higher-order abstract syntax to de Bruijn notation as follows. We walk over the structure of the term keeping track of the number of abstractions we have descended through. Whenever we come to an abstraction we use the context to record a new variable for that abstraction and the abstraction depth at which it was encountered.

When we encounter a variable occurrence, we subtract the current abstraction depth from the corresponding depth in the context to determine the index for that variable occurrence. Using the predicates $add : nt \rightarrow nt \rightarrow nt \rightarrow o$, $depth : tm \rightarrow nt \rightarrow o$, and $ho2db : tm \rightarrow nt \rightarrow db \rightarrow o$, the specification of the translation is presented in Figure 7.6.

Now there is a derivation of $ho2db M z M'$ if and only if M is a higher-order abstract syntax representation of the de Bruijn notation term M' . Moreover, note that the translation is symmetric: we could start with either M or M' and construct a derivation of $ho2db M z M'$ to determine a value for the other.

Now we want to show that the above translation is deterministic in both directions. In doing this, we will need to make certain properties of natural numbers explicit. For this we make use of the following two definitions.

$$\begin{array}{ll} nat\ z \stackrel{\mu}{=} \top & le\ A\ A \stackrel{\mu}{=} \top \\ nat\ (s\ A) \stackrel{\mu}{=} nat\ A & le\ A\ (s\ B) \stackrel{\mu}{=} le\ A\ B \end{array}$$

Along with these we prove the following arithmetic properties by straightforward induction.

$$\begin{array}{l} \forall A, B. le\ (s\ A)\ B \supset le\ A\ B \\ \forall A. nat\ A \supset le\ (s\ A)\ A \supset \perp \\ \forall A, B, C. (\Vdash add\ A\ B\ C) \supset le\ B\ C \\ \forall A_1, A_2, B, C. nat\ C \supset (\Vdash add\ A_1\ B\ C) \supset (\Vdash add\ A_2\ B\ C) \supset (A_1 = A_2) \\ \forall A, B_1, B_2, C. (\Vdash add\ A\ B_1\ C) \supset (\Vdash add\ A\ B_2\ C) \supset (B_1 = B_2) \end{array}$$

Note that we have made the assumption nat explicit in some of these to provide a target for induction.

Derivations of $ho2db$ will construct contexts of the form

$$depth\ x_n\ (s^n\ z) :: \dots :: depth\ x_2\ (s\ (s\ z)) :: depth\ x_1\ (s\ z) :: depth\ x_0\ z :: nil$$

where each x_i is unique. Moreover, the numbers associated with each x_i will also be unique since they are sequential. Each of these uniqueness properties will be needed to

show determinacy for one or the other direction of the translation. We can describe these contexts with the following definition.

$$dctx\ nil\ z \stackrel{\mu}{=} \top \quad (\nabla x. dctx\ (\text{depth } x\ D :: L)\ (s\ D)) \stackrel{\mu}{=} dctx\ L\ D$$

The corresponding lemma for *dctx* is as follows

$$\forall E, L, D. dctx\ L\ D \supset member\ E\ L \supset \exists X, D_X. (E = \text{depth } X\ D_X) \wedge name\ X$$

The proof is by induction on the *member* judgment. One complication related to contexts arises when we call *add* from within *ho2db*: the *add* judgments inherits the context from *ho2db*. This is a problem since all of our lemmas about *add* assume that it has an empty context. We can fix this by proving the following lemma.

$$\forall L, D, A, B, C. dctx\ L\ D \supset (L \Vdash add\ A\ B\ C) \supset (\Vdash add\ A\ B\ C)$$

This is proved by a simple induction on the *add* judgment.

Now let us consider the determinacy proof going from higher-order abstract syntax to de Bruijn notation. For this, we need the following lemma which says that each variable in the context has a unique index associated with it.

$$\begin{aligned} \forall L, D, X, D_1, D_2. dctx\ L\ D \supset \\ member\ (\text{depth } X\ D_1)\ L \supset member\ (\text{depth } X\ D_2)\ L \supset (D_1 = D_2) \end{aligned}$$

This is proved by a straightforward induction on one of the *member* hypotheses. Then we can prove the generalized determinacy result:

$$\begin{aligned} \forall L, M, M'_1, M'_2, D. dctx\ L\ D \supset \\ (L \Vdash ho2db\ M\ D\ M'_1) \supset (L \Vdash ho2db\ M\ D\ M'_2) \supset (M'_1 = M'_2). \end{aligned}$$

This is proved by induction on one of the *ho2db* judgments. We then apply this generalization with $L = nil$ and $D = z$ to get the specific determinacy result we care about.

To prove determinacy in the other direction we need a lemma which says that each index in the context has a unique variable associated with it. We can state this as

$$\forall L, D, X_1, X_2, D_X. \text{dctx } L D \supset \\ \text{member } (\text{depth } X_1 D_X) L \supset \text{member } (\text{depth } X_2 D_X) L \supset (X_1 = X_2).$$

This is proved by induction on one of the *member* hypotheses, however we need an additional result about the restrictions on indices in the context for the proof to go through. Specifically, the following lemma is required.

$$\forall L, D, D_X, X. \text{dctx } L D \supset \text{member } (\text{depth } X D_X) L \supset \text{le } D D_X \supset \perp$$

This is proved by induction on the *member* hypothesis and in turn requires the following result which follows by a simple induction.

$$\forall L, D. \text{dctx } L D \supset \text{nat } D$$

With these lemmas in place, the generalized determinacy result is as follows.

$$\forall L, M_1, M_2, D, M'. \text{dctx } L D \supset \\ (L \Vdash \text{ho2db } M_1 D M') \supset (L \Vdash \text{ho2db } M_2 D M') \supset (M_1 = M_2)$$

This is now proved by straightforward induction on one of the *ho2db* hypotheses, and again we can substitution $L = \text{nil}$ and $D = z$ to obtain the specialized result.

7.5 Formalizing Tait-Style Proofs for Strong Normalization

Tait introduced the idea of a logical relation and showed how this could be used to provide an elegant proof of the strong normalization property for the typed λ -calculus [Tai67]. Girard subsequently generalized this idea to obtain a strong normalization result for the computationally much richer second-order λ -calculus or System F [Gir72]. This style of argument has both an elegance and a sophistication that would be interesting to see captured

$$\begin{aligned}
& \text{type } i \\
& \text{type } A \supset \text{type } B \supset \text{type } (\text{arrow } A B) \\
& \text{of } M (\text{arrow } A B) \supset \text{of } N A \supset \text{of } (\text{app } M N) B \\
& \text{type } A \supset (\forall x. \text{of } x A \supset \text{of } (R x) B) \supset \text{of } (\text{abs } A R) (\text{arrow } A B) \\
& \text{type } A \supset \text{of } c A \\
& \text{step } M M' \supset \text{step } (\text{app } M N) (\text{app } M' N) \\
& \text{step } N N' \supset \text{step } (\text{app } M N) (\text{app } M N') \\
& \text{step } (\text{app } (\text{abs } A R) M) (R M) \\
& (\forall x. \text{step } (R x) (R' x)) \supset \text{step } (\text{abs } A R) (\text{abs } A R')
\end{aligned}$$

Figure 7.7: Specification of typing and one-step reduction

in formalizations. We show in this section that our framework is up to the task by considering an encoding of the argument for the simply typed λ -calculus drawn from [GTL89]. One note, however, is that the strong normalization argument requires a definition for a logical relation which does not satisfy our current stratification restriction. We strongly believe that the stratification condition on definitions in \mathcal{G} could be weakened to allow this definition while preserving cut-elimination, but at present we have no corresponding cut-elimination proof.

To encode the simply-typed λ -calculus we use the familiar types ty and tm along with their constructors i , $arrow$, app , and abs . In Girard's argument he assumes that we are always working with open terms and can therefore always select a free variable at any type. Rather than explicitly representing this style of reasoning, we opt to introduce a constant $c : tm$ which we allow to take on any type. This does not impair the adequacy of our final result: if a term does not contain c then none of the terms it reduces to will contain it, and

therefore c has no effect on normalization. The specification of typing ($of : tm \rightarrow ty \rightarrow o$) and one-step reduction ($step : tm \rightarrow tm \rightarrow o$) is given in Figure 7.7. The specification includes a predicate a predicate $type : ty \rightarrow o$ to recognize types, which we use in the abstraction typing rule since this will be needed for later arguments. Also, we add a typing clause for c to allow it to take on any type.

Strong normalization says that all reduction paths eventually terminate. We can succinctly encode this property in the following definition.

$$sn\ M \stackrel{\mu}{=} \forall M'. (\Vdash step\ M\ M') \supset sn\ M'$$

Note that there is no explicit base case for sn , but if M has no reductions then $(\Vdash step\ M\ M')$ will be impossible and therefore $sn\ M$ will hold. Also, we will see that structural induction on the definition of sn corresponds to induction on the structure of the possible reductions from a term. The adequacy of sn can be established in the same manner as adequacy for the path equivalence application (Section 7.3). We can now state the goal of this section:

$$\forall M, A. (\Vdash of\ M\ A) \supset sn\ M$$

The rest of this section describes definitions and lemmas necessary to prove this formula.

7.5.1 Typing and One-step Reduction

In order to reason about typing judgments, we need to make explicit the structure of the contexts of such judgments. They are described by the following definition.

$$ctx\ nil \stackrel{\mu}{=} \top \qquad (\forall x. ctx\ (of\ x\ A :: L)) \stackrel{\mu}{=} (\Vdash type\ A) \wedge ctx\ L$$

We then prove the corresponding lemma about context membership:

$$\forall E, L. ctx\ L \supset member\ E\ L \supset \exists X, A. (E = of\ X\ A) \wedge name\ X \wedge (\Vdash type\ A)$$

The proof is by induction the the *member* hypothesis. Another auxiliary lemma we need about typing says that we can extract *type* judgments from *of* judgments.

$$\forall L, M, A. ctx\ L \supset (L \Vdash of\ M\ A) \supset (\Vdash type\ A)$$

This is proved by induction on the *of* judgment and requires the following lemma which says that *type* judgments ignore typing contexts.

$$\forall L, A. \text{ctx } L \supset (L \Vdash \text{type } A) \supset (\Vdash \text{type } A)$$

This is proved by induction on the *type* judgment.

Now, the first real result we need is that one-step reduction preserves typing:

$$\forall L, M, M', A. \text{ctx } L \supset (L \Vdash \text{of } M \ A) \supset (\Vdash \text{step } M \ M') \supset (L \Vdash \text{of } M' \ A).$$

The proof is by induction on the *step* judgment. Note that we have to generalize the typing context since one-step reduction can take place underneath abstractions. Another useful lemma is the following.

$$\forall M. \text{sn } (\text{app } M \ c) \supset \text{sn } M$$

The proof is by induction on *sn*.

7.5.2 The Logical Relation

The difficulty with proving strong normalization directly is that it is not closed under application, *i.e.*, *sn* *M* and *sn* *N* does not imply *sn* (*app* *M* *N*). Instead, we must strengthen the normalization property to one which includes a notion of closure under application. This strengthened condition is called *reducibility* and is originally due to Tait [Tai67]. We say that a term *M* reduces at type *A* if *reduce* *M* *A* holds where *reduce* is defined as follows:

$$\begin{aligned} \text{reduce } M \ i &\stackrel{\mu}{=} (\Vdash \text{of } M \ i) \wedge \text{sn } M \\ \text{reduce } M \ (\text{arrow } A \ B) &\stackrel{\mu}{=} (\Vdash \text{of } M \ (\text{arrow } A \ B)) \wedge \\ &(\forall U. \text{reduce } U \ A \supset \text{reduce } (\text{app } M \ U) \ B) \end{aligned}$$

Note that *reduce* is defined with a negative use of itself and therefore does not satisfy the current stratification condition on definition. However, the second argument to *reduce* is smaller in the negative occurrence, and thus there are no logical loops introduced by this

definition. Intuitively, we can think of $(\lambda x. \text{reduce } x \ A)$ as defining a separate fixed-point for each type A , and that these fixed-points are constructed based on induction on A .

An auxiliary notion used when discussing reducibility is called *neutrality*: a term is called *neutral* if it is not an abstraction. We can define this directly as follows.

$$\text{neutral } M \triangleq \forall A, R. (M = \text{abs } A \ R) \supset \perp$$

Now Girard lays out three properties of reducibility which we can formalize as follows.

$$\text{(CR 1)} \ \forall M, A. (\Vdash \text{type } A) \supset \text{reduce } M \ A \supset \text{sn } M$$

$$\text{(CR 2)} \ \forall M, M', A. (\Vdash \text{type } A) \supset \text{reduce } M \ A \supset (\Vdash \text{step } M \ M') \supset \text{reduce } M' \ A$$

$$\text{(CR 3)} \ \forall M, A. (\Vdash \text{type } A) \supset \text{neutral } M \supset (\Vdash \text{of } M \ A) \supset$$

$$(\forall M'. (\Vdash \text{step } M \ M') \supset \text{reduce } M' \ A) \supset \text{reduce } M \ A$$

Each of these follows by induction on the *type* judgment. The proof of (CR 2) is straightforward, but the proofs (CR 1) and (CR 3) are more complicated. In particular, (CR 1) depends on (CR 3) at types structurally smaller than A while (CR 3) depends on (CR 1) at the same type A . As in the POPLmark application (Section 7.2) we can handle this by stating a combined lemma and using $\wedge \mathcal{R}^*$ within the induction:

$$\forall A. (\Vdash \text{type } A) \supset$$

$$(\forall M. \text{reduce } M \ A \supset \text{sn } M) \wedge$$

$$(\forall M. \text{neutral } M \supset (\Vdash \text{of } M \ A) \supset$$

$$(\forall M'. (\Vdash \text{step } M \ M') \supset \text{reduce } M' \ A) \supset \text{reduce } M \ A)$$

The proof is by induction on the *type* judgment, and the (CR 1) portion of the proof is relatively straightforward. In the (CR 3) portion, when A is an arrow type, say *arrow* $A_1 \ A_2$, we need to show

$$\forall U. \text{reduce } U \ A_1 \supset \text{reduce } (\text{app } M \ U) \ A_2.$$

From the (CR 1) inductive hypothesis on type A_1 we can determine that $\text{sn } A_1$ holds, and then proof is by an inner induction on $\text{sn } A_1$.

The last reducibility lemma we need says that if for all reducible U of type A , $M[U/x]$ is reducible, then so is $\lambda x : A. M$. For $\lambda x : A. M$ to be reducible requires showing that for all reducible V that $M V$ is reducible. Girard proves this by induction on the sum of the lengths of the longest reduction paths from M and V . We can state this unfolded reducibility lemma as follows.

$$\begin{aligned} \forall V, M, A, B. (\Vdash \text{of } (abs\ A\ M)\ (\text{arrow}\ A\ B)) \supset \\ sn\ V \supset sn\ (M\ c) \supset reduce\ V\ A \supset \\ (\forall U. reduce\ U\ A \supset reduce\ (M\ U)\ B) \supset \\ reduce\ (app\ (abs\ A\ M)\ V)\ B \end{aligned}$$

The proof of this formula is by induction on $sn\ V$ with a nested induction on $sn\ (M\ c)$.

Clearly *reduce* is closed under application and by (CR 1) it implies strong normalization, thus we strengthen our desired normalization result to the following:

$$\forall M, A. (\Vdash \text{of } M\ A) \supset reduce\ M\ A.$$

In order to prove this formula we will have to induct on the height of the proof of the typing judgment. However, when we consider the case that M is an abstraction, we will not be able to use the inductive hypothesis since *reduce* is defined only on closed terms, *i.e.*, those typeable in the empty context. The standard way to deal with this issue is to generalize the desired formula to say that if M , a possibly open term, has type A then each closed instantiation for all the free variables in M , say N , satisfies *reduce* $N\ A$. This requires a formal description of simultaneous substitutions that can “close” a term.

7.5.3 Arbitrary Cascading Substitutions and Freshness Results

Given $(L \Vdash \text{of } M\ A)$, *i.e.*, an open term and its typing context, we define a process of substituting each free variable in M with a value V which satisfies the logical relation for

the appropriate type. We define this *subst* relation as follows:

$$\begin{aligned} \text{subst nil } M \ M &\stackrel{\mu}{=} \top \\ (\nabla x.\text{subst } ((\text{of } x \ A) :: L) \ (R \ x) \ M) &\stackrel{\mu}{=} \exists U. \text{reduce } U \ A \wedge \text{subst } L \ (R \ U) \ M \end{aligned}$$

By employing nominal abstraction in the second clause, we are able to use the notion of substitution in the meta-logic to directly and succinctly encode substitution in the object language. Also note that we are, in fact, defining a process of cascading substitutions rather than simultaneous substitutions. Since the substitutions we define (using closed terms) do not affect each other, these two notions of substitution are equivalent. We will have to prove some part of this formally, of course, which in turn requires proving results about the (non)occurrences of nominal constants in our judgments.

One consequence of defining cascading substitutions via the notion of substitution in the meta-logic is that we do not get to specify where substitutions are applied in a term. In particular, given an abstraction $\text{abs } A \ R$ we cannot preclude the possibility that a substitution for a nominal constant in this term will affect the type A . Instead, we must show that well-formed types cannot contain free variables which we formalize as

$$\forall A.\nabla x. (\Vdash \text{type } (A \ x)) \supset \exists A'. (A = \lambda y.A')$$

This formula essentially states any dependencies a type has nominal constants must be vacuous. A related result is that in any provable judgment of the form $(L \Vdash \text{of } M \ A)$, any nominal constant (denoting a free variable) in M must also occur in L , *i.e.*,

$$\forall L, M, A.\nabla x. \text{ctx } L \supset (L \Vdash \text{of } (M \ x) \ (A \ x)) \supset \exists M'. (M = \lambda y.M')$$

This is proved by induction on the *of* judgment.

Given these results about the (non)occurrences of nominal constants in judgments, we can now prove fundamental properties of arbitrary cascading substitutions. The first property states that closed terms, those typeable in the empty context, are not affected by substitutions, *i.e.*,

$$\forall L, M, N, A. (\Vdash \text{of } M \ A) \supset \text{subst } L \ M \ N \supset (M = N).$$

The proof here is by induction on *subst* which corresponds to induction on the length of the list L . The key step within the proof is using the lemma that any nominal constant in the judgment (\Vdash of $M A$) must also be contained in the context of that judgment. Since the context is empty in this case, there are no nominal constants in M and thus the substitutions from L do not affect it.

We must show that our cascading substitutions act compositionally on terms in the simply-typed λ -calculus. For the term c this is almost trivial,

$$\forall L, M. \text{subst } L c M \supset (M = c).$$

The proof is by induction on *subst*. For application we have the following.

$$\begin{aligned} \forall L, M, N, U. \text{ctx } L \supset \text{subst } L (\text{app } M N) U \supset \\ \exists M_U, N_U. (U = \text{app } M_U N_U) \wedge \text{subst } L M M_U \wedge \text{subst } L N N_U \end{aligned}$$

This is proved by induction on *subst*. Finally, for abstractions we prove the following, also by induction on *subst*:

$$\begin{aligned} \forall L, A, R, U. \text{ctx } L \supset \text{subst } L (\text{abs } A R) U \supset (\Vdash \text{type } A) \supset \\ \exists R_U. (U = \text{abs } A R_U) \wedge \\ (\forall V. \text{reduce } V A \supset \nabla x. \text{subst } ((\text{of } x A) :: L) (R x) (R_U V)) \end{aligned}$$

Here we have the additional hypothesis of (\Vdash type A) to ensure that the substitutions created from L do not affect A . At one point in this proof we have to show that the order in which cascading substitutions are applied is irrelevant. The key to showing this is realizing that all substitutions are for closed terms. Since closed terms cannot contain any nominal constants, substitutions do not affect each other.

Finally, we must show that cascading substitutions preserve typing. Moreover, after applying a full cascading substitution for all the free variables in a term, that term should now be typeable in the empty context:

$$\forall L, M, N, A. \text{ctx } L \supset \text{subst } L M N \supset (L \Vdash \text{of } M A) \supset (\Vdash \text{of } N A).$$

This formula is proved by induction on *subst*.

7.5.4 The Final Result

Using cascading substitutions we can now formalize the generalization of strong normalization that we described earlier: given a (possibly open) well-typed term, every closed instantiation for it satisfies the logical relation *reduce*:

$$\forall L, M, N, A. \text{ ctx } L \supset (L \Vdash \text{ of } M \ A) \supset \text{ subst } L \ M \ N \supset \text{ reduce } N \ A$$

The proof of this formula is by induction on the typing judgment. The inductive cases are fairly straightforward using the compositional properties of cascading substitutions and various results about reducibility. In the base case, we must prove

$$\forall L, M, N, A. \text{ ctx } L \supset \text{ member (of } M \ A) \ L \supset \text{ subst } L \ M \ N \supset \text{ reduce } N \ A,$$

which is done by induction on *member*. Strong normalization is now a simple corollary where we take L to be *nil*. Thus we have proved

$$\forall M, A. (\Vdash \text{ of } M \ A) \supset \text{ sn } M.$$

Chapter 8

Related Work

There are many frameworks which can be used to specify, to prototype, and to reason about computational systems. Some of these are designed specifically for this purpose while others have a different motivation, but can achieve a similar result. In this chapter we present a selection of these frameworks and contrast their capabilities with the framework put forth in this thesis. As the contributions of this thesis are primarily in the reasoning part of the framework, we shall give extra attention to this component in the comparisons.

Our framework is based on a two-level logic approach to reasoning. We have found this to be very effective in practice, but one could use the logic \mathcal{G} in a single-level logic fashion as well. The frameworks in this chapter come in both varieties: some use a two-level logic approach to which we can compare directly, while others use a single-level logic approach. In either case, the differences due to the reasoning approach used are often overshadowed by the differences in the treatment of binding. Thus we shall often say very little about the reasoning approach except when comparing against another two-level logic framework.

We organize our comparison of frameworks around the techniques used to represent the binding structure of objects. This is by far the most salient characteristic of the frameworks, and has the largest effect on the succinctness and the quality of the corresponding reasoning. Thus we will focus on issues such as the representation of binding, determining equality modulo renaming of bound variables, capture-avoiding substitution, and representing judgments with side-conditions related to binding. We will use the example of the simply-typed λ -calculus from Section 1.2 to illustrate these issues. We will order our comparisons based on the kind of support for binding provided by the framework. Specifically, we will look at frameworks based on first-order, nominal, and higher-order representations.

8.1 First-order Representations

First-order representations provide no special treatment for binders. As a result, variables must be encoded using strings or integers and binding aspects must be captured through constructors. Further, mechanisms for manipulating and reasoning about binders must be developed by interpreting the constructors representing them on a case-by-case basis by users of the framework. On the other hand, the benefit of first-order representations is that many mature frameworks exist which support this type of representation. For example, languages like SML and Prolog can effectively prototype specifications written using a first-order representation, while in the reasoning phase, theorem provers like Coq [BC04], ACL2 [KMM00], and HOL [Har96] can operate directly on first-order representations. Our discussion in this section will focus not on any particular framework but rather on the benefits and costs of various first-order representations. In particular, we look at the three most common first-order representations: named, nameless, and locally nameless.

8.1.1 Named Representation

The most direct and naive approach to encoding binders is to assign each variable a fixed name. For instance, the term $(\lambda x:i. x)$ might be encoded as $(abs\ "x"\ i\ (var\ "x"))$. Here we have picked a particular name, x , to denote the otherwise arbitrary variable in the function. This representation is very natural, but it creates at least three major problems for users.

First, equality modulo the renaming of bound variables is not reflected in the representation. For example, the terms $(\lambda x:i. x)$ and $(\lambda y:i. y)$ have two different representations, $(abs\ "x"\ i\ (var\ "x"))$ and $(abs\ "y"\ i\ (var\ "y"))$. Thus users of a named representation must explicitly define a notion of equivalence for each syntactic class with binding. This becomes particularly painful in reasoning where the user must establish many equivalence lemmas.

Second, no support is provided for capture-avoiding substitution over binding, and instead users must define this substitution on their own. Naive capture-avoiding substitution is not structurally recursive, and thus one must resort to well-founded recursion or instead

use simultaneous capture-avoiding substitution. Either choice results in additional overhead during reasoning when the user must prove various substitution lemmas. Moreover, substitution must be defined for each class of syntactic objects with binding, and the proofs of related lemmas must be repeated.

Third, no logical support is provided for treating side-conditions related to variable binding structure. An example of such a side-condition is manifest in the following rule for typing abstractions in the λ -calculus:

$$\frac{\Gamma, x : a \vdash r : b}{\Gamma \vdash (\lambda x : a. r) : a \rightarrow b} \quad x \notin \text{dom}(\Gamma).$$

With the named representation, users must devise their own mechanisms for treating such side-conditions. A naive approach in the case of the rule above is to select any fresh variable name, but this can lead to structural induction principles which are too weak to be usable in practice. Moreover, one must still prove that the choice for a variable name is truly arbitrary.

Large-scale developments have been constructed using the named representation, and the result is often that the binding issues overwhelm the development. For instance, VanInwegen used a named representation to encode and reason about SML in the HOL theorem prover [Van96]. She noted:

Proving theorems about substitutions (and related operations such as alpha-conversion) required far more time and HOL code than any other variety of theorems.

8.1.2 Nameless Representation

A more sophisticated first-order representation encodes each variable occurrence with an integer denoting the location of its binder relative to the binding structure around it. Commonly, one uses the distance from the variable occurrence to its binder, measured in terms of other binders above it in the abstract syntax tree. For example, the term $(\lambda x : i. (\lambda y : i. x))$ would be encoded as $(\text{abs } i (\text{abs } i (\text{var } 2)))$. Here the 2 denotes that the binder for this

variable occurrence is two binders away. This kind of representation originates from de Bruijn [dB72] and hence is often referred to as the de Bruijn representation.

The benefit of a nameless representation over a named representation is that α -equivalent terms, *i.e.*, those that differ only in the names of bound variables, are syntactically identical. Thus in the reasoning phase the user does not need to prove additional properties about α -equivalence.

The nameless representation shares many problems with the named representation and has some additional ones as well. The nameless representation still requires users to define capture-avoiding substitution themselves, and now this makes it necessary to reason about the correctness of the arithmetical operations that have to be carried out for maintaining the consistency of the representation when effecting substitutions. A new difficulty introduced by the nameless treatment of variables is that representations become hard for humans to read, since different occurrences of the same variable in them may be rendered into different integers depending on the contexts in which they appear. This also has an impact on the statements of lemmas and theorems that often need to explicitly talk about re-numberings and other arithmetical operations over terms, thereby diminishing clarity.

The nameless representation has been used in large-scale developments. Hirschhoff, for instance, used it to formalize the π -calculus in the Coq theorem prover [Hir97]. He found that the nameless representation simplified much of the work with bound variables versus the named representation, but the treatment of binding within it still overwhelmed the development. He concluded:

Technical work, however, still represents the biggest part of our implementation, mainly due to the managing of De Bruijn indexes [...] Of our 800 proved lemmas, about 600 are concerned with operators on free names.

8.1.3 Locally Nameless Representation

The most promising first-order representation is a hybrid approach which uses the nameless representation for bound variables and the named representation for free variables. This is

called the locally nameless representation [ACP⁺08, Cha09].

The locally nameless representation has advantages over both the named and nameless representations. First, α -equivalent terms are syntactically equal, as in the nameless representation. Second, the statement of lemmas and theorems rarely need to talk about arithmetical operations over terms. Third, since free and bound variables are syntactically distinguished, capture-avoiding substitution can be defined in a straightforward and structurally recursive way.

Like other first-order approaches, the locally nameless representation still requires users to define capture-avoiding substitution and prove various lemmas about it. A drawback specific to this representation is that users must provide functions which bind and unbind variables (*i.e.*, implementing the interface between the named and nameless representations). Constructing or deconstructing a term with binding requires going through these functions in order to ensure that certain invariants regarding free and bound variables are maintained. Finally, users must show that these binding and unbinding functions interact with substitution in appropriate ways. Recent progress has been made in automatically generating this type of infrastructure [AW09].

The locally nameless representation has some analogs to our own representation in the following sense: we represent bound variables using λ -terms and free variables using nominal constants. However, we provide capture-avoiding substitution for free to the user. Unbinding and binding of terms (*e.g.*, switching between λ -binders and nominal constants) is handled using application and nominal abstraction, respectively. In the locally nameless approach one occasionally needs to prove that free variables can be renamed while preserving provability, while that is an innate property of our framework due to our treatment of nominal constants. The fundamental contrast is that the locally nameless representation allows one to use an existing theorem prover, but requires significant binding infrastructure to be constructed, while our representation requires a new theorem prover, but incorporates binding infrastructure into the theory underlying the prover.

8.2 Nominal Representations

The nominal representation of binding is a mild extension of first-order abstract syntax with support for α -equivalence classes. The basis of the nominal representation is an infinite collection of names called atoms together with a freshness predicate—denoted by the infix operator $\#$ —between atoms and other objects and a swapping operation involving a pair of atoms and a term. Binding is represented by means of a term constructor $\langle \cdot \rangle$ which takes an atom and a term. The nominal representation then assumes certain properties of swapping and freshness with respect to this constructor so that α -equivalence classes are respected. This representation is also referred to as nominal abstract syntax.

Nominal representations were first introduced through the nominal logic of Pitts [Pit03], which is an extension of first-order logic. When working with nominal abstract syntax in a logical setting it is often desirable to quantify over fresh atoms. In this regard, a useful consequence of the properties assumed for freshness and swapping is that the following equivalence holds for any formula ϕ whose free variables are a, x_1, \dots, x_n where a is of atom type:

$$\exists a.(a\#x_1 \wedge \dots \wedge a\#x_n \wedge \phi) \quad \equiv \quad \forall a.(a\#x_1 \wedge \dots \wedge a\#x_n \supset \phi)$$

Nominal logic introduces the \mathcal{N} -quantifier by defining $\mathcal{N}a.\phi$ as one of the above formulas. This is very reminiscent of the properties shown for the ∇ -quantifier in Section 3.5.1, and in general, the ∇ -quantifier and the \mathcal{N} -quantifier behave very similarly.

The most prominent specification and prototyping language based on nominal representations is α Prolog, an extension of Prolog that accords a proof search interpretation of a version of Horn clauses in nominal logic [CU03]. In particular, α Prolog allows the \mathcal{N} -quantifier to appear in the heads of clauses. This allows α Prolog to describe specifications which involve a finer treatment of names than what is possible in our specification logic of hH^2 . However, it seems that α Prolog clauses bear a close resemblance to the patterned form of definitions in \mathcal{G} which allow the ∇ -quantifier in the head (see Section 3.4). While a formal encoding of α Prolog clauses as definitions in \mathcal{G} is left to future work, we note that

such definitions can be animated using a system similar to Bedwyr [BGM⁺07], a specification tool based on a simple proof search procedure for the Linc logic (one of the precursors to \mathcal{G}).

Nominal logic does not have a parallel to the fixed-point interpretation of definitions in \mathcal{G} , and thus nominal logic cannot be used directly to reason about specifications written within it. Instead, such reasoning must be carried out indirectly by first formalizing the relevant nominal logic specification in a richer logic such as that underlying a system like Coq or Isabelle/HOL and then using the capabilities of that logic [ABW06, UT05]. The most prominent development in this area is the Nominal package for Isabelle/HOL. This package allows for an easy definition of syntactic objects with α -equivalence classes. This construction is conducted completely within the HOL logic and can thus be trusted. Moreover, the construction of these α -equivalence classes and some boilerplate results about them are provided automatically via the macro-like features of Isabelle. This includes a strong induction principle which matches the one used in typical “pencil and paper” proofs, and it includes a recursion combinator which allows capture-avoiding substitution to be defined structurally.

The nominal approach has a number of drawbacks. First, binding is only simulated by means of a distinguished constructor and thus substitution is not automatically provided. Instead, users must define it on their own for both specification and reasoning, and consequently, must prove substitution lemmas relative to their definition of substitution. Second, in order to use functions and predicates in the reasoning phase, one must prove properties which state that name swapping does not change the results of a function or the provability of a predicate—a property which is enforceable statically for definitions of predicates in \mathcal{G} . Third, to effectively use the nominal representation in reasoning, one really needs an existing package which automates the construction of α -equivalence classes and proves the related lemmas. Although such a mature package exists for Isabelle/HOL, other theorem provers may not have the automation capabilities necessary to effectively construct such a package. Finally, an often trumpeted benefit of nominal representations is that they allow

a first-class treatment of names, but the analyses enabled by that treatment seem no more powerful than what is now provided by nominal abstraction. A formal validation of this observation is left to future work.

8.3 Higher-order Representations

Higher-order representations use the meta-level function space to encode binding in object languages, *e.g.*, by using data constructors such as $abs : (tm \rightarrow tm) \rightarrow tm$. This allows the object representation to inherit all the properties of binding from the meta-level. However, traditional tools often have a very strong notion of equality (*e.g.*, incorporating case analysis or fixed-point combinators) which makes them ill-suited to encoding higher-order representations. For this reason, we choose to focus here on frameworks based on the λ -tree syntax representation of binding which assumes only $\alpha\beta\eta$ -conversion in determining equality [Mil00]. This allows an adequate representation of object languages with binding, and provides free α -conversion and capture-avoiding substitution for those languages. The cost is that usually new frameworks must be developed which support the λ -tree syntax representation. In this section we discuss such frameworks which have been implemented.

8.3.1 Hybrid

Hybrid is a system which aims to support reasoning over higher-order abstract syntax specifications using traditional theorem provers such as Coq and Isabelle/HOL [FM09a]. The basic idea of the system is translate higher-order abstract syntax descriptions into an underlying de Bruijn representation. The logic of the theorem prover then serves as the meta-logic in which reasoning is conducted. This approach necessarily produces more overhead during reasoning due to the need occasionally to reason about the effects of the translation. However, there is good reason to believe that most of this can be automated in the future. Also, Hybrid is often used in a two-level logic approach using a specification logic which is essentially identical to our own hH^2 specification language.

The Hybrid system, by design, lacks a meta-logic with the tools to elegantly reason over

higher-order abstract syntax descriptions. Most notably, the meta-logics used by Hybrid lack a device like the ∇ -quantifier for reasoning about open terms and generic judgments. Recent work has suggested that such a device is not necessary for simple reasoning tasks such as type uniqueness arguments [FM09b]. Yet, it is unclear how the naive approach used in this work will scale to problems such as those proposed by the POPLmark Challenge [ABF⁺05]. In such problems one needs to recognize as equivalent those judgments which differ only in the renaming of free variables. Such a property is built into our meta-logic by representing such free variables by nominal constants, while in Hybrid one will have to manually develop and prove properties about notions of variable permutations.

8.3.2 Twelf

Twelf [PS99] is a system for specifying and reasoning with λ -tree syntax using LF, a dependently typed lambda calculus [HHP93]. In the LF methodology, object language judgments are encoded as LF types, and rules for making judgments are encoded as LF constructors for the corresponding types. The LF terms inhabiting these types are then derivations of judgments. Thus LF constitutes a specification language. Twelf implements an operational semantics for constructing LF terms which provides a means of animating LF specifications.

Since dependent types can be exploited in LF specifications, these can often be more elegant than those described in our simply-typed setting. For example, one can provide a definition of simply-typed λ -terms where the type of a λ -term is reflected in the type of its LF representation. When it is done in this way, one does not need to talk about pre-terms and provide a separate typing judgment for selecting well-typed terms. Moreover, this allows some properties to be obtained for free. For example, we can define evaluation over this representation of simply-typed λ -calculus so that type preservation is a direct consequence of the type of the evaluation judgment (*i.e.*, evaluation is defined to take a λ -term with a particular type and return another λ -term with the same type). However, in terms of expressive power, the simply-typed and dependently-typed specification languages are equivalent [Fel91]. Thus when referring to the example of the simply-typed λ -calculus

we will assume that it is encoded in LF in the same style as in our framework.

Since derivations of judgments are LF terms, we can think of defining further judgments over such terms. For example, suppose that we encode the simply-typed λ -calculus in LF including the type constructors *of* and *eval* corresponding to typing and evaluation judgments and the corresponding term constructors for forming those judgments. Then we could define a judgment named *preserve* which holds of a derivation of (*of t a*), a derivation of (*eval t v*), and a derivation of (*of v a*). Viewing this judgment as one which takes the first two arguments and produces the third, we could provide term constructors for *preserve* which describe how derivations of (*of t a*) and (*eval t v*) are used to reconstruct a derivation of (*of v a*). Twelf can then check that this judgment is total in its first two arguments, *i.e.*, it is defined and terminates for all inputs. If so, we can think of *preserve* as a proof of the meta-property that evaluation preserves typing in the simply-typed λ -calculus. This style of encoding is known as a Twelf meta-theorem.

The Twelf approach of encoding meta-theorems as LF judgments has some serious limitations. For example, consider the following statement of the type preservation theorem: “*forall* derivations of (*of t a*) and *forall* derivations of (*eval t v*) there *exists* a derivation of (*of v a*).” This theorem was encoded in an LF judgment which took the first two derivations as input and produced the last one as output. In general, a judgment representing a Twelf meta-theorem has inputs corresponding to \forall quantifiers and outputs corresponding to \exists quantifiers. Therefore, meta-theorems are restricted to a $\forall\exists$ quantification structure.

A related issue with the Twelf approach is that Twelf does not have a definition mechanism. Instead one has to use LF judgments to describe the properties of a specification. This is severely limiting since LF judgments can only describe behaviors that *may* happen and cannot describe those which *must* happen. For example, to state the strong normalization property for the simply-typed λ -calculus in Section 7.5, we used the following definition:

$$sn\ M \stackrel{\mu}{=} \forall M'. (\Vdash\ step\ M\ M') \supset sn\ M'$$

This says that in order for *sn M* to hold, every term to which *M* can convert *must* also satisfy *sn*. Such a definition is not possible with Twelf. A similar issue arises if one tries to

encode the path equivalence property for λ -terms from Section 7.3. The hypothesis in this case is that every path in one λ -term *must* occur in the other λ -term.

There is also a practical issue of relying on Twelf's totality checks in order to ensure that a meta-theorem is correct. It is possible, for example, for one to fill out the details of a meta-theorem so that totality holds, but for Twelf's checker to be unable to determine totality. In such a case, one must confront various options: 1) try to rewrite the meta-theorem so that totality is more evident, 2) wait for a new version of Twelf's totality checker that may be more powerful, or 3) do a careful hand proof of totality. The first option is not always possible, and the latter two are fairly undesirable.

An interesting comparison between the Twelf approach and our own is in the treatment of judgment contexts. In our approach, the definition of *seq* includes a list argument which keeps track of the context of a judgment and makes it explicit during reasoning. We then define a predicate like *ctx* which will recognize the structure of such a context, and we prove various inversion lemmas about membership in that context. In Twelf, such contexts are called regular worlds, and although they are declared explicitly, they are kept implicit during reasoning. The Twelf machinery automatically provides the associated inversion properties of regular worlds. Like most automation, this is very useful when it works and rather bothersome when it does not. For instance, in the conversion between higher-order abstract syntax and de Bruijn representations from Section 7.4, we work with a context which has an arithmetical property which depends on the judgment being made. Specifically, the context must not contain de Bruijn indices which are greater than the depth at which the conversion judgment is being made. This is needed to ensure uniqueness of de Bruijn indices when descending underneath abstractions. The regular worlds mechanism of Twelf does not allow the description of a context to depend on the arguments of the judgments made in that context. Thus one cannot express this property directly and must instead find a way to work around this limitation, *e.g.*, by making the context explicit [Cra08].

8.3.3 Delphin

Delphin is a higher-order functional programming language which operates over LF terms and can serve as a meta-logic for LF specifications [Pos08]. Delphin makes a distinction between LF functions which are purely representational (*i.e.*, that must be parametric in their argument) and Delphin functions which are computational (*i.e.*, that may perform case analysis on their argument). A Delphin meta-theorem is a Delphin function which is total. For example, the property of type preservation for the simply-typed λ -calculus is encoded as a function which takes LF terms denoting derivations of $(of\ t\ a)$ and $(eval\ t\ v)$ and returns an LF term denoting a derivation of $(of\ v\ a)$. Like Twelf, it is possible for Delphin not to be able to automatically determine totality of a meta-theorem, and then one must either rewrite the meta-theorem, wait for a stronger totality checker, or perform the totality check by hand.

The central way in which Delphin improves on Twelf is that it treats Delphin functions as first-class, and thus more sophisticated properties can be encoded during reasoning. For example, the path equivalence of λ -terms from Section 7.3 can be encoded fairly directly in Delphin. The property that all the paths in the λ -term s must also exist in the λ -term t can be represented in Delphin by a function which takes a judgment like $(path\ s\ p)$ and returns a judgment like $(path\ t\ p)$, and such a function can be an input (*i.e.*, hypothesis) to a Delphin meta-theorem stating the path equivalence property.

Delphin also uses first-class functions to treat the contexts of specification judgments. When a Delphin meta-theorem is written, it may make a recursive call to itself underneath some additional abstractions. These abstractions create new variables for which the Delphin meta-theorem must be defined. To achieve this, the Delphin meta-theorem carries around an argument which is a function mapping such variables to an appropriate invariant. This approach to representing contexts is more flexible than the regular worlds approach of Twelf. Specifically, in the example of conversion between higher-order abstract syntax and de Bruijn representations from Section 7.4, the dependency between the judgment and the context in the judgment can be made explicit in Delphin. Thus one can prove that the

conversion is deterministic in a fairly straightforward way in Delphin.

Despite the additional flexibility that Delphin provides in working with the contexts of judgments, it still does not make those contexts explicit as in our approach. Thus, some operations over contexts which we can perform easily in our framework are difficult or impossible in the Delphin approach. For example, in our formalization of Girard’s proof of strong normalization for the simply-typed λ -calculus in Section 7.5, we defined a process of closing a term by instantiating all free variables with closed terms of the appropriate types. This definition was based on walking over the context of the typing judgment of such a term, something that is not possible to do in Delphin.

8.3.4 Tac

Tac is a general framework for implementing logics. For the purposes of our present discussion, we will focus on the particular logic μLJ which is the most popular logic implemented in Tac [BMSV09b, Bae08a]. The logic μLJ comes from the same line of logics as \mathcal{G} and differs primarily in the semantics attributed to the ∇ -quantifier. We recall that the interpretation of ∇ in \mathcal{G} is derived from adding to $FO\lambda^{\Delta\mathbb{N}}$ the exchange and strengthening properties related to this quantifier that are embodied in the following equivalences:

$$\nabla x.\nabla y.F \equiv \nabla y.\nabla x.F \qquad \nabla x.F \equiv F, \text{ if } x \text{ does not occur in } F$$

The μLJ logic eschews these additions, strengthening the interpretation of the ∇ -quantifier instead through a capability to lift its predicative effect over types. At a practical, proof construction level, whereas the ∇ -quantifier can be treated in \mathcal{G} using nominal constants, in μLJ it must be treated by using explicit local contexts for each formula in a sequent. The size and ordering of the local context is always respected and instantiations for existentially or universally quantified variables may only use those generic variables which appear in the local context.

The μLJ logic does not have an operation like nominal abstraction and instead treats only equality. The issue with extending μLJ to treat nominal abstraction is that the pro-

cess of nominal capture-avoiding substitution (through which the nominal abstraction rules are defined) is based on carrying substitution information from one formula into all other formulas in a sequent. In the minimal setting, however, such information may be invalid in other formulas because the local signatures do not match. For example, a substitution which replaces M by a variable x from the local context does not make any sense in a formula which contains M but has an empty local context. As a result of this lack of nominal abstraction, the descriptions of properties such as the binding structure of specification judgment contexts in μLJ is less direct and thus harder to work with (see Figure 7.1 for an example). Furthermore, without nominal abstraction, one cannot directly formulate the invariants necessary to perform induction underneath ∇ (see Section 5.3.1). An ability of equivalent power is obtained in μLJ instead through the lifting capability mentioned earlier [Bae08b]. From a practical perspective, however, we find that reasoning based on lifting is often much more complicated than reasoning based on traditional induction combined with nominal abstraction.

The benefit of minimal treatment of the ∇ -quantifier is that the local context of a formula can be used to provide an adequate encoding for certain types of similar contexts in an encoding. This allows certain encodings to be shallower or to have fewer adequacy side-conditions than their counterparts in our setting. For example, in the statement of adequacy for our encoding of the specification logic into the predicate *seq* in Section 6.5.1 we have the requirement that ∇ -quantification is allowed only at inhabited types. This is necessary since if τ were an un-inhabited type then $\exists_{\tau}x.\top$ should not be provable in the specification logic, and yet its encoding as a *seq* judgment is provable if ∇ -quantification is allowed at type τ . The issue is that the specification logic existential quantifier is mapped to the meta-logic existential quantifier and the latter allows instantiations containing any nominal constants even if there are no other inhabitants at that type. If we take the definition of *seq* as being in μLJ then it should be an adequate encoding of the specification logic without any conditions. Thus the local context in the minimal approach provides an adequate representation of the variable signature of an hH^2 sequent. To achieve the same

condition-less adequacy for \mathcal{G} would require explicitly carrying around a representation of the specification logic signature and using this to restrict the type of instantiations for meta-logic universal and existential quantifiers. This approach would require more work due to the need to establish properties about the signature, but this is the same work which is already required in the minimal approach. Moreover, this explicit encoding of the signature would allow one to directly analyze and interact with the signature (*e.g.*, quantifying over all signatures of a certain type) which is not possible in the minimal approach.

Conclusion and Future Work

This thesis has concerned the development of a framework for specifying, prototyping, and reasoning about formal systems. The specific framework that has been of interest has two defining characteristics. First, it has been based on an intertwining of two distinct logics for specification and for reasoning about specifications. The specification logic has the property of also being executable, thereby rendering descriptions written in it transparently into prototypes of the formal systems that are encoded. The reasoning logic has the capability of directly embedding the specification logic; specifications themselves are represented indirectly through this medium. This is, in fact, the style of encoding that is developed here. The benefits of this approach are that the same specifications can be used for prototyping and reasoning and generic properties of the specification logic can be proved and used to advantage in reasoning. The second important characteristic of our framework is that uses a higher-order treatment of binding constructs, supporting this approach in both the specification and the reasoning levels through targeted logical devices.

The focus in this thesis has been on the reasoning component of the above framework. In this context, we have developed the logic \mathcal{G} that provides the mechanism of fixed-point definitions that can also be interpreted inductively or co-inductively and that has sophisticated devices for dealing with higher-order representations of syntactic constructs. An important component of this logic is the notion of nominal abstraction that allows for the reflection into definitions of properties of objects introduced into proofs in the course of treating binding constructs. We have used \mathcal{G} as the basis of an interactive theorem prover called Abella and have explored a two-level logic approach to reasoning about formal systems in its context. This system has been applied to several interesting reasoning examples

and has yielded appealing solutions in most of these situations.

While several promising results have been obtained in this thesis, there remain many more interesting things still to be done. We sketch below some possible ways in which the framework for specification, prototyping, and reasoning that has been considered can be further enriched. The kind of work involved in realizing these possibilities ranges from foundational considerations for increasing the expressive power of the meta-logic to more implementation oriented efforts to better facilitate the reasoning process.

9.1 More Permissive Stratification Conditions for Definitions

The stratification condition for definitions in \mathcal{G} is fairly simplistic, and it rules out seemingly well-behaved definitions such as the reducibility relation used in logical relations arguments (see Section 7.5). One could imagine a more sophisticated condition which would allow definitions to be stratified based on an ordering relation over the arguments of the predicate being defined. The proof theoretic arguments needed to prove cut-elimination for a logic with such definitions seem rather delicate, particularly since we allow substitutions which may interfere with any ordering based on term structure. From the perspective of developing the theory for such an extension, a first step might be to realize the addition to the Linc^- logic [TM09]. Given the way the cut-elimination proof for \mathcal{G} has been obtained from cut-elimination for Linc^- , if we can successfully carry out such an extension to Linc^- , the desired result relative to \mathcal{G} might then follow easily.

There is also an interaction of this line of research with the development of induction and co-induction. The strict notion of stratification that \mathcal{G} uses ensures that each definition describes a single fixed-point and the induction and co-induction rules operate on this structure. However, if we weaken the stratification condition, then each definition can be viewed as a possibly infinite collection of fixed-points. The rules for induction and co-induction must be carefully adapted in light of this fact.

$$\begin{aligned}
(\nabla x. \text{typeof}(L x) x A) &\triangleq \nabla x. \text{member}(\text{assm } x A) (L x) \\
\text{typeof } L (\text{app } M N) B &\triangleq \exists A. \text{typeof } L M (\text{arr } A B) \wedge \text{typeof } L N A \\
\text{typeof } L (\text{abs } A R) (\text{arr } A B) &\triangleq \nabla x. \text{typeof}((\text{assm } x A) :: L) (R x) B
\end{aligned}$$

Figure 9.1: Typing judgment directly within \mathcal{G}

9.2 Context Inversion Properties

When reasoning about specification judgments we often need to describe and utilize properties of the contexts in which those judgments are formed. This takes the form of stating a definition describing those contexts, proving various inversion lemmas about membership in those contexts, and then applying these lemmas at the appropriate times. Manually stating, proving, and using these lemmas introduces a fair amount of overhead which seems mundane enough that we might want to avoid it.

One option is to attack this problem with automation. One could imagine automatically generating and proving inversion properties for those definitions which can be seen as describing contexts. The inversion properties follow directly from the definitions, and the proofs are by simple inductive arguments. These lemmas could then be automatically applied anytime we have a member of such a context. However, it is unlikely that such automation of these properties would be able to cope with more complicated properties of contexts such as those used in the conversion between higher-order abstract syntax and the de Bruijn representation (see Section 7.4).

Another option would be to devise an alternate version of the specification logic or of its encoding in the meta-logic so that such context inversion properties are not needed as often. It is unclear how such alternatives would be developed, but as an analogy, consider the following. Typing for the simply-typed λ -calculus can be defined directly within \mathcal{G} via a definition of $(\text{typeof } L M A)$ which holds when M has type A in the typing context L . The clauses for this definition are presented in Figure 9.1. Using nominal abstraction, this

definition of typing directly precludes the possibility of looking anything up in the context which is not of the form $(asm\ x\ A)$ for some nominal constant x . Thus one does not need to deal with superfluous cases when performing case analysis on a typing judgment. Note, however, that uniqueness properties regarding the typing context would still need to be handled manually.

9.3 Types and Explicit Typing

The types in \mathcal{G} play no role in reasoning except to restrict the valid instantiations of quantifiers. Thus, for example, one cannot directly perform induction or case analysis on a term based on its type. Instead, one must create a definition which recognizes terms of that type, and then use induction or case analysis on that definition. This requires that one knows that the definition holds on the term, which in turn may require carrying around more explicit typing information in the specification or reasoning. All of this creates overhead just to work effectively with types. For example, in formalizing Girard's proof of strong normalization for the simply-typed λ -calculus (Section 7.5) we had to create a specification logic judgment which recognized well-formed types. This judgment was then carried around during reasoning, and it even had to be put into the specification of the object language typing judgment. We then had to prove a lemma which said that an object language type could not contain any nominal constants.

One possible solution is to attach explicit typing information to every variable in the specification and in reasoning. Ideally this should be done in such a way that the end user would not need to deal with explicit typing information, but would be able to perform operations like induction and case analysis based on the type of a term. A major difficulty in such automation would be dealing with the contexts needed to recognize terms which use higher-order abstract syntax. Multiple terms may have different contexts which have a particular relationship to each other which needs to be maintained. It is not clear how such information could be succinctly expressed.

9.4 Alternate Specification Logics

One motivation for the two-level logic approach to reasoning is that it lets us use general properties of a specification logic in reasoning about particular specifications. This approach has been successful relative to the second-order hereditary Harrop formula logic. However, different problem domains might require different specification logics. For example, a *linear specification logic* that allows for transient judgments has been found useful in characterizing properties of hardware [Chi95] and programming languages with references [MM02]. One can imagine an extension of the Abella system which allows different specification logics to be plugged in and used as particular reasoning tasks demand. Given the way our framework is designed, judgments from these different specification languages would be able to co-exist during reasoning.

9.5 Focusing and Proof Search

Recent research has been looking at techniques for guiding proof search in \mathcal{G} -like logics based on the notion of *focusing* [BM07, BMSV09a]. These techniques allow the automation of a significant portion of the reasoning process by pruning redundant choices. For example, it was proven that if an atomic judgment is to be inducted on during a proof, then this induction can be done immediately. These techniques have been effectively realized in the Tac theorem prover [BMSV09b]. The Abella system could also be extended to support this type of automation. Moreover, one should investigate how this automation interacts with the two-level logic approach to reasoning.

9.6 An Integrated Framework

The Teyjus system allows for animating descriptions in our specification logic and the Abella system allows for reasoning about such descriptions. It would be worthwhile to combine these systems into an integrated framework which enables a more fluid relationship between the processes of specification and reasoning. In its simplest form, such an integration would

allow the different aspects of prototyping and reasoning to be invoked seamlessly from a common description of a formal system. As an example of a deeper kind of integration looked at from the perspective of the reasoning component, uses of the *def \mathcal{R}* and *def \mathcal{L}* rules relative to the encodings of specifications within \mathcal{G} can draw benefit from computations within the specification logic. An important issue to be tackled in implementing such relationships would be that of designing an interface that allows a smooth transition between the different functionalities that Teyjus and Abella, the two currently separate components of our framework, provide.

Bibliography

- [ABF⁺05] Brian E. Aydemir, Aaron Bohannon, Matthew Fairbairn, J. Nathan Foster, Benjamin C. Pierce, Peter Sewell, Dimitrios Vytiniotis, Geoffrey Washburn, Stephanie Weirich, and Steve Zdancewic. Mechanized metatheory for the masses: The POPLmark challenge. In *Theorem Proving in Higher Order Logics: 18th International Conference*, number 3603 in LNCS, pages 50–65. Springer-Verlag, 2005.
- [ABW06] Brian Aydemir, Aaron Bohannon, and Stephanie Weirich. Nominal reasoning techniques in Coq. In *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP)*, Seattle, WA, USA, August 2006.
- [ACP⁺08] Brian Aydemir, Arthur Charguéraud, Benjamin C. Pierce, Randy Pollack, and Stephanie Weirich. Engineering formal metatheory. In *35th ACM Symposium on Principles of Programming Languages*, pages 3–15. ACM, January 2008.
- [AW09] Brian Aydemir and Stephanie Weirich. LNgén: Tool support for locally nameless representations. Available from <http://www.cis.upenn.edu/~baydemir/papers/lngen.pdf>, 2009.
- [Bae08a] David Baelde. *A linear approach to the proof-theory of least and greatest fixed points*. PhD thesis, Ecole Polytechnique, December 2008.
- [Bae08b] David Baelde. On the expressivity of minimal generic quantification. In A. Abel and C. Urban, editors, *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2008)*, number 228 in Electronic Notes in Theoretical Computer Science, pages 3–19, 2008.
- [Bar84] Hank Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, revised edition, 1984.
- [BC04] Yves Bertot and Pierre Castéran. *Interactive Theorem Proving and Program Development. Coq’Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. Springer Verlag, 2004.
- [BGM⁺06] David Baelde, Andrew Gacek, Dale Miller, Gopalan Nadathur, and Alwen Tiu. *A User Guide to Bedwyr*, November 2006.

- [BGM⁺07] David Baelde, Andrew Gacek, Dale Miller, Gopalan Nadathur, and Alwen Tiu. The Bedwyr system for model checking over syntactic expressions. In Frank Pfenning, editor, *21th Conference on Automated Deduction (CADE)*, number 4603 in LNAI, pages 391–397. Springer, 2007.
- [BM07] David Baelde and Dale Miller. Least and greatest fixed points in linear logic. In N. Dershowitz and A. Voronkov, editors, *International Conference on Logic for Programming and Automated Reasoning (LPAR)*, volume 4790 of LNCS, pages 92–106, 2007.
- [BMSV09a] David Baelde, Dale Miller, Zach Snow, and Alexandre Viel. Focused inductive theorem proving. Available from <http://www.lix.polytechnique.fr/Labo/Dale.Miller/papers/tac-draft.pdf>, 2009.
- [BMSV09b] David Baelde, Dale Miller, Zach Snow, and Alexandre Viel. Tac: A generic and adaptable interactive theorem prover. <http://slimmer.gforge.inria.fr/tac/>, 2009.
- [CG94] Pierre-Louis Curien and Giorgio Ghelli. Coherence of subsumption, minimum typing and type-checking in $F \leq$. *Theoretical aspects of object-oriented programming: types, semantics, and language design*, pages 247–292, 1994.
- [Cha09] Arthur Charguéraud. The locally nameless representation. Available from <http://arthur.chargueraud.org/research/2009/ln/>, July 2009.
- [Chi95] Jawahar Chirimar. *Proof Theoretic Approach to Specification Languages*. PhD thesis, University of Pennsylvania, February 1995.
- [Chu40] Alonzo Church. A formulation of the simple theory of types. *Journal of Symbolic Logic*, 5:56–68, 1940.
- [Chu41] Alonzo Church. *The Calculi of Lambda-Conversion*. Princeton University Press, 1941.
- [CMMS94] Luca Cardelli, John C. Mitchell, Simone Martini, and Andre Scedrov. An extension of System F with subtyping. *Information and Computation*, 109(1–2):4–56, 1994.
- [Cra08] Karl Cray. Explicit contexts in LF (extended abstract). In A. Abel and C. Urban, editors, *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2008)*, number 228 in Electronic Notes in Theoretical Computer Science, pages 53–68, 2008.
- [CU03] James Cheney and Christian Urban. System description: Alpha-Prolog, a fresh approach to logic programming modulo alpha-equivalence. In J. Levy, M. Kohlhase, J. Niehren, and M. Villaret, editors, *Proc. 17th Int. Workshop on*

- Unification, UNIF'03*, pages 15–19, Valencia, Spain, June 2003. Departamento de Sistemas Informáticos y Computación, Universidad Politécnica de Valencia. Technical Report DSIC-II/12/03.
- [CU08] James Cheney and Christian Urban. Nominal logic programming. *ACM Transactions on Programming Languages and Systems*, 30(5):1–47, 2008.
- [dB72] Nicolaas Govert de Bruijn. Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser Theorem. *Indagationes Mathematicae*, 34(5):381–392, 1972.
- [Eri91] Lars-Henrik Eriksson. A finitary version of the calculus of partial inductive definitions. In L.-H. Eriksson, L. Hallnäs, and P. Schroeder-Heister, editors, *Proceedings of the Second International Workshop on Extensions to Logic Programming*, volume 596 of *LNAI*, pages 89–134. Springer-Verlag, 1991.
- [Fel89] Amy Felty. *Specifying and Implementing Theorem Provers in a Higher-Order Logic Programming Language*. PhD thesis, University of Pennsylvania, August 1989.
- [Fel91] Amy Felty. Transforming specifications in a dependent-type lambda calculus to specifications in an intuitionistic logic. In Gérard Huet and Gordon D. Plotkin, editors, *Logical Frameworks*. Cambridge University Press, 1991.
- [FM09a] Amy Felty and Alberto Momigliano. Hybrid: A definitional two-level approach to reasoning with higher-order abstract syntax, 2009. Available from <http://arxiv.org/abs/0811.4367>.
- [FM09b] Amy Felty and Alberto Momigliano. Reasoning with hypothetical judgments and open terms in Hybrid. In *Proceedings of PPDP 2009: 11th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*, 2009. To appear.
- [Gac08] Andrew Gacek. The Abella interactive theorem prover (system description). In Alessandro Armando, Peter Baumgartner, and Gilles Dowek, editors, *Fourth International Joint Conference on Automated Reasoning*, volume 5195 of *LNCS*, pages 154–161. Springer, 2008.
- [Gac09] Andrew Gacek. The Abella system and homepage. <http://abella.cs.umn.edu/>, 2009.
- [GHN⁺08] Andrew Gacek, Steven Holte, Gopalan Nadathur, Xiaochu Qi, and Zach Snow. The Teyjus system – version 2. <http://teyjus.cs.umn.edu/>, March 2008.
- [Gir72] Jean-Yves Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. PhD thesis, Université Paris 7, 1972.

- [Gir92] Jean-Yves Girard. A fixpoint theorem in linear logic. An email posting to the mailing list `linear@cs.stanford.edu`, February 1992.
- [GMN08a] Andrew Gacek, Dale Miller, and Gopalan Nadathur. Combining generic judgments with recursive definitions. In F. Pfenning, editor, *23th Symposium on Logic in Computer Science*, pages 33–44. IEEE Computer Society Press, 2008.
- [GMN08b] Andrew Gacek, Dale Miller, and Gopalan Nadathur. Reasoning in Abella about structural operational semantics specifications. In A. Abel and C. Urban, editors, *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2008)*, number 228 in *Electronic Notes in Theoretical Computer Science*, pages 85–100, 2008.
- [GMN09] Andrew Gacek, Dale Miller, and Gopalan Nadathur. Nominal abstraction. Submitted. Available from <http://arxiv.org/abs/0908.1390>, 2009.
- [GTL89] Jean-Yves Girard, Paul Taylor, and Yves Lafont. *Proofs and Types*. Cambridge University Press, 1989.
- [Har96] John Harrison. HOL Light: A tutorial introduction. In *FMCAD*, pages 265–269, 1996.
- [HHP93] Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- [Hir97] Daniel Hirschhoff. A full formalization of pi-calculus theory in the Calculus of Constructions. In E. Gunter and A. Felty, editors, *International Conference on Theorem Proving in Higher Order Logics (TPHOLs'97)*, number 1275 in *LNCS*, pages 153–169, Murray Hill, New Jersey, August 1997.
- [Hue75] Gérard Huet. A unification algorithm for typed λ -calculus. *Theoretical Computer Science*, 1:27–57, 1975.
- [KMM00] Matt Kaufmann, J. Strother Moore, and Panagiotis Manolios. *Computer-Aided Reasoning: An Approach*. Kluwer Academic Publishers, Norwell, MA, USA, 2000.
- [Mil91] Dale Miller. A logic programming language with lambda-abstraction, function variables, and simple unification. *Journal of Logic and Computation*, 1(4):497–536, 1991.
- [Mil92] Dale Miller. Unification under a mixed prefix. *Journal of Symbolic Computation*, 14(4):321–358, 1992.
- [Mil00] Dale Miller. Abstract syntax for variable binders: An overview. In John Lloyd and et. al., editors, *Computational Logic - CL 2000*, number 1861 in *LNAI*, pages 239–253. Springer, 2000.

- [MM00] Raymond McDowell and Dale Miller. Cut-elimination for a logic with definitions and induction. *Theoretical Computer Science*, 232:91–119, 2000.
- [MM02] Raymond McDowell and Dale Miller. Reasoning with higher-order abstract syntax in a logical framework. *ACM Transactions on Computational Logic*, 3(1):80–136, 2002.
- [MN87] Dale Miller and Gopalan Nadathur. A logic programming approach to manipulating formulas and programs. In Seif Haridi, editor, *IEEE Symposium on Logic Programming*, pages 379–388, San Francisco, September 1987.
- [MNPS91] Dale Miller, Gopalan Nadathur, Frank Pfenning, and Andre Scedrov. Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51:125–157, 1991.
- [MT03a] Dale Miller and Alwen Tiu. A proof theory for generic judgments: An extended abstract. In Phokion Kolaitis, editor, *18th Symposium on Logic in Computer Science*, pages 118–127. IEEE, June 2003.
- [MT03b] Alberto Momigliano and Alwen Tiu. Induction and co-induction in sequent calculus. In Mario Coppo, Stefano Berardi, and Ferruccio Damiani, editors, *Post-proceedings of TYPES 2003*, number 3085 in LNCS, pages 293–308, January 2003.
- [MT05] Dale Miller and Alwen Tiu. A proof theory for generic judgments. *ACM Transactions on Computational Logic*, 6(4):749–783, October 2005.
- [Nip93] Tobias Nipkow. Functional unification of higher-order patterns. In M. Vardi, editor, *Proc. 8th IEEE Symposium on Logic in Computer Science (LICS 1993)*, pages 64–74. IEEE, June 1993.
- [NL05] Gopalan Nadathur and Natalie Linnell. Practical higher-order pattern unification with on-the-fly raising. In *ICLP 2005: 21st International Logic Programming Conference*, volume 3668 of LNCS, pages 371–386, Sitges, Spain, October 2005. Springer.
- [NM88] Gopalan Nadathur and Dale Miller. An overview of λ Prolog. In *Fifth International Logic Programming Conference*, pages 810–827, Seattle, August 1988. MIT Press.
- [PE88] Frank Pfenning and Conal Elliott. Higher-order abstract syntax. In *Proceedings of the ACM-SIGPLAN Conference on Programming Language Design and Implementation*, pages 199–208. ACM Press, June 1988.
- [Pit03] Andrew M. Pitts. Nominal logic, A first order theory of names and binding. *Information and Computation*, 186(2):165–193, 2003.

- [Plo81] Gordon D. Plotkin. A structural approach to operational semantics. DAIMI FN-19, Aarhus University, Aarhus, Denmark, September 1981.
- [Pos08] Adam Poswolsky. *Functional Programming with Logical Frameworks*. PhD thesis, Yale University, December 2008.
- [PS99] Frank Pfenning and Carsten Schürmann. System description: Twelf — A meta-logical framework for deductive systems. In H. Ganzinger, editor, *16th Conference on Automated Deduction (CADE)*, number 1632 in LNAI, pages 202–206, Trento, 1999. Springer.
- [Qi09] Xiaochu Qi. *An Implementation of the Language λ Prolog*. PhD thesis, University of Minnesota, September 2009. Forthcoming.
- [SH93] Peter Schroeder-Heister. Rules of definitional reflection. In M. Vardi, editor, *Eighth Annual Symposium on Logic in Computer Science*, pages 222–232. IEEE Computer Society Press, IEEE, June 1993.
- [Tai67] William W. Tait. Intensional interpretations of functionals of finite type I. *Journal of Symbolic Logic*, 32(2):198–212, 1967.
- [Tiu04] Alwen Tiu. *A Logical Framework for Reasoning about Logical Specifications*. PhD thesis, Pennsylvania State University, May 2004.
- [Tiu06] Alwen Tiu. A logic for reasoning about generic judgments. In A. Momigliano and B. Pientka, editors, *International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP'06)*, 2006.
- [Tiu08] Alwen Tiu. Cut elimination for a logic with generic judgments and induction. Technical report, CoRR, January 2008. Extended version of LFMTP'06 paper. Available from <http://arxiv.org/abs/0801.3065>.
- [TM08] Alwen Tiu and Dale Miller. Proof search specifications of bisimulation and modal logic for the π -calculus. Submitted, May 2008.
- [TM09] Alwen Tiu and Alberto Momigliano. Induction and co-induction in sequent calculus. Available from <http://arxiv.org/abs/0812.4727>, 2009.
- [UT05] Christian Urban and Christine Tasson. Nominal techniques in Isabelle/HOL. In R. Nieuwenhuis, editor, *20th Conference on Automated Deduction (CADE)*, volume 3632 of *LNCS*, pages 38–53. Springer, 2005.
- [Van96] Myra VanInwegen. *The Machine-Assisted Proof of Programming Language Properties*. PhD thesis, University of Pennsylvania, May 1996.