



Delft University of Technology

Document Version

Final published version

Citation (APA)

van der Rest, C. R. (2026). *Reusable Programming Language Components*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:48c18134-87fd-4d1b-b609-2a9d6e242cc2>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

reusable programming language components

Cas van der Rest

REUSABLE PROGRAMMING LANGUAGE COMPONENTS

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft
op gezag van de Rector Magnificus prof.dr.ir. H. Bijl
voorzitter van het college voor promoties
in het openbaar te verdedigen op dinsdag 21 april 2026 om 17:30 uur

DOOR

Casper Roland van der REST

Dit proefschrift is goedgekeurd door de promotoren en copromotor.

SAMENSTELLING PROMOTIECOMMISSIE BESTAAT UIT:

| | |
|-------------------------|---|
| Rector Magnificus | <i>voorzitter</i> |
| Prof.dr. A. van Deursen | Technische Universiteit Delft, <i>promotor</i> |
| Dr. N. Yorke-Smith | Technische Universiteit Delft, <i>promotor</i> |
| Dr. C. Bach | Syddansk Universitet, Odense, Denemarken, en Technische Universiteit Delft, <i>copromotor</i> |

ONAFHANKELIJKE LEDEN:

| | |
|---------------------------|---|
| Prof.dr. K. G. Langendoen | Technische Universiteit Delft |
| Prof.dr. G. K. Keller | Universiteit Utrecht |
| Prof.dr.ir. T. Schrijvers | Katholieke Universiteit Leuven, Leuven, België |
| Prof.dr. A.P. Tolmach | Portland State University, Portland, Verenigde Staten |
| Prof.dr. M.M. de Weerd | Technische Universiteit Delft, <i>reservelid</i> |

Copyright © 2026 Cas van der Rest

Ontwerp omslag: Theresa Schokkenbroek

ISBN: 978-94-6384-952-4

Uitgever: Linkt Media | linktmedia.nl

Het ontwerp van dit proefschrift is gebaseerd op de 'Classic Thesis'-stijl van André Miede. De typografie van de omslag en de hoofdstuknummering is gebaseerd op een digitale interpretatie door Scott Watson van het 'Stedelijk'-lettertype van Wim Crouwel. De liedtekst op de volgende pagina komt uit het nummer *Love You More* van Racon.



Het werk in dit proefschrift is uitgevoerd aan de Technische Universiteit Delft onder auspiciën van de onderzoeksschool IPA (Instituut voor Programmatuurkunde en Algoritmiek). Het onderzoek in dit proefschrift is gedeeltelijk gefinancierd door het NWO VENI-project "Composable and Safe-by-Construction Programming Language Definitions" (VI.Veni.192.259).

*Clouds above go sailing by
I found my meaning in this life
Clear white is flying in my eyes
Underneath a blue, blue sky
The waves come rolling in
With the tide*

*All the people rushing by, by, by
Looking for meaning in this life
So used up, and blinded by lies
They're underneath the blue, blue sky
The way they seldom seem to smile
I don't know why*

ACKNOWLEDGEMENTS

I didn't get here alone, and I want to take a brief moment to reflect on the people who supported me.

Thank you to my supervisor, Casper Bach, who founded the line of research that ultimately led to the position for which I was hired. I am proud of the work that we did together and of how we learned to collaborate effectively and in good spirit, even under high pressure. Thank you for your mentorship, support, and the freedom to explore my own ideas

Thank you to my original promotor, Eelco Visser, who provided invaluable guidance to the project until his untimely passing in the spring of 2022. Neil Yorke-Smith and Arie van Deursen stepped up to act as my promotors, and helped tremendously in improving the manuscript and navigating bureaucratic procedures. I am grateful to the other committee members, Gabrielle Keller, Andrew Tolmach, Tom Schrijvers, and Koen Langendoen, who took the time to read, evaluate, and provide feedback on an earlier version of this manuscript. And, thank you to the various members of the support staff who helped me with countless administrative issues. I also want to mention Wouter Swierstra, who encouraged me to pursue a PhD, and Peter Mosses, who has a wealth of knowledge about (the history of) PL research and gave me advice on numerous occasions.

Thank you, Kobe Wullaert, Daniël Pelsmaeker, Lucas Escot, Dennis Sprokholt, and Dana Liesnikov, for your friendship, both in and outside of work. Let's try to find ways to keep climbing, playing boardgames, and having drinks together. I wish you the best, and I look forward to crashing your defenses. Thank you, Jeff Smits and Aron Zwaan, for our adventures in Hobbiton and elsewhere. Thank you, Arjen Rouvoet and Jules Jacobs, for being my first office mates and showing me an incredible standard to aspire to. Thank you to them and to the other past

and present members of the PL group. I think fondly of the time I spent with you and the amazing places we got to explore together—Oregon, Ljubljana, Auckland, Valencia, and more. It's you that made my time at TU Delft memorable and worthwhile.

Thank you to my kind and wise family. My parents, Theresa Schokkenbroek and Roel van der Rest, are there for me no matter what. I would never have been where I am now without the innumerable ways in which you supported me. My sister, Iris, has on more occasions than I can remember offered me either kind words or a stern talking-to, whichever she felt I needed most. Thank you, Rein and Peter, and all my other family in the Netherlands, for being there

Thank you, Sibren, for your good spirit and putting up with my communication habits; Gijs, for nerding out over ice skating and somehow getting my table upstairs; and Matthijs, for history lessons and musical escapades. Thank you, Sára, for organizing everything and letting me borrow all your stuff, and Jelco, for darts, beers, and being good company on and off the bike. Thank you, Max, for your level-headedness and advice, and Luka, for dinner at X and being a fine GM. Thank you, to all my friends at Arktika, for good times in Utrecht, Collalbo, Inzell, Ardennes, Baselga, and other places. I cherish our shared memories, and your friendship has been essential to my sanity over the past six years.

Thank you to my significant others. Your love and support meant a great deal to me and contributed immensely to this journey.

CONTENTS

| | | |
|-------|---|----|
| 1 | Introduction | 1 |
| 1.1 | The Problem: a Lack of Types and Type Safety | 2 |
| 1.2 | Solution Direction and Thesis Statements | 6 |
| 1.3 | Research Methods in Programming Languages | 8 |
| 1.4 | Approach and Thesis Structure | 9 |
| 1.5 | Origin of the Chapters | 13 |
| 1.6 | A Note on Artifacts | 14 |
| | | |
| 1 | Modular Semantics in Agda | |
| 2 | Intrinsically-Typed Definitional Interpreters à la Carte | 19 |
| 2.1 | Introduction | 20 |
| 2.1.1 | Background: Intrinsically-Typed Interpreters | 22 |
| 2.1.2 | Challenge: Intrinsically-Typed Programming Language Fragments | 24 |
| 2.1.3 | Contributions | 25 |
| 2.2 | Data Types à la Carte | 26 |
| 2.2.1 | Composing Data Types | 27 |
| 2.2.2 | Composing Functions | 31 |
| 2.2.3 | Discussion | 33 |
| 2.3 | Indexed Data Types à la Carte, for Defining Composable Intrinsically-Typed Interpreters | 34 |
| 2.3.1 | Composing Index Types | 35 |
| 2.3.2 | Composing Intrinsically-Typed Values | 35 |
| 2.3.3 | Composing Intrinsically-Typed Expressions | 36 |
| 2.3.4 | Composing Index-Preserving Functions | 39 |
| 2.3.5 | Discussion | 42 |
| 2.4 | Intrinsically-Typed Language Fragments | 44 |
| 2.4.1 | Canons and Language Fragments | 44 |
| 2.4.2 | Fragment Composition and the Need for Partially-Overlapping Canons | 46 |
| 2.4.3 | Fragment Composition with Partially-Overlapping Canons | 47 |

| | | |
|-------|---|----|
| 2.5 | Language Fragments with Lexical Variables and Effects | 51 |
| 2.5.1 | Fragments for a Class of Semantic Domains | 52 |
| 2.5.2 | Simply-Typed Lambda Calculus | 56 |
| 2.5.3 | Exceptions | 58 |
| 2.5.4 | ML-Style References | 59 |
| 2.5.5 | Case Study | 61 |
| 2.5.6 | Discussion | 63 |
| 2.6 | Related Work | 64 |
| 2.6.1 | Meta-Theory à la Carte. | 64 |
| 2.6.2 | Generic Programming and Meta-Theory | 65 |
| 2.6.3 | Other Approaches to Modular Semantics and their Proofs | 67 |
| 2.7 | Conclusion | 68 |
| 3 | Hefty Algebras: Modular Elaborations and Reasoning for Programs with Higher-Order Effects | 71 |
| 3.1 | Introduction | 72 |
| 3.1.1 | Background: Algebraic Effects and Handlers | 72 |
| 3.1.2 | The Modularity Problem with Higher-Order Operations | 74 |
| 3.1.3 | Solving the Modularity Problem: Elaboration Algebras | 78 |
| 3.1.4 | Contributions | 80 |
| 3.2 | Algebraic Effects and Handlers in Agda | 81 |
| 3.2.1 | Algebraic Effects and The Free Monad | 82 |
| 3.2.2 | Row Insertions and Smart Constructors | 84 |
| 3.2.3 | Fold and Monadic Bind for Free | 86 |
| 3.2.4 | Effect Handlers | 87 |
| 3.2.5 | The Modularity Problem with Higher-Order Effects, Revisited | 89 |
| 3.2.6 | Scoped Effects and Handlers | 91 |
| 3.3 | Hefty Trees and Algebras | 96 |
| 3.3.1 | Generalizing Free to Support Higher-Order Operations | 97 |

| | | |
|-------|---|-----|
| 3.3.2 | Programs with Algebraic and Higher-Order Effects | 101 |
| 3.3.3 | Higher-Order Operations with Polymorphic Return Types | 102 |
| 3.3.4 | Hefty Algebras | 103 |
| 3.3.5 | Discussion and Limitations | 106 |
| 3.4 | Examples | 108 |
| 3.4.1 | λ as a Higher-Order Operation | 108 |
| 3.4.2 | Optionally Transactional Exception Catching | 112 |
| 3.4.3 | Logic Programming | 115 |
| 3.4.4 | Concurrency | 116 |
| 3.5 | Modular Reasoning for Higher-Order Effects | 119 |
| 3.5.1 | Theories of Algebraic Effects | 121 |
| 3.5.2 | Modal Necessity | 122 |
| 3.5.3 | Effect Theories | 124 |
| 3.5.4 | Syntactic Equivalence of Effectful Programs | 127 |
| 3.5.5 | Handler Correctness | 130 |
| 3.5.6 | Theories of Higher-Order Effects | 131 |
| 3.5.7 | Equivalence of Programs with Higher-Order Effects | 134 |
| 3.5.8 | Correctness of Elaborations | 136 |
| 3.5.9 | Proving Correctness of Elaborations | 138 |
| 3.6 | Related Work | 140 |
| 3.7 | Conclusion | 143 |

II Meta Language Design

| | | |
|-------|--|-----|
| 4 | Towards a Language for Defining Reusable Programming Language Components | 147 |
| 4.1 | Introduction | 147 |
| 4.2 | CS by Example | 151 |
| 4.2.1 | Data Types and Functions | 151 |
| 4.2.2 | Effects and Handlers | 152 |
| 4.2.3 | Order of Evaluation, Suspension, and Enactment | 154 |
| 4.2.4 | Modules and Imports | 156 |

| | | |
|-------|--|-----|
| 4.2.5 | Composable Data Types and Functions | 157 |
| 4.3 | Defining Reusable Language Components in CS | 160 |
| 4.3.1 | A Signature for Reusable Components | 160 |
| 4.3.2 | A Language Component for Arithmetic Expressions | 161 |
| 4.3.3 | Implementing Functions as a Reusable Effect | 162 |
| 4.3.4 | Example Usage | 167 |
| 4.4 | Related Work | 167 |
| 4.4.1 | Effect Semantics | 167 |
| 4.4.2 | Implementations of Algebraic Effects and Handlers | 169 |
| 4.4.3 | Semantics of Composable Data Types and Functions | 170 |
| 4.4.4 | Row Types | 171 |
| 4.5 | Future Work | 172 |
| 4.6 | Conclusion | 173 |
| 5 | Types and Semantics for Extensible Data Types | 175 |
| 5.1 | Introduction | 175 |
| 5.1.1 | Contributions | 177 |
| 5.2 | Programming with Extensible Data Types, by Example | 179 |
| 5.2.1 | Notation | 180 |
| 5.2.2 | Modular Interpreters in the style of Data Types à la Carte | 180 |
| 5.2.3 | Modular Algebraic Effects using the Free Monad | 181 |
| 5.2.4 | Modular Higher-Order Effects | 182 |
| 5.3 | The Calculus | 184 |
| 5.3.1 | Well-Formed Types | 186 |
| 5.3.2 | Well-Typed Terms | 188 |
| 5.3.3 | Type Equivalence | 189 |
| 5.4 | Categorical Semantics | 191 |
| 5.4.1 | Interpreting Kinds and Kind Environments | 192 |
| 5.4.2 | Interpreting Types | 193 |
| 5.4.3 | On the Existence of Initial Algebras | 199 |

| | | |
|-------|-------------------------------------|-----|
| 5.4.4 | Arrow Types Correspond to Morphisms | 200 |
| 5.4.5 | Interpreting Terms | 202 |
| 5.5 | Operational Semantics | 204 |
| 5.5.1 | Reduction Rules | 204 |
| 5.5.2 | Relation to the Denotational Model | 206 |
| 5.6 | Related Work | 208 |
| 5.7 | Conclusion and Future work | 210 |

Conclusions

| | | |
|-------|--|-----|
| 6 | Conclusions | 215 |
| 6.1 | Summary of the Contributions | 215 |
| 6.2 | Hypothesis 1: intrinsically-typed interpreters | 216 |
| 6.3 | Hypothesis 2: Meta Language Design | 218 |
| 6.4 | Future Work | 219 |
| 6.4.1 | Modular semantics | 220 |
| 6.4.2 | Meta Language Design | 221 |
| 6.4.3 | Connecting the Dots | 221 |

Appendix

| | |
|--------------|-----|
| Bibliography | 225 |
| Summary | 255 |
| Samenvatting | 259 |



INTRODUCTION

Sometimes a computer does not do what it is supposed to do. The colloquial term for a computer mishap is *bug*, named after a famous incident in 1947 involving the Harvard Mark II computer, where operators discovered a moth that had worked its way into the internals of the machine causing it to malfunction. Although the moth in question is unlikely to cause any more trouble—it is kept safely in the National Museum of American History in Washington D.C.—the term “bug” has stuck around, and is used to refer to any unintended or unexpected behavior of a computer or program. Nowadays, with the ever-increasing pervasiveness of computers in almost all aspects of our daily lives, the number of opportunities for bugs to inflict damage is growing as well. Indeed, the economical and societal consequences of software malfunctions are potentially enormous.¹

Bugs can have a wide variety of causes, but often they arise as a result of what we might classify as miscommunication between the programmer and the machine. This class of bugs is fittingly summarized by the acronym *PEBCAK*: Problem Exists Between Chair And Keyboard. The programmer assumes they have produced a piece of code that does one thing, but when the program is executed by the machine something else happens. The program could produce a wrong result, behave unexpectedly, or even crash altogether. At first glance it is tempting to ascribe such errors to a lack of skill or training on behalf of the programmer. After all, a computer does nothing apart from diligently executing the instructions we give it. This is, however,

¹ For example, in 2022 CISQ estimated the “cost of poor software quality” in the US to be 2.41 Trillion US dollar.

only part of the story. The tool used by the programmer to communicate with the machine, the *programming language*, plays an important role too. It shapes the way programmers think about a problem, and sets the boundaries for what constitutes a valid program text.

There are vast differences across languages when it comes to how well they shield programmers from their own mistakes, and how effective they allow them to communicate with the machine. Undeniably, some programming languages make it easier for programmers to “shoot themselves in the foot”. Over time programming languages have become more effective at protecting programmers from their own mistakes. Where the first computers were programmed by punching holes into a paper card, with careful examination of the punched card being the only safeguard against bugs, modern languages provide much more support for catching mistakes during the development process. Languages may do this by abstracting away from details of the underlying machine that are otherwise very error-prone to work with, such as memory management or concurrency, or performing sophisticated compile-time checks to eliminate certain classes of mistakes altogether. Despite these advances, bugs remain prevalent and there remains much to be gained from further improving programming language design.

1.1 THE PROBLEM: A LACK OF TYPES AND TYPE SAFETY

One of the most important techniques in programming language design for preventing bugs are *type systems*. [Pierce \[2002\]](#) coins the following definition:

A type system is a tractable syntactic method for proving the absence of certain program behaviors by classifying phrases according to the kinds of values they compute.

This quote describes a somewhat classical perspective on type systems, where each (sub)term is assigned a *type* that classifies

the kind of value that it evaluates to. These kinds of type systems rule out errors resulting from a mismatch in the *type of data*. For example, if we write a program that feeds a piece of text to a function that is meant to take a number as input, the *type checker*² should catch the mistake before the program is ever executed.

Nowadays, type systems have the potential to statically enforce much more sophisticated properties beyond tracking the kind of values computed by a term. For example, *type-and-effect systems* keep track of the side effects of evaluating a term [Nielson and Nielson, 1999], *quantitative type theory* monitors the number of times a variable is consumed [Atkey, 2018], and *session types* control the interaction of processes in a distributed setting [Honda et al., 1998].

IT IS IMPORTANT to realize that a type system merely defines a partition of the set of all program texts into two groups based on whether a program is well-typed according to the system or not. In practice, this judgment is only of interest to us if we can relate it to what happens when programs are executed. That is, if a program is judged to be well-typed by the system, this should imply that execution of the program is indeed devoid of certain kinds of erroneous behavior. This is an important meta-property of type systems, aptly summarized by Milner [1978] in his famous quote:

Well-typed programs cannot go wrong.

If a type system successfully rules out all programs that exhibit a particular kind of wrong behavior,³ it is said to be *type sound*. Dually, we could say that a type system is *complete* if it designates all programs that do not exhibit wrong behavior as well-typed. However, as a consequence of *Rice's theorem*,⁴ it is impossible for a type system to be both sound and complete. Since their main purpose is to prevent software errors, type systems are typically designed to be sound, giving up completeness.

Evidently, type soundness is a key property of a type system to be a suitable tool for helping programmers write correct pro-

² A type checker is a computer program that implements a type system. It is a decision procedure that determines whether a program is well-formed according to the rules of the system.

³ The precise meaning of “wrong” here depends on the kind of errors the type system is intended to prevent.

⁴ Rice's theorem states that any non-trivial semantic property of a program is undecidable: no systematic procedure can exist that decides whether such a property holds or not.

grams. How can we establish that a type system is sound? The most airtight method is to mathematically verify type soundness, for example using the syntactic approach of [Wright and Felleisen \[1994\]](#). However, this is rarely done outside the context of academic research since producing the required proofs is a complex and incredibly time-consuming task. As a result, trust in the judgment of type checkers for mainstream programming languages is typically based on empirical instead of mathematical evidence.

Moreover, the absence of wrong behavior in well-typed programs is limited to those mistakes that the type system was designed to rule out. Type systems of mainstream programming languages generally model only limited aspects of a program's execution, ignoring, for example, interaction with the file system or a shared piece of memory. Still, these may very well cause erroneous behavior, so by not modelling them in the type system bugs may arise despite the type checker's blessing.

This leaves us with the following two ways in which (implementations of) mainstream programming languages are not helping programmers prevent bugs as much as they could:

- languages and their type system tend not to be formally specified, let alone verified to be sound. This diminishes the extent to which we can trust a type checker to deliver a meaningful judgment, and
- type checkers tend to model only limited aspects of a program's execution, usually the type of values computed by a term. This leaves the programmer alone in catching errors that arise from any other aspect of a program's execution.

These shortcomings do not exist because the state-of-the-art in programming language theory does not offer us ways to address them. Still, if we look at the difference between what is possible at the cutting edge of programming language design and verification, and the programming languages that people actually use, we find a gap spanning, in extreme cases, several

decades. To illustrate: a form of *pattern matching* was added to both Java and Python⁵ in 2021, while Haskell, which is by no means the source of the technique, included it as of its original specification released back in 1992 [Hudak et al., 1992]. This begs the question: why is there such a large gap between the cutting edge in programming language research and daily practice of programming language design?

⁵ More details on pattern matching in Java and Python can be found in the release notes for [Java 16](#) and [Python 3.10](#), respectively.

THE SIMPLE, YET unsatisfying, explanation is that it is a matter of accessibility. The development of new programming languages requires a tremendous amount of time, knowledge, and effort on behalf of the language designers. The already significant cost involved with developing a programming language is multiplied when we include a formal specification of the type system and proofs of type soundness. This makes it unrealistic for most, if not all, real-world language projects to benefit from the full potential of what the state-of-the-art in programming language technology has to offer: the associated cost is simply too high. The few examples of the formal specification and verification of programming languages and their implementations at scale that do exist, such as CompCert [Leroy, 2009] or CakeML [Kumar et al., 2014], are large and multi-year collaborative efforts by experts in the field of programming languages. Considering the vast amount of time and expertise required for such projects, they present an impossible standard for the vast majority of language projects.

We summarize the situation as follows:

The development of formally specified and verified programming languages is extremely expensive, and therefore an unrealistic standard for real-world language projects. As a result, there exists a gap between the level of safety that the cutting edge in programming language design has to offer, and what is provided by mainstream programming languages.

This thesis works towards closing this gap, by developing tools and techniques for reducing the cost of developing formally

specified and verified programming languages. Ultimately, the goal is to bring the highest standard of correctness and safety offered by the state-of-the-art in programming language research to the languages that programmers use to write software every day.

1.2 SOLUTION DIRECTION AND THESIS STATEMENTS

The main strategy we adopt in this thesis for reducing the costs associated with the development of formally specified and verified programming languages is *component reuse*. That is to say: the development cost of formally verified programming languages would be greatly reduced if we could benefit from previous verification efforts by grabbing off-the-shelf components for those features that are shared with other languages. This would allow us to direct more resources towards specifying and verifying the novel features of a language project. Indeed, the POPLMark challenge [Aydemir et al., 2005], which provides a benchmark for measuring how well state-of-the-art technology in formal verification is up to the task of formalizing programming language metatheory, identifies component reuse as one of the key issues inhibiting a more widespread adoption of formal verification techniques in the design and development of new programming languages.

In a perfect world, we would specify and verify programming languages on a per-feature basis. The definition of each language construct should be such that it is a mathematically rigorous and executable specification of that construct's *semantics*,⁶ but simultaneously concise and perspicuous enough that it can serve as documentation to readers with a background in computer science. Furthermore, we should be able to reuse these specifications in the context of a different language project without having to alter the original specification. This combination of qualities will allow for rapid development of and experimentation with new designs of formally specified and verified programming languages.

⁶ A language's semantics gives a mathematical description of how a program behaves when executed.

Before we can achieve modular, formal, and readable specifications of programming language components, we identify two fronts on which we need to make progress. The first challenge is to formally specify and verify programming language components in a way that the specification remains independent of the context of the language in which they are used, while remaining concise and readable. Especially when we require a type-safety proof to be part of this specification, and when we want to modularly define a component's side effects. Second, we recognize that the functional and/or dependently-typed *meta languages*⁷ used to define and verify programming languages offer little to no support for modular definitions.⁸ As a result, once we set out to define reusable programming language components in these meta languages, we find that modularity adds significant syntactic and interpretive overhead to our definitions.

⁷ Examples of this would be languages like Haskell or Agda.

⁸ Wadler [1998] coined the term expression problem for this lack of modularity in (functional) languages.

To MAKE THINGS concrete, this thesis explores the following two hypotheses:

Hypothesis 1: reusable programming language components should be concise, readable, and safe-by-construction. Intrinsically-typed definitional interpreters are an excellent match for these requirements.

Hypothesis 2: modularity adds significant syntactic and interpretative overhead when using state-of-the-art (dependently-typed) programming languages to define reusable programming language components. Incorporating modular inductive data types in the design of functional languages is therefore an essential first step in the development of meta languages for the purpose of defining reusable programming language components.

It is important to be mindful of the limitations of intrinsically-typed interpreters as a means to specify a language's semantics. Recent work by Bach et al. [2018], Rouvoet et al. [2020] showed

that the technique is suitable for defining the semantics of imperative languages and languages with session types, but the technique has yet to be applied to more advanced type systems. While mechanizations of both System F [Chapman et al., 2019] and dependent type theory [Altenkirch and Kaposi, 2016] exist, it is not clear how their semantics could be expressed as intrinsically-typed interpreters in the style of Augustsson and Carlsson [1999]. This thesis is concerned with defining reusable programming language components for those languages to which can already be described using intrinsically-typed definitional interpreters. Before discussing in more detail how this thesis works towards this goal, we must first say a few words about how research on programming languages, including the research presented in this thesis, tends to be conducted.

1.3 RESEARCH METHODS IN PROGRAMMING LANGUAGES

⁹ Here, “idea” can be any sort of contribution, be it a novel language design, a type-checking algorithm, techniques for specifying a language’s semantics, or something else.

Typically, papers introduce a new idea or concept⁹ somewhat informally, using (code) examples and appealing to the reader’s intuition. It is usually only in the later sections that papers develop a more mathematically rigorous description of their contributions, using mathematical notation and proofs. This “top-down” approach, which focusses on fostering a high-level understanding of a paper’s contents, is somewhat necessary in many cases due to a general lack of quantifiable measures for evaluating research. For instance, how would one quantify the usefulness of a new programming feature or proof technique? Rather than relying on the value of contributions to be self-evident, some responsibility is shifted to researchers to explain why their contributions are relevant to the community.

When it comes to mathematical proofs, the programming languages community distinguishes *pen-and-paper proofs* from *mechanized proofs*. The former documents a mathematical argument using a combination of natural language and mathematical notation, and is intended to be read by humans. The latter embeds the entire argument in a chosen formalism, and is intended

to be consumed by a computer that verifies correctness of the argument.

Conferences may offer authors to submit an *artifact* in conjunction with their paper. The purpose of an artifact is to provide additional evidence for the claimed contributions of a submission, commonly in the form of a prototype implementation or a machine-checked proof. A paper and its artifact can, to some extent, be mixed using *literate programming* [Knuth, 1992]. When using literate programming, a paper intersperses pieces of human language with computer language, and the sources of the paper can be consumed by a compiler or theorem prover to validate that the programs and/or proofs are well-formed.

To evaluate the applicability of an idea design, papers may resort to small *case studies*, where ideas are applied to more involved examples compared to the minimal examples used to illustrate ideas. Such evaluations are qualitative, and often appeal to informal qualities of a solution, arguing that it is more usable, modular, or elegant. Quantitative measures are sometimes used for evaluation in papers too, if the merit of a piece of work can be judged using more easily quantifiable attributes, such as performance or code size. It is good practice for the sources of prototypes and mechanized proofs to be publicly available, such that other members can inspect and learn from them.

The individual chapters of thesis have been published at various programming language conferences, and thus follow the general practice and methodology of research in the field of programming languages as described above.

1.4 APPROACH AND THESIS STRUCTURE

The key technique for defining programming languages we embrace in this thesis are *definitional interpreters* [Reynolds, 1998]. The gist of the approach is to specify a language's semantics by defining an interpreter, and equating the dynamic semantics of the language to the behavior of this interpreter. In the context of

definitional interpreters, it is important to distinguish the *object language*, the language we are defining a semantics for, and the *meta language* (sometimes also referred to as the *host language*), which is the language in which we define the semantics.

Broadly speaking, there are three main approaches when it comes to formally specifying a programming language’s behavior:

1. *axiomatic semantics* [Hoare, 1969], in which we use logical predicates to make assertions about a program’s effect on the program state,
2. *operational semantics* [Plotkin, 2004], in which we use an inductively-defined transition relation to describe how a program evolves over time, and
3. *denotational semantics* [Scott and Strachey, 1971], in which we map programs onto a mathematical object that models their semantics.

Definitional interpreters are most closely related to denotational semantics, where the “mathematical” object that we use to model programs of the object language is a program in the meta language. Of course, the mathematical rigor of such a specification depends on the meta languages used. A definitional interpreter written in Haskell, for example, cannot be taken as a mathematical specification of a language due to the presence of non-termination, exceptions, and other unsafe language features such as `unsafePerformIO`. When using a language like Agda, on the other hand, the function that maps object language terms to Agda programs must be *total*,¹⁰ hence we can view definitional interpreters written in Agda as a denotational semantics that denotes programs of the object language as terms in Agda’s type theory. We highlight that for the purpose of defining reusable programming language components, there is a balance to be struck between these extremes. While mathematical rigor is crucial, so is readability and usability of the specifications.

When working in a dependently typed language, such as Agda, *intrinsically-typed* definitional interpreters [Augustsson

¹⁰ Here, *total* means that for every possible input value, an output will be computed in finite time.

and Carlsson, 1999] offer a combination of readability and safety. The crux of the technique is to encode a typing invariant in the abstract syntax tree of the object language, such that it is only inhabited by well-typed terms.¹¹ The corresponding interpreter then maps well-typed terms to values of the same type. While the resulting definition looks like a “normal” interpreter, its type is precise enough such that we can omit any code for dealing with ill-formed terms. Due to totality of the function that implements the interpreter, the fact that it exists immediately implies a type-safety theorem about the typing invariant encoded in the abstract syntax. For these reasons intrinsically-typed interpreters are an attractive candidate for specifying and verifying reusable language components.

THIS THESIS IS organized into two parts that respectively support the two hypotheses described in Section 1.2. In part I (Chapters 2 and 3), we work on furthering the collection of techniques available to languages designers for writing type-safe modular language specifications in Agda. While this is possible, doing things modularly does add some overhead to our definitions. In part II (Chapters 4 and 5), we work towards a meta language design that has modularity built in, reducing this overhead while retaining type safety.

In Chapter 2, we discuss how to define modular intrinsically-typed interpreters in Agda. We give a definition of *intrinsically-typed language fragments* that specify the syntax and semantics of a language component in a way that is type-safe by construction, such that language fragments can be composed and reused freely to build larger languages. While modular type safety proofs have been studied before by Delaware et al. [2013a], Keuchel and Schrijvers [2013], and Delaware et al. [2013c], their work focusses on an extrinsic style of proofs, inducing considerable syntactic overhead compared to the intrinsic approach.¹²

In Chapter 3, we discuss how to define modular elaborations of *higher-order effects*, along with modular reasoning principles that allow us to verify that elaborations respect equational theories of the effect(s) they implement. While there is a considerable

¹¹ The converse, more traditional approach, is to use extrinsic typing, where well-typedness is defined separately as a predicate over an untyped abstract syntax type.

¹² The amount of code needed to specify a verified language component differs roughly by one order of magnitude.

amount of existing work on (modular) higher-order effects (for example, by [Yang et al. \[2022\]](#), [Van den Berg et al. \[2021a\]](#), and [Wu et al. \[2014\]](#)), these focus on defining handlers for higher-order effects directly. We take a slightly different approach by defining the semantics of higher-order effects in terms of algebraic effects, effectively following the approach by [Plotkin and Pretnar \[2009b\]](#) of implementing higher-order effects as handlers, but adding an additional syntactic layer to regain modularity.

In Chapter 4, we present the design of a meta language for developing reusable programming language components, that has built-in support for modular data types and effects. A crucial feature of the language is that we can solve the expression problem [[Wadler, 1998](#)] without resorting to embedding a data type’s initial algebra semantics [[Goguen, 1976](#)], reducing the syntactic overhead of modularity. Furthermore, it features a type-and-effect system based on *latent effects* [[Van den Berg et al., 2021a](#)] that support, among other things, λ -abstraction as a user-defined effect.

In Chapter 5 puts the modular data types presented in Chapter 4 on formal footing by presenting a core calculus together with a type system and semantics that features built-in support for initial algebra semantics [[Goguen, 1976](#)]. Many familiar programming abstractions for modularity can be encoded in the calculus, such as modular data types and interpreters in the style of *Data Types á la Carte* [[Swierstra, 2008](#)], algebraic effects and handlers [[Plotkin and Power, 2002](#), [Plotkin and Pretnar, 2009a](#)], as well as various implementations of higher-order effects such as *Scoped Effects* [[Yang et al., 2022](#)] or *Hefty Algebras* [[Bach and Van der Rest, 2023](#)]. While the calculus lacks any syntactical conveniences for working with these constructions, it provides a basis for understanding how modular data types can be added to a language in a principled way.

Finally, in Chapter 6 we will look back on the work presented in these two parts, and reflect on the contributions of the thesis in light of the goals outlined in this introduction. Moreover, we

will discuss future work, and how to connect the contributions presented in the two parts of this thesis.

1.5 ORIGIN OF THE CHAPTERS

The contents of this thesis correspond to previously published work as follows.

CHAPTER 2 is based on:

Cas van der Rest, Casper Bach, Arjen Rouvoet, Eelco Visser and Peter D. Mosses. “Intrinsically-Typed Definitional Interpreters à la Carte. In: *Proceedings of the ACM on Programming Languages* 6.OOPSLA, 2022. DOI: 10.1145/3563355.

CHAPTER 3 is based on:

Casper Bach and **Cas van der Rest**. “Hefty Algebras: Modular Elaboration of Higher-Order Algebraic Effects”. In: *Proceedings of the ACM on Programming Languages* 7.POPL, 2023. DOI: 10.1145/3571255.

and¹³

Cas van der Rest and Casper Bach. “Hefty Algebras: Modular Elaboration of Higher-Order Algebraic Effects”. In: *Journal of Functional Programming* volume 35. DOI: 10.1017/S0956796825100142.

¹³ The journal version extends the original conference paper with techniques for modular reasoning about programs with higher-order effects.

CHAPTER 4 is based on:

Cas van der Rest and Casper Bach. “Towards a Language for Defining Reusable Programming Language Components - (Project Paper). In: *Selected Papers from the 23rd International Symposium in Trends in Functional Programming*, 2022. DOI: 10.1007/978-3-031-21314-4_2.

¹⁴ An extended version of the paper was published on arXiv [Van der Rest and Bach, 2023b]

CHAPTER 5 is based on:¹⁴

Cas van der Rest and Casper Bach. “Types and Semantics for Extensible Data Types”. In: *Proceedings of the 21st Asian Symposium on Programming Languages and Systems*, 2023. DOI: 10.1007/978-981-99-8311-7_3.

The contents of the papers that form Chapters 2 to 5 of this thesis have been included verbatim in their corresponding chapter. At the start and end of each chapter, a (small) unnumbered section may have been added that discusses and reflects on the contents of the paper in the context of this thesis, and/or discussing its relation to the preceding and succeeding chapters.

THE AUTHOR ALSO contributed to the following papers:

Casper Bach, **Cas van der Rest**, and Tom Schrijvers. “Staged Effects and Handlers for Modular Languages with Abstraction”. In: *Workshop on Partial Evaluation and Program Manipulation*, 2021.

Cas van der Rest and Wouter Swierstra. “A Completely Unique Account of Enumeration”. In: *Proceedings of the ACM on Programming Languages* 6.ICFP, 2022. DOI: 10.1145/3547636.

1.6 A NOTE ON ARTIFACTS

Part of the contributions in this thesis have been mechanized in Agda. Specifically, the contributions outlined in part I of this thesis are supported by Agda developments that are publicly available [Van der Rest et al., 2022a, Van der Rest and Bach, 2024]. While the code examples in Part I of this thesis are type-checked, it is not the exact code contained in the accompanying developments. For expositional purposes, definitions are occasionally simplified, and (parts of) proofs and definitions may have been omitted where they would clutter the explanation. The code in the Agda development contains the full definitions.

For the language design presented in Chapter 4, there exists a prototype interpreter written in Haskell [Hudak et al., 1992]

which is publicly available. The calculus design presented in Chapter 5 has not been formalized or implemented yet.

Part I

MODULAR SEMANTICS IN AGDA



INTRINSICALLY-TYPED DEFINITIONAL INTERPRETERS À LA CARTE

PREFACE

As noted in the introduction, intrinsically-typed definitional interpreters appear to be an excellent candidate for specifying reusable programming language components due to their conciseness and readability. Although recent work by , [Bach et al. \[2018\]](#), and [Rouvoet et al. \[2020, 2021\]](#) has pushed the technique far beyond the typical examples of interpreters for simply-typed expression languages, they remain non-modular. There are several hurdles to overcome before we can achieve modularity of intrinsically-typed definitional interpreters, such as how to maintain type-safety in light of extensibility, how ensure that components imposing different requirements on their domain remain compatible, and how to alleviate the syntactic and interpretive overhead incurred by modularity.

This chapter begins our journey towards modular intrinsically-typed definitional interpreters by answering the first question: how do we achieve modularity of intrinsically-typed interpreters while maintaining type safety of the definitions? Since intrinsically-typed interpreters can generally be defined as a *catamorphism* over an (indexed) abstract syntax type, they are amenable to the usual techniques for encoding type-safe modularity for data types in functional languages (Section 2.2). However, in the presence of potential future extensions of the type

syntax, it turns out that maintaining type safety in modular definitions requires some additional care (Section 2.3).

2.1 INTRODUCTION

Type safety is a crucial aspect of designing typed programming languages. According to [Pierce \[2002\]](#), “a type system is a tractable syntactic method for proving the absence of certain program behaviors by classifying phrases according to the kinds of values they compute.” The *type safety property* of a language defines precisely what program behaviors its type system is supposed to rule out. But it is challenging to define one’s type system and interpreter¹⁵ in a way that it satisfies the intended type safety property—or, put differently, that it rules out the “bad” program behaviors it is supposed to. For this reason, programming language researchers often rely on mathematical proofs to verify that a type system and interpreter satisfies the type safety property. However, constructing a type safety proof can be a labor intensive and complex task.

¹⁵ It is possible to define the dynamic semantics of a language in many different ways. In this paper, we assume that the dynamic semantics is given by an interpreter.

Our research objective in this paper is to make it as easy as possible for DSL developers to develop and verify the type safety of typed domain-specific languages (DSLs). We propose two sub-objectives that could address this goal together:

1. Support *reuse* of common programming language components, so that DSL developers can focus on developing type safe DSL components.
2. Make it easy to develop and debug type safe DSL components by automating the task of verifying that a language definition is type safe.

There is previous research that goes towards addressing these sub-objectives *individually*, but no previous research that we are aware of which provides a viable solution to *both* at the same time.

Previous work by [Keuchel and Schrijvers \[2013\]](#) and [Delaware et al. \[2013b,c\]](#) suggests a promising direction for addressing the

first sub-objective, by modularizing *extrinsic* type safety proofs and interpreters. An extrinsic proof is an inductive proof on the structure of the syntax or typing rules of a language. By using the modular extrinsic proof techniques of [Keuchel and Schrijvers \[2013\]](#) and [Delaware et al. \[2013b,c\]](#), domain-specific language designers can compose their interpreter, type system, and type safety proof from pre-proven cases for off-the-shelf components, allowing them to focus on defining and proving the type safety of domain-specific components. However, for the second sub-objective that we gave above, the extrinsic proof style has a number of shortcomings. Most importantly, it is unclear how to automate the task of constructing modular, extrinsic type safety proof cases. We also argue that the extrinsic specification style is not as easy to work with as the alternative *intrinsically-typed* style which we discuss shortly. In particular, interpreters in an extrinsic specification style contain redundant cases for bad behavior that can never happen in a type safe language because the type system rules it out. Furthermore, understanding *when and why* an extrinsic type safety proof case does not hold, is key to finding and fixing type safety errors. But it requires previous experience with inductive proofs to identify what the error is. Domain-specific language designers, however, usually do not have the necessary experience to verify type safety.

A more concise and declarative style of verifying type safety is to write an *intrinsically-typed interpreter* [[Augustsson and Carlsson, 1999](#)] in a *dependently-typed host language*. Such interpreters save language designers from having to read and write redundant cases for ill-typed expressions, as the host language type checker can automatically verify that these cases are unreachable in practice. This results in shorter, more declarative specifications that are *safe-by-construction*, in the sense that the type safety of the interpreter follows from the well-typedness of its definition. We do not have to establish type safety in a separate proof: the interpreter *is* the type safety proof. [Wadler et al. \[2020\]](#) observe that, in Agda [[Norell, 2009](#)], “extrinsically-typed terms

¹⁶ A pun referring to the ratio in code size between intrinsically and extrinsically typed, which approximates the golden ratio.

require about 1.6 times as much code as intrinsically-typed”, leading them to suggest that “intrinsic typing is golden”.¹⁶ Another appeal of intrinsically-typed interpreters is that language designers can debug type safety issues by using compiler errors produced by the dependently-typed host language as a guide. While the quality of compile time errors depends on the host language, it does not require previous experience with inductive proof techniques. This makes intrinsically-typed interpreters an attractive approach to developing and debugging type safe languages, because it reduces the amount of work on behalf of DSL developers by taking extrinsic proof obligations for type safety and making them intrinsic to interpreter well-typing. However, intrinsically-typed interpreters fail to address our first sub-objective, since they do not in general support reuse.

In this paper we adapt and combine techniques for modular meta-theory with the intrinsically-typed approach, and develop a new notion of *intrinsically-typed language fragments* and *language fragment composition* that makes it possible to reuse off-the-shelf pre-verified components. Intrinsically-typed language fragments are as concise and declarative and similarly easy to develop and debug as plain, monolithic intrinsically-typed interpreters. Unlike monolithic interpreters, language fragments can be developed and checked in isolation and combined with other fragments to compose type safe languages from reusable components.

2.1.1 Background: Intrinsically-Typed Interpreters

Figure 1 shows an intrinsically-typed interpreter for a simple language with arithmetic (`lit`, `add`) and Boolean (`tt`, `ff`, `ite`) expressions, implemented in Agda [Norell, 2009]. It consists of:

1. a data type of object language types, `Ty : Set`;
2. a function that maps each object language type to a type, `Val : Ty → Set`;

```

data Ty : Set where
  nat  : Ty
  bool : Ty

data Expr : Ty → Set where
  lit  : ℕ → Expr nat
  add  : Expr nat → Expr nat → Expr nat
  tt   : Expr bool
  ff   : Expr bool
  ite  : Expr bool → Expr t →
        Expr t → Expr t

Val : Ty → Set
Val nat = ℕ
Val bool = Bool

interp : Expr t → Val t
interp (lit n)      = n
interp (add e1 e2) = interp e1 + interp e2
interp tt          = true
interp ff          = false
interp (ite e e1 e2)
  = if interp e then interp e1
    else interp e2

```

3. an indexed data type representing the well-typed object language expressions, $\text{Expr} : \text{Ty} \rightarrow \text{Set}$;
4. an *index-preserving evaluation function* that embeds a type safety theorem by mapping well-typed expression to a value of the same type, $\text{interp} : \text{Expr } t \rightarrow \text{Val } t$ (i.e., “well-typed expressions cannot go wrong”, Milner [1978]).

The key ingredient that allows Agda to verify that the interpreter in Figure 1 is type safe is *dependent pattern matching* [Cockx, 2017, Coquand, 1992], enabling Agda to infer a precise type for each variable bound by a pattern match clause. For instance, in the clause $\text{interp} (\text{add } e_1 e_2) = \dots$, Agda infers that e_1 and e_2 have type Expr nat in the right hand side of the definition. Indeed, if e_1 or e_2 had any other type, the pattern match would be ill-formed according to the definition of Expr . Since e_1 and $e_2 : \text{Expr nat}$, Agda can deduce that the recursive calls in $\text{interp } e_1 + \text{interp } e_2$ must yield natural numbers, as the return type of the calls, Val nat , normalizes to \mathbb{N} . Thus no error handling of type mismatching is needed: thanks to Agda’s type checker we know that *this will not happen*, and we do not have to spell out any redundant cases for going wrong.

Figure 1: An intrinsically-typed interpreter for a small expression language

| | |
|---|--|
| <pre> data Ty : Set where {- ...1 -} nat : Ty Val : Ty → Set {- ...2 -} Val nat = ℕ data Expr : Ty → Set where {- ...3 -} lit : ℕ → Expr nat add : Expr nat → Expr nat → Expr nat interp : Expr t → Val t {- ...4 -} interp (lit n) = n interp (add e₁ e₂) = interp e₁ + interp e₂ </pre> | <pre> data Ty : Set where {- ...5 -} bool : Ty Val : Ty → Set {- ...6 -} Val bool = Bool data Expr : Ty → Set where {- ...7 -} tt ff : Expr bool ite : Expr bool → Expr t → Expr t → Expr t interp : Expr t → Val t {- ...8 -} interp tt = true interp ff = false interp (ite e e₁ e₂) = if interp e then interp e₁ else interp e₂ </pre> |
|---|--|

Figure 2: Two intrinsically-typed interpreters for arithmetic expressions (left) and Boolean expressions (right).

2.1.2 Challenge: Intrinsically-Typed Programming Language Fragments

The interpreter in Figure 1 mixes arithmetic and Boolean expressions. If we want to extend or reuse (parts of) this language, we have no option but to modify or copy-paste existing code. A better approach is to assemble interpreters from reusable components. Figure 2 provides an informal illustration of how we might define and check such reusable components in isolation, and compose them with other fragments, to incrementally develop a verified and type safe interpreter. We can compose such fragments by *concatenating* their constructor declarations and function clauses. The Agda comments `{- ...i -}` indicate program points where new constructors or clauses will be inserted during composition. For example, we recover the language from

Figure 1 by inserting the constructors of `Ty` (right) at `{- ...1 -}`, the clauses of `Val` (right) at `{- ...2 -}`, and so forth. Throughout this paper we develop a semantics of intrinsically-typed language fragments that supports this kind of composition, without having to re-type-check existing language fragments.

The challenge with defining fragment composition is that not all extensions are well behaved. In particular, extensions that change the *canonical forms* of a type (i.e., removing or adding a new value constructor for an existing type) are problematic. For example, if we add `Val bool = Maybe Bool` at the `{- ...6 -}` position in Figure 2, then the `interp (ite e e1 e2)` case on the right becomes ill-typed. If we were to re-type-check the composed definition, Agda would correctly reject this extended interpreter for being ill-typed. Such re-checking is, however, contrary to the goal of reusing pre-verified components. This raises the question: under which conditions is language fragment composition guaranteed to be well-typed?

This paper answers this question by introducing a subtyping relation for witnessing that canonical forms are preserved when values are extended. That is, the set of constructors for a value of a given type never changes. By using this subtyping relation in our definition of language fragments and language fragment composition, we automatically rule out bad extensions, such as the `Val bool = Maybe Bool` extension discussed above.

2.1.3 Contributions

Working in Agda¹⁷, we make the following technical contributions:

- We extend (in Section 2.3) the techniques (which we recall in Section 2.2) from *data types à la carte* [Swierstra, 2008] to intrinsically-typed interpreters, and introduce a subtyping relation for canonical forms. This relation allows us to define cases of modular intrinsically-typed interpreters in a way that supports type safe composition.

¹⁷ The code in his paper is available in the accompanying artifact [Van der Rest et al., 2022a].

- We introduce (in Section 2.4) *intrinsically-typed language fragments*, which bundle the syntax and semantics of one or more language constructs. Using a general type union, we define *canonical form unions* to support composition of language fragments with overlapping values.
- We generalize (in Section 2.5) our framework from Agda’s `Set` to a broader class of semantic domains. With this generalization, we can modularly define languages with effects such as name binding, exceptions, and mutable state, provided that we choose a semantic domain that supports these effects upfront.
- We demonstrate (in Section 2.5) how language fragments support reuse by developing a small library of pre-verified language components, and reusing these to compose different languages.

These contributions demonstrate that modern dependently typed languages such as Agda or Idris can take us far toward addressing the two sub-objectives from the introductory paragraphs of this paper. However, it would be an overstatement to say that our contributions address the research objective of making it “as easy as possible for DSL developers to develop and verify the type safety of typed domain-specific languages (DSLs)”. Our long-term ambition is to take the model that underpins our generic Agda framework, and implement it in a new meta-language that lowers the barrier for entry, and allows language designers to build verified DSLs from reusable components. The language fragment composition operation that we introduce in Section 2.4 provides a promising model for how language component reuse could work in such a meta-language.

2.2 DATA TYPES À LA CARTE

We recall how *data types à la carte* [Swierstra, 2008] lets us define *open* data types and functions. The remaining sections of this

paper extend and build upon this framework. Our exposition closely follows the original exposition by Swierstra, with one difference: we encode types à la carte in Agda using *containers* [Abbott et al., 2005a, Altenkirch et al., 2015].¹⁸

2.2.1 Composing Data Types

The idea behind data types à la carte is to encode data type definitions as data. By treating data type definitions as data we can explicitly manipulate them, but also recover their meaning by mapping to Agda data types. We explain how to encode data type definitions as *signatures* that can be mapped to plain Agda data types (Section 2.2.1.1), how to compose them using *signature composition* (Section 2.2.1.2), and how to define *open data type constructors* using *signature subtyping* (Section 2.2.1.3).

2.2.1.1 Signatures

A signature describes a set of data type constructors. The following *record type* in Agda defines a type of signatures. We dub the record type *Signature*, but it corresponds to what is commonly known as a *finitary container* [Abbott et al., 2005a, Altenkirch et al., 2015].¹⁹ We use finite containers because they make the presentation of data types more uniform (sub-expressions are always given by a vector as we illustrate next). We could have used

```
record Signature : Set where
  constructor _▷_
  field Symbols : Set
       Arity    : Symbols → ℕ
```

Signature records can be constructed using `_▷_`²⁰, as indicated by the keyword “constructor”. Such records comprise a set²¹ of constructor *Symbols*, and a function that associates an *Arity* (given by a natural number) with each constructor symbol.

To illustrate how signatures encode inductive data types, we compare the plain inductive definitions (top) with their encod-

¹⁸ The reason we use containers instead of the signature functors that Swierstra [2008] uses is that the fixpoint of signature functors is *not strictly positive* and hence rejected by Agda (and other dependently-typed languages). The idea of using containers to implement data types à la carte in a dependently-typed language is due to Keuchel and Schrijvers [2013].

¹⁹ Most sections of this paper could also be defined in terms of plain containers whose *Arity* is not restricted to be finite.

²⁰ The underscores in the name `_▷_` indicate the argument positions of the infix operator.

²¹ *Set* is the type of types in Agda. To rule out inconsistencies, Agda has an infinite hierarchy of *Sets* (i.e., *Set* : *Set*₁ : ...), and the *Signature* type that we define really lives in *Set*₁, since one of its fields is itself a *Set*. For presentation purposes, we abstract from universe levels and write *Set* everywhere.

ing as signatures (bottom). Starting with arithmetic expressions:

22

²² The notation tt
 $ff : BoolExpr$ is
 syntactic sugar for two
 separate constructor
 declarations
 $tt : BoolExpr$ and
 $ff : BoolExpr$.

```
data ArithExpr : Set where
  lit  : ℕ → ArithExpr
  add  : ArithExpr → ArithExpr → ArithExpr

data ArithExprSymbols : Set where
  lit  : ℕ → ArithExprSymbols
  add  : ArithExprSymbols

ArithExprΣ = ArithExprSymbols ▷ (λ where
  (lit n) → 0; add → 2)
```

And Boolean expressions:

```
data BoolExpr : Set where
  tt ff : BoolExpr
  ite  : (e e1 e2 : BoolExpr) → BoolExpr

data BoolExprSymbols : Set where
  tt ff : BoolExprSymbols
  ite  : BoolExprSymbols

BoolExprΣ = BoolExprSymbols ▷ (λ where
  ite → 3; tt → 0; ff → 0)
```

We recover an inductive data type from a signature by taking a fixpoint of the corresponding *signature functor*. The function $\llbracket _ \rrbracket$ defines this functor, and the type μ its fixpoint—that is, the type of syntax trees whose constructors are given by σ .²³

²³ The notation
 $\Sigma[x : A](B x)$ denotes
 a dependent pair of a
 value $x : A$ and a
 value of type $B x$ for
 some $B : A \rightarrow Set$.

```
\llbracket \_ \rrbracket : Signature → (Set → Set)
\llbracket σ \rrbracket A = Σ[ s : Symbols σ ] (Vec A (Arity σ s))

data μ (σ : Signature) : Set where
  ⟨ \_ \rangle : \llbracket σ \rrbracket (μ σ) → μ σ
```

Taking the least fixpoint of the $ArithExprΣ$ signature yields a type that is equivalent to $ArithExpr$:

```

example0 : ArithExpr
example0 = lit 42

example0' : μ ArithExprΣ
example0' = ⟨ lit 42 , [] ⟩

example1 : ArithExpr
example1 = add (lit 11) (lit 31)

example1' : μ ArithExprΣ
example1' = ⟨ add , ⟨ lit 11 , [] ⟩ :: ⟨ lit 31 , [] ⟩ :: [] ⟩

```

2.2.1.2 Signature Composition

Signatures can be composed by taking the disjoint union of their symbols and arities. The function `+` defines this disjoint composition using the usual sum type (\uplus):

```

_ :+ : Signature → Signature → Signature
σ1 :+ : σ2 = (Symbols σ1  $\uplus$  Symbols σ2) ▷ (λ where
  (inj1 s) → Arity σ1 s
  (inj2 s) → Arity σ2 s)

data _ $\uplus$ _ (A B : Set) : Set where
  inj1 : A → A  $\uplus$  B
  inj2 : B → A  $\uplus$  B

```

Using disjoint signature composition we can define signatures in isolation and compose them without having to re-check them. For example, we can compose the `ArithExprΣ` signature with `BoolExprΣ`, accommodating expressions that mix arithmetic and Boolean expressions. For example, the term below encodes a simple if-then-else expression (`ite tt 42 0`):

```

example2 : μ (ArithExprΣ :+ : BoolExprΣ)
example2 =
  ⟨ inj2 ite ,
    ⟨ inj2 tt , [] ⟩
    :: ⟨ inj1 (lit 42) , [] ⟩
  ⟩

```

```

:: < inj1 (lit 0) , [] >
:: [] >

```

The repeated applications of `inji` make working with composed signatures cumbersome. Following Swierstra [2008], we address this using signature subtyping and smart constructors.

2.2.1.3 Signature Subtyping and Smart Constructors

Smart constructors construct instances of a data type whose full set of constructors is left open. For example, the following function constructs a literal in the syntax tree of any signature σ that contains the symbol/arity pairs of `ArithExprΣ`.²⁴

```
lit' : { ArithExprΣ ≤ σ } → ℕ → μ σ
```

The σ in the type of `lit'` is thus decided by the context that it is used in, enabling us to flexibly reuse `lit'` in different contexts. The source of this flexibility is signature subtyping.²⁵

```

record _≤_ (σ1 σ2 : Signature) : Set where
  field inj      : [[ σ1 ]] A → [[ σ2 ]] A
        proj     : [[ σ2 ]] A → Maybe ( [[ σ1 ]] A )
        proj-inj : { x : [[ σ1 ]] A }
                  → proj (inj x) ≡ just x
        inj-proj : { x : [[ σ1 ]] A } { y : [[ σ2 ]] A }
                  → proj y ≡ just x → inj x ≡ y

```

The type $\sigma_1 \leq \sigma_2$ witnesses that it is always possible to *inject* elements in the interpretation of σ_1 into the interpretation of σ_2 , whereas the converse *projection* is only partial. The `proj-inj` and `inj-proj` fields establish that injection and projection are partial inverses.

It is possible to automatically search for injections into co-products using instance parameters [Devriese and Piessens, 2011]. We elide the definition of the necessary instances, but they are entirely analogous to the instances found in the original data types à la carte framework [Swierstra, 2008]. The code accompanying this paper also contains the implementation.

²⁴ The double braces `{ ... }` declare an instance parameter. Instance parameters are similar to type class constraints in Haskell or implicit parameters in Scala: when invoking a function with an instance parameter, an automatic search is performed during type checking time at the call-site of the function. The search will either automatically find an instance that has the correct type, or cause type checking to fail if Agda cannot find an instance parameter of the expected type.

²⁵ The curly braces in `proj-inj : { x : F A } → ...` are implicit function arguments that Agda will attempt to automatically infer when we construct a record element. The `≡` type is propositional equality.

Using signature subtyping we can implement the smart constructor for `lit'` mentioned above. We use the smart `inject` function on below to implement a smart constructor for literals:

```
inject : { σ1 ≼ σ2 } → [ [ σ1 ] ] (μ σ2) → μ σ2
inject x = ⟨ inj x ⟩

lit' : { ArithExprΣ ≼ σ } → ℕ → μ σ
lit' n = inject (lit n , [ ])
```

By defining similar smart constructors for Boolean expressions, `example2'` from above can be implemented more concisely as follows:

```
example2' : μ (ArithExprΣ :+: BoolExprΣ)
example2' = ite' tt' (lit' 42) (lit' 0)
```

2.2.2 Composing Functions

We recall how to define a function by cases using data types à la carte.

2.2.2.1 Algebras

The function `fold` transforms a tree of type `μ σ` into a value of type `A`:

```
fold : ([ [ σ ] ] A → A) → μ σ → A
```

The parameter of type `([[σ]] A → A)` is called an *algebra*, and determines how to turn a constructor whose sub-trees are already folded into a value of type `A`, into a value of type `A`. We abbreviate the type of algebras using the following alias:

```
Algebra : Signature → Set → Set
Algebra σ A = [ [ σ ] ] A → A

ArithAlg : Algebra ArithExprΣ ℕ
ArithAlg (lit n , [ ]) = n
ArithAlg (add , n1 :: n2 :: [ ]) = n1 + n2
```

The `ArithAlg` function defines an algebra that evaluates arithmetic expressions to natural numbers. The patterns n and m bound by matching on the `add` constructor are not expressions, but rather the numbers that result from evaluating the expressions in those positions. The function `fold` takes care of evaluating sub-expressions, and is defined as follows:

`mutual`

```

fold : Algebra σ A → μ σ → A
fold f ⟨ s , v ⟩ = f (s , map-fold f v)

map-fold : Algebra σ A → Vec (μ σ) n → Vec A n
map-fold f [] = []
map-fold f (x :: v) = fold f x :: map-fold f v

```

To pass Agda’s termination checker, we must inline the definition of `map` for lists (`map-fold`), which applies `fold` to recursive sub-expressions.

2.2.2.2 Algebra Composition

We can sum algebras using the \oplus operator given below.

```

_⊕_ : Algebra σ1 A → Algebra σ2 A → Algebra (σ1 :+: σ2) A
(f ⊕ g) (inj1 s , v) = f (s , v)
(f ⊕ g) (inj2 s , v) = g (s , v)

```

Summing two algebras over two signatures σ_1 and σ_2 thus yields a larger algebra for the signature composition $\sigma_1 :+: \sigma_2$: This algebra sum operator only works for algebras with the *same carrier type* $A : \text{Set}$. This implies that the `ArithAlg` algebra can *only* be composed with algebras that also use `ℕ` as their carrier. This excludes, for example, the composition of `ArithAlg` with an algebra for Boolean expressions with `Bool` as its carrier type.

We can allow such compositions by defining algebras with an *open* carrier type; i.e., using signature subtyping. The idea is to represent values as signatures, and use signature subtyping

to assert what value constructors each algebra *at least* requires; e.g.:

$$\begin{aligned} \text{ArithAlg}_0 &: \{ \text{NatVal}\Sigma \preceq \sigma \} \rightarrow \{ \text{StuckVal}\Sigma \preceq \sigma \} \\ &\rightarrow \text{Algebra ArithExpr}\Sigma (\mu \sigma) \end{aligned}$$

The carrier of this algebra is the fixpoint of *some* signature σ , about which we only know that it contains at least the constructors described by `NatVal` and `StuckVal`. As the names suggest, `NatVal` describes natural number values, and `StuckVal` represents a *stuck* value. Stuck values are needed because the recursive positions of `ArithExpr` are not intrinsically guaranteed to return numbers, and if they do not, interpretation gets stuck. By defining a similar algebra for `BoolExpr`, we can assemble an interpreter that we can use to evaluate `example2'`:

$$\begin{aligned} \text{interpArithBool} &: \mu (\text{ArithExpr}\Sigma \text{ :+} : \text{BoolExpr}\Sigma) \\ &\rightarrow \mu (\text{NatVal}\Sigma \text{ :+} : \text{BoolVal}\Sigma \text{ :+} : \text{StuckVal}\Sigma) \\ \text{interpArithBool} &= \text{fold} (\text{ArithAlg}_0 \oplus \text{BoolAlg}) \end{aligned}$$

$$\begin{aligned} \text{test} &: \text{interpArithBool } \text{example2}' \\ &\equiv \text{inject} (\text{nat } 42, []) \\ \text{test} &= \text{refl} \end{aligned}$$

2.2.3 Discussion

The interpreters that we can write using the techniques shown in this section are inherently weakly typed. This weak typing means that we must define values as a data type with separate constructors for each kind of a value. As a result, the algebras we define describe partial functions, returning stuck if a recursively evaluated value is not tagged with the right constructor. By contrast, the intrinsically-typed interpreter shown in the introduction is tagless: Agda uses the type index to figure out what kind of value we are dealing with, allowing us to work with bare values instead. As a result this interpreter is type-safe

by construction. To define such interpreters modularly we lift data types à la carte to indexed types, defining the following types in a composable way:

$$\begin{aligned} \text{Ty} & : \text{Set} \\ \text{Val} & : \text{Ty} \rightarrow \text{Set} \\ \text{Expr} & : \text{Ty} \rightarrow \text{Set} \\ \text{interp} & : \forall \{t\} \rightarrow \text{Expr } t \rightarrow \text{Val } t \end{aligned}$$

The key challenge is that the clauses of `interp` may use dependent pattern matching on values `Val t` at an index type $t : \text{Ty}$. If we know exactly what t is, we only have to consider the cases that `Val` associates with that type. But if `Val` and `Ty` are open-ended, how do we know that these will remain to be the only possible values for t ?

The answer to this question is to ensure that values have *canonical forms*. Canonical forms lemmas are a widely-used technique for making type safety proofs robust under extension and composition [Wright and Felleisen, 1994, Pierce, 2002, Delaware et al., 2013b]. In the next section we show that this technique, in conjunction with *indexed* data types à la carte, provides exactly the abstraction we need to encode composable intrinsically-typed interpreters.

2.3 INDEXED DATA TYPES À LA CARTE, FOR DEFINING COMPOSABLE INTRINSICALLY-TYPED INTERPRETERS

We extend the data types à la carte framework to *indexed* data types, and encode *intrinsically-typed interpreters* in this framework as follows:

- We encode types `Ty : Set` as a plain *signature* (Section 2.3.1).
- We encode values `Val : Ty → Set` as an *algebra* over object language type signatures (Section 2.3.2).
- We encode expressions `Expr : Ty → Set` as an *indexed signature* with an *open index type* (Section 2.3.3) such that we can

add new expression- or type constructors, without having to modify or re-check existing expression constructors.

- We encode interpreters $\text{interp} : \text{Expr } t \rightarrow \text{Val } t$ as an *indexed algebra* over the indexed signatures of object language expressions (Section 2.3.4). These indexed algebras have an *open carrier type* and an *open index type*, allowing us to can add new expression constructors, interpreter cases, types, and values, without modifying or re-checking existing code. It is crucial that the carrier of these indexed algebras is only open to extensions that *preserve canonical forms*.

2.3.1 Composing Index Types

Since the index type $\text{Ty} : \text{Set}$ is a plain data type, we can use plain data types à la carte to encode it as a signature. For example, below on the left is a data type NatTy representing a notion of object language type with a single type constructor, and on the right is its signature encoding:

```
data NatTy : Set where
  nat : NatTy
data NatTyShape : Set where
  nat : NatTyShape
NatTyΣ = NatTyShape ▷ (λ _ → 0)
```

By similarly encoding a Boolean type constructor as a signature BoolTyΣ , we can compose a signature that encodes the object language types of the interpreter from Figure 1:

```
ArithBoolTyΣ = NatTyΣ :+: BoolTyΣ
```

2.3.2 Composing Intrinsically-Typed Values

The intrinsically-typed interpreter in Figure 1 defines $\text{Val} : \text{Ty} \rightarrow \text{Set}$ as a function. This function maps object language types to their

canonical forms (i.e., the set of possible value constructors) of that type. (Note that it is also possible to model `Val` as a data type, but a benefit of modeling `Val` as a function is that values are *tag-less* [Augustsson and Carlsson, 1999], which avoids the need to tag and untag values in the interpreter.) Since `Ty : Set` is encoded as a signature, we can encode `Val` as an algebra over that signature. The following algebras define the canonical forms of `nat` and `bool`:

```
NatVal : Algebra NatTyΣ Set      BoolVal : Algebra BoolTyΣ Set
NatVal (nat , []) = ℕ           BoolVal (bool , []) = Bool
```

We can compose these algebras using the algebra sum operation from Section 2.2.2:

```
ArithBoolVal : Algebra ArithBoolTyΣ Set
ArithBoolVal = NatVal ⊕ BoolVal
```

Folding `ArithBoolVal` over `ArithBoolTyΣ` yields a function that is isomorphic to `Val` from Figure 1.

2.3.3 Composing Intrinsically-Typed Expressions

The expressions of the intrinsically-typed interpreter in Figure 1 are defined as an *indexed data type* `Expr : Ty → Set` whose index type is `Ty`. To define this type in a composable way, we lift the notion of signature from the data types à la carte framework discussed in Section 2.2 to *indexed signatures*, and obtain a framework for *indexed data types* à la carte.

2.3.3.1 Indexed Signatures

Below is the type of *indexed signatures* that describe I -indexed data types. We dub this type `ISignature I`, but this type is also commonly known as a *finitary indexed container* [Altenkirch et al., 2015]:

```
record ISignature (I : Set) : Set₁ where
  constructor _►_
```

```

field ISymbols : I → Set
Indices       : {i : I} → ISymbols i → List I

```

The `ISymbols` field relates each index to a set of symbols, and the `Indices` field associates the symbols at each index with a list whose length represents the arity of each constructor symbol, and whose elements describe what the index (of type I) of each recursive position is. We can interpret indexed signatures as indexed data types, just as we interpreted signatures as plain data types (Section 2.2):

```

I[ ] : ISignature I → (I → Set) → (I → Set)
I[ ζ ] P i = Σ [ s : ISymbols ζ i ] (All P (Indices ζ s))

data Iμ (ζ : ISignature I) : I → Set where
  I[ ] : {i : I} → I[ ζ ] (Iμ ζ) i → Iμ ζ i

```

The implementation of the `I[]` function uses the following `All` relation on lists, which asserts that each element in a given `List I` satisfies a given proposition $P : I \rightarrow \text{Set}$:

```

data All (P : I → Set) : List I → Set where
  [] : All P []
  _::_ : {i : I} {xs : List I} → P i → All P xs → All P (i :: xs)

```

We can think of `I[ζ]` as mapping an indexed signature ζ to its corresponding signature functor on I -indexed types, similarly to how `[σ]` maps a plain signature to its corresponding signature functor on plain types.

2.3.3.2 Indexed Signature Composition

Two indexed signatures with the same index type I can be summed into a larger signature that comprises the shapes and indices of both:

```

_::+_ : ISignature I → ISignature I → ISignature I
ISymbols (ζ1 ::+ ζ2) i = ISymbols ζ1 i ⊔ ISymbols ζ2 i
Indices (ζ1 ::+ ζ2) (inj1 s) = Indices ζ1 s
Indices (ζ1 ::+ ζ2) (inj2 s) = Indices ζ2 s

```

Using these ingredients we can now define composable indexed expression types corresponding to the $\text{Expr} : \text{Ty} \rightarrow \text{Set}$ from Figure 1 where the index type Ty is *fixed*. However, since new language fragments may introduce new object language type constructors, we need to model Ty in a way that allows such extensions.

2.3.3.3 Indexed Signatures with Open Index Types

We can define signatures whose index type is open by using the subtyping relation $_ \preceq _$ (Section 2.2.1.3) to witness a lower bound on the index type, just like we did in Section 2.2.2.2. The indexed signature below thus describes intrinsically-typed arithmetic expressions whose type constructors are described by any signature σ that contains *at least* the constructors described by $\text{NatTy}\Sigma$:

```
data IArithExprSymbols { _ : NatTyΣ ≼ σ } : μ σ → Set where
  val : ℕ → IArithExprSymbols (inject (nat , []))
  add :      IArithExprSymbols (inject (nat , []))

IArithExprΣ : { _ : NatTyΣ ≼ σ } → ISignature (μ σ)
IArithExprΣ = IArithExprSymbols ▶ (λ where
  (val _) → []
  add     → inject (nat , []) :: inject (nat , []) :: [])
```

By defining Boolean expressions in a similar manner, we can compose indexed signatures with open index types:

```
IArithBoolExprΣ : { _ : NatTyΣ ≼ σ } → { _ : BoolTyΣ ≼ σ }
                → ISignature (μ σ)
IArithBoolExprΣ = IArithExprΣ :++: IBoolExprΣ
```

By defining a similar subtyping relation for indexed signatures as we did for plain signatures, we can define smart constructors for indexed types, similarly to how we defined smart constructors for plain types in Section 2.2.1.3. We elide this relation for brevity, and refer the interested reader to the code accompanying the paper where it can be found.

2.3.4 Composing Index-Preserving Functions

Our goal is to use indexed algebras to encode interpreters of type $\text{interp} : \text{Expr } t \rightarrow \text{Val } t$ whose index type $\text{Ty} : \text{Set}$ and value type $\text{Val} : \text{Ty} \rightarrow \text{Set}$ are *open*. That is, we should define interp in a way that we can add new constructors to $\text{Ty} : \text{Set}$ and $\text{Val} : \text{Ty} \rightarrow \text{Set}$ and ensure that pattern matches inside interp on values remain exhaustive. We realize this goal by defining a subtyping relation in Section 2.3.4.2 that characterizes such safe extensions. But first, we need indexed algebras.

2.3.4.1 Indexed Algebras

We can generically fold a tree of type $\text{l}\mu \zeta i$ into a value of type $P i$ using the following function, where $P : I \rightarrow \text{Set}$:

$$\text{lfold} : \forall [I [\zeta] P \Rightarrow P] \rightarrow \forall [\text{l}\mu \zeta \Rightarrow P]$$

Its implementation is analogous to the implementation of fold from Section 2.2.2.1, and its type uses the following abbreviations for indexed types (both from the Agda Standard Library²⁶):

$$\begin{aligned} \forall [_] & : (I \rightarrow \text{Set}) \rightarrow \text{Set} \\ \forall [_] \{ I \} P & = \{ i : I \} \rightarrow P i \\ _ \Rightarrow _ & : (I \rightarrow \text{Set}) \rightarrow (I \rightarrow \text{Set}) \rightarrow (I \rightarrow \text{Set}) \\ (P \Rightarrow Q) i & = P i \rightarrow Q i \end{aligned}$$

We call the function parameter $\forall [I [\zeta] P \Rightarrow P]$ an *indexed algebra*, which we will abbreviate using the following alias:

$$\begin{aligned} \text{IAlgebra} & : (\zeta : \text{ISignature } I) (P : I \rightarrow \text{Set}) \rightarrow \text{Set} \\ \text{IAlgebra } \zeta P & = \forall [I [\zeta] P \Rightarrow P] \end{aligned}$$

Just like plain algebras, indexed algebras are closed under indexed signature sums:

$$\begin{aligned} _ : \oplus _ & : \text{IAlgebra } \zeta_1 P \rightarrow \text{IAlgebra } \zeta_2 P \\ & \rightarrow \text{IAlgebra } (\zeta_1 : ++ : \zeta_2) P \end{aligned}$$

26

<https://github.com/agda/agda-stdlib>

$$\begin{aligned} (f : \oplus) g) (\text{inj}_1 s, a) &= f(s, a) \\ (f : \oplus) g) (\text{inj}_2 s, a) &= g(s, a) \end{aligned}$$

This sum operation assumes that both summands have the *same* carrier, P . To sum indexed algebras corresponding to intrinsically-typed interpreters with *different* notions of types and values, we need a subtyping relation that witnesses what values (or *canonical forms*) a carrier type *at least* has.

2.3.4.2 Canonical Forms Subtyping

Our goal is to define indexed algebras with open carrier types in the style illustrated by `interpArithAlg` below:

```
interpArithAlg : { _ : NatVal ⊆ W }
               → IAlgebra IArithExprΣ (fold W)
interpArithAlg (val n,          []) = ↑ n
interpArithAlg (add , n1 :: n2 :: []) = ↑ (↓ n1 + ↓ n2)
```

Here, `_ ⊆ _` denotes value subtyping. In order to define composable intrinsically-typed interpreters, this subtyping relation should witness that we can *safely* convert between `NatVal` and W . In the definition of `interpArith`, we use \uparrow to represent a safe “upcast” from `NatVal` to W , and \downarrow to represent a safe “downcast” from W to `NatVal`. We define `_ ⊆ _` in terms of a *type isomorphism*:

```
record _ ↔ _ (A B : Set) : Set where
  field inj↔      : A → B
        proj↔     : B → A
        proj-inj↔ : {a : A} → proj↔ (inj↔ a) ≡ a
        inj-proj↔ : {b : B} → inj↔ (proj↔ b) ≡ b
```

The `inj↔` and `proj↔` fields lets us convert any A into a B and vice versa. The `proj-inj↔` and `inj-proj↔` fields restrict `inj↔` and `proj↔` to proper inverses.

Intuitively, a witness of the form $W_1 \subseteq W_2$ tells us that W_1 and W_2 define *the same values for the same types*, but W_2 may define values for more types than W_1 .

```

record _⊆_ (W1 : Algebra σ1 Set)
      (W2 : Algebra σ2 Set) : Set where
  field { ≲-type } : σ1 ≲ σ2
      canonical : {V : T → Set} → (t : [ σ1 ] T)
      → W1 (fmap V t) ↔ W2 (fmap V (inj t))

```

The \preceq -type²⁷ field establishes that W_2 is defined on the same (or more) types as W_1 . The **canonical** field asserts that W_1 and W_2 define *the same value* (up to isomorphism) for every shared type. We express this fact by requiring that W_1 and W_2 are isomorphic for all shared types t , invariant of the type of its sub-trees (T), or the way sub-trees are mapped to values (V). To apply V to the sub-trees in t , we use the function **fmap**:

```

fmap : (f : A → B) → [ σ ] A → [ σ ] B
fmap f (s , v) = (s , map f v)

```

Finally, we can implement the safe up-casting (\uparrow) and down-casting (\downarrow) operations:

```

↑ : { _ : V ⊆ W } → V (fmap (fold W) t) → fold W (inject t)
↓ : { _ : V ⊆ W } → fold W (inject t) → V (fmap (fold W) t)

```

The implementation of these function uses the type isomorphism in the **canonical** field of the $\{ _ : V \subseteq W \}$ instance argument to safely convert between the values computed by the sub- and super algebras. Comparing with Section 2.2, we can think of \uparrow as the intrinsically-typed counterpart to **inject**, and of \downarrow as the intrinsically-typed counterpart to a weakly-typed projection function. The crucial difference is that \downarrow is a total function, whereas **project** may return **nothing**.

2.3.4.3 Indexed Algebra Composition

Indexed algebras can be composed in the same way as plain algebras. For example, we can compose **interpArithAlg** with the following indexed algebra for evaluating intrinsically-typed Boolean expressions:

²⁷ By wrapping the \preceq -type field of this record type in instance argument brackets, instance parameter search will be able to automatically resolve this field projection. For example, the **inj** field projection that occurs in the **canonical** field of the $_ \subseteq _$ record type is implicitly projecting from \preceq -type.

$$\begin{aligned}
\text{interpBoolAlg} &: \{ _ : \text{BoolVal} \subseteq W \} \\
&\quad \rightarrow \text{Algebra } \text{IBoolExpr}\Sigma \text{ (fold } W) \\
\text{interpBoolAlg} \text{ (tt } _, _) &= \uparrow \text{ true} \\
\text{interpBoolAlg} \text{ (ff } _, _) &= \uparrow \text{ true} \\
\text{interpBoolAlg} \text{ (ite } _, v :: v_1 :: v_2 :: _) &= \text{if } \downarrow v \text{ then } v_1 \text{ else } v_2 \\
\text{interpArithBoolAlg} &: \{ _ : \text{NatVal} \subseteq W \} \rightarrow \{ _ : \text{BoolVal} \subseteq W \} \\
&\quad \rightarrow \text{Algebra } \text{IArithBoolExpr}\Sigma \text{ (fold } W) \\
\text{interpArithBoolAlg} &= \text{interpArithAlg} : \oplus : \text{interpBoolAlg}
\end{aligned}$$

Using these indexed algebras, we can now evaluate expressions given by the fixpoint of an indexed signature, using `lfold`:

$$\text{lfold} : \text{Algebra } \zeta P \rightarrow \forall [\mu \zeta \Rightarrow P]$$

This yields an interpreter that is analogous to the interpreter from Figure 1 in the introduction:

$$\begin{aligned}
\text{interpArithBool} &: \mu \text{IArithBoolExpr}\Sigma t \\
&\quad \rightarrow \text{fold } (\text{NatVal} \oplus \text{BoolVal}) t \\
\text{interpArithBool} &= \text{lfold } \text{interpArithBoolAlg}
\end{aligned}$$

Unlike the interpreter from the introduction, this interpreter is assembled from separately-defined, reusable components—i.e., `interpArithAlg` and `interpBoolAlg`.

2.3.5 Discussion

This section showed how to assemble intrinsically-typed interpreters using *indexed* data types à la carte and canonical forms subtyping. While this approach allows us to assemble languages, the notion of *language fragment* that we illustrated in Figure 2 in the introduction remains informal. There are three reasons why it is useful to define language fragments as a first-class abstraction instead. The first reason is that it would allow language designers to compose languages using a single, uniform notion of *language fragment composition*, instead of the four we used in this section to compose respectively types, values, expressions, and interpreters. The second reason is that

intrinsically-typed interpreters as defined in this section are not closed under composition, since indexed algebra composition “grows” the number of canonical forms subtype constraints, as `interpArithBoolAlg` illustrates. The last reason is that it is possible for interpreters to make *conflicting assumptions* about values, as we illustrate below.

Consider the encoding below of an intrinsically-typed interpreter which uses *ternary* Booleans, akin to the `Val bool = Maybe Bool` clause discussed in Section 2.1.2. The `null` expression below constructs the third kind of Boolean value:

```
TernaryBoolVal : Algebra BoolTyΣ Set
TernaryBoolVal (bool , []) = Maybe Bool

data INullExprShape { _ : BoolTyΣ ≼ σ } : μ σ → Set where
  null : INullExprShape (inject (bool , []))

INullExprΣ : { BoolTyΣ ≼ σ } → ISignature (μ σ)
INullExprΣ = INullExprShape ▶ λ _ → []

interpNullAlg : { _ : TernaryBoolVal ⊆ W }
               → IAlgebra INullExprΣ (fold W)
interpNullAlg (null , []) = ↑ nothing
```

Using algebra composition, we can compose `interpBoolAlg` with `interpNullAlg`, growing the number of subtype constraints:

```
- Causes unsolved instances, even with
- overlapping-instances enabled!
{- badAlgebra : { _ : BoolVal ⊆ V }
               → { _ : TernaryBoolVal ⊆ V }
               → IAlgebra
                  (IBoolExprΣ :++: INullExprΣ)
                  (fold V)
badAlgebra = interpBoolAlg :⊕: interpNullAlg -}
```

This is unsatisfactory: the subtype constraints of `badAlgebra` represent an *unsound* identification of types, because `BoolVal` and `TernaryBoolVal` define different canonical forms for the `bool` type. In other words, the subtype constraints *conflict*—i.e., cannot be

proven—for any $V : \text{Algebra BoolTy}\Sigma \text{ Set}$. The next section introduces *intrinsically-typed language fragments*, with a composition operator that avoids the issues discussed above, and enables sound intrinsically typed fragment composition.

2.4 INTRINSICALLY-TYPED LANGUAGE FRAGMENTS

In this section we introduce *language fragments*: an abstraction that bundles a set of intrinsically-typed syntax constructors with the associated cases of an interpreter. This abstraction comes with a single composition operation, *language fragment composition*, that has nice closure properties, and that subsumes the four different notions of composition that we introduced in the previous section (for object language type signatures, value typing algebras, expression data type signatures, and interpreter algebras). It thus goes towards addressing both the first and the second sub-objective stated in the introduction of this paper, by making it easier to develop type safe DSL components in a way that supports reuse.

We first discuss (Section 2.4.1) how to bundle these four components in a way that each component is defined as being *open*. We then discuss (Section 2.4.2) how to compose language fragments, and why it is necessary to allow compositions with partially-overlapping canonical forms. Finally, we present (Section 2.4.3) language fragment composition, which makes language fragments closed under composition. The idea of language fragments, however, transcends their formulation in this section, which concerns the definition and composition fragments of simply-typed expression languages. In Section 2.5, we will consider how to apply the same techniques to a more expressive semantic domain.

2.4.1 *Canons and Language Fragments*

Language fragments (**Fragment**) bundle a piece of intrinsically-typed syntax with its interpretation, and are parameterized over

the available *canonical forms* (`Canon`)—i.e., a signature of object language types together with an algebra over this signature. The type of language fragments is given by the `Fragment` record type below, which in turn is parameterized by a set of canonical forms as given by the `Canon` record type:

```
record Canon : Set where
  constructor canon
  field
    ty : Signature
    val : Algebra ty Set

record Fragment (c : Canon) : Set where
  field expr : ISignature (μ (ty c))
  interp : IAlgebra expr (fold (val c))
```

A value of type `Fragment c` is a *self-contained description* of an intrinsically-typed interpreter. It is, however, defined in terms of the *fixed* set of canonical forms given by c , meaning that the composition of such fragments is limited to fragments that depend on the *same* canonical forms. To compose language fragments with *different* canonical forms, we must define them with extension of the canon in mind, similar to how we defined indexed algebras with open carriers in Section 2.3.4.2.

Because language fragments are a self-contained description of intrinsically-typed interpreters, we can view extensibility of canonical forms independent from the definition of fragments. To do this, we define the \square modifier, which *closes* a canon-indexed type over all possible future extensions:²⁸

```
□ : (Canon → Set) → Canon → Set
□ P c = ∀ {c'} → { val c ⊆ val c' } → P c'
```

We use the \square modifier to define *open* language fragments for the intrinsically-typed interpreter components from Section 2.3:

```
ArithFrag : □ Fragment (canon NatTyΣ
                        NatVal)
```

²⁸ The \square notation is borrowed from the necessity modality of modal logic. Its implementation is essentially a shallow embedding of Kripke Semantics [Kripke, 1963] and is inspired by Allais et al. [2018].

```

expr ArithFrag = IArithExprΣ
interp ArithFrag = interpArithAlg

BoolFrag : □ Fragment ( canon BoolTyΣ
                        BoolVal )

expr BoolFrag = IBoolExprΣ
interp BoolFrag = interpBoolAlg

```

The \square modifier provides exactly the subtyping proof that `interpBoolAlg` and `interpArithAlg` need. Furthermore, it is always possible to transform an open fragment to a closed fragment, and extract an interpreter from a closed fragment:

```

extract : ∀[ □ P ⇒ P ]
extract φ = φ { ⊆-refl }

toInterp : ( φ : Fragment c ) → ∀[ lμ ( expr φ ) ⇒ fold ( val c ) ]
toInterp φ = lfold ( interp φ )

```

2.4.2 Fragment Composition and the Need for Partially-Overlapping Canons

We define a composition operation for open fragments with the same canonical forms as follows:

```

fcompose-eq : ∀[ □ Fragment ⇒ □ Fragment ⇒ □ Fragment ]
expr ( fcompose-eq φ1 φ2 ) = expr φ1 :++: expr φ2
interp ( fcompose-eq φ1 φ2 ) = interp φ1 :⊕: interp φ2

```

In many scenarios, however, `fcompose-eq` alone is insufficient: it is often necessary to compose fragments with *different* canonical forms. Indeed, `ArithFrag` and `BoolFrag` have different canonical forms, so we cannot compose them into a new open fragment with `fcompose-eq`.

It may seem tempting to use a composition operation that *sums* the canonical forms of fragments, but this is also problematic. Say we have a fragment that defines an interpreter for a binary less-than-or-equals expression that compares the results of evaluating its arguments:

```

LeqFrag : □ Fragment (canon
  (NatTyΣ :+: BoolTyΣ)
  (NatVal ⊕ BoolVal))

```

A fragment that combines `LeqFrag` and `BoolFrag` by summing their canons has the following type:

```

Fragment (canon
  (NatTyΣ :+: BoolTyΣ :+: BoolTyΣ)
  (NatVal ⊕ (BoolVal ⊕ BoolVal)))

```

The resulting fragment has two *distinct* notions of Boolean types and values! As a result, an expression such as `ite (leq 0 1) 42 0` would be ill-typed; the `bool` type of `leq 0 1` is not the same `bool` type that `ite` expects. Rather, we should identify the Boolean types of the two fragments. In other words: their canonical forms are *partially overlapping*.

2.4.3 Fragment Composition with Partially-Overlapping Canons

We introduce a fragment composition operation for language fragments with partially-overlapping canons in three stages. First, we introduce *type unions* (Section 2.4.3.1), which precisely characterize how two types overlap. Then, we define a similar *union for canonical forms* (Section 2.4.3.2) in terms of this type union, which describes how two canons overlap. Finally, we define *fragment composition* (Section 2.4.3.3), which combines two open fragments, given a canon union that witnesses how their canonical forms are combined.

2.4.3.1 Overlapping Unions for Types

Figure 3 defines the `Union` relation, which describes how the elements of two types $A : \text{Set}$ and $B : \text{Set}$ map to elements in a third type $C : \text{Set}$, such that each element in C corresponds to an element of either A , B , or both. The correspondence is witnessed by the functions `inja`, `injb`, and `from`, which should be injective.²⁹

²⁹ For readers wondering how general this type union is: the `Union` type is a pushout in the category of Agda `Sets`, meaning it is union-like indeed.

| | |
|--|---|
| <pre> record Union (A B C : Set) : Set where field inja : A → C injb : B → C from : C → These A B inja-inv : ∀ {a} → ⟨⟨ _ ≡ a , ∅ , (_ ≡ a) ∘ proj₁ ⟩⟩ (from (inja a)) injb-inv : ∀ {b} → ⟨⟨ ∅ , _ ≡ b , (_ ≡ b) ∘ proj₂ ⟩⟩ (from (injb b)) from-inv : ∀ {c} → ⟨⟨ (_ ≡ c) ∘ inj_a , (_ ≡ c) ∘ inj_b , (λ (a , b) → inj_a a ≡ c × inj_b b ≡ c) ⟩⟩ (from c) </pre> | <pre> data These (A B : Set) : Set where this : A → These A B that : B → These A B these : A → B → These A B ⟨⟨ _ , _ , _ ⟩⟩ : (A → X) → (B → X) → (A × B → X) → These A B → X ⟨⟨ f , g , h ⟩⟩ (this a) = f a ⟨⟨ f , g , h ⟩⟩ (that b) = g b ⟨⟨ f , g , h ⟩⟩ (these a b) = h (a , b) </pre> |
|--|---|

Figure 3: The proof-relevant relation `Union A B C` specifies that `C` is the union of `A` and `B`. The types `A` and `B` can overlap in `C`, in which case `inja a ≡ injb b` for some elements `a : A` and `b : B`. The ternary relation uses the type `These` and its `⟨⟨ _ , _ , _ ⟩⟩` eliminator shown on the right.

There are two trivial unions which can always be constructed:

```

union-copy : Union A A A
union-disjoint : Union A B (A ⊔ B)

```

The function `union-copy` constructs a union of `A` with itself such that all its elements overlap with themselves. Conversely, `union-disjoint` constructs a union of two sets of types `A` and `B` such that none of their elements are identified in `A ⊔ B`. We elide the proofs of the inverse laws (`inja-inv`, `injb-inv`, and `from-inv`) here; the code accompanying this paper contains the details.

2.4.3.2 Overlapping Canons

Using the `Union` relation, we define a similar ternary union relation for canons:

```

record _•_≡_ (c1 c2 c : Canon) : Set where
  field
    { ty-union } : Union
      ( [ (ty c1) ] T ) ( [ (ty c2) ] T ) ( [ (ty c) ] T )
    canonicall : { V : T → Set } { t : [ (ty c1) ] T }
      → (val c1) (fmap V t) ↔ (val c) (fmap V (inja t))

```

$$\begin{aligned} \text{canonical}_\top &: \{V : T \rightarrow \text{Set}\} \{t : \llbracket (\text{ty } c_2) \rrbracket T\} \\ &\rightarrow (\text{val } c_2) (\text{fmap } V t) \leftrightarrow (\text{val } c) (\text{fmap } V (\text{injb } t)) \end{aligned}$$

The canon union is directed by a type union: **ty-union** witnesses that the type constructors of c are a union of the type constructors in c_1 and c_2 . The key, however, are the additional proofs **canonical_l** and **canonical_r** that witness that the values of c_1 and c_2 agree in the overlap (as described by the type union) of their types. This amounts to a modularization of canonicity lemmas.

Since canon union is based on type union, we can construct similar trivial unions, where either all type constructors are overlapping (**•-copy**), or no type constructors are overlapping (**•-disjoint**, which uses the auxiliary disjoint canon union on the left):

$$\begin{aligned} _ \uplus^c _ &: (c_1 \ c_2 : \text{Canon}) \rightarrow \text{Canon} \\ \text{ty } (c_1 \uplus^c c_2) &= (\text{ty } c_1) :+ (\text{ty } c_2) \\ \text{val } (c_1 \uplus^c c_2) &= (\text{val } c_1) \oplus (\text{val } c_2) \\ \bullet\text{-copy} &: c \bullet c \quad \equiv c \\ \bullet\text{-disjoint} &: c_1 \bullet c_2 \quad \equiv (c_1 \uplus^c c_2) \end{aligned}$$

We show how to use **•-disjoint** and **•-copy** to compose **ArithFrag**, **BoolFrag**, and **LeqFrag** in Section 2.4.3.4.

2.4.3.3 Fragment Composition Operation

Using canon union, we can now define a general **fcompose** operation which lets us compose fragments with partially-overlapping canons:

$$\begin{aligned} \text{fcompose} &: \square \text{Fragment } c_1 \rightarrow \square \text{Fragment } c_2 \rightarrow c_1 \bullet c_2 \equiv c \\ &\rightarrow \square \text{Fragment } c \end{aligned}$$

Perhaps surprisingly, we can implement this operation in terms of the **fcompose-eq** operation, by (1) exploiting the *comonadic structure* of \square , and (2) observing that we can recover subtyping ($_ \subseteq _$) witnesses from union ($_ \bullet \equiv _$) witnesses. For \square , we already showed that it has an **extract** function Section 2.4.1.

³⁰ The *extract* and *duplicate* functions correspond to the operations of a comonad.

Additionally, we can define a *duplicate* function.³⁰, following from transitivity of $_ \subseteq _$:

$$\begin{aligned} \text{duplicate} &: \forall [\Box P \Rightarrow \Box (\Box P)] \\ \text{duplicate } px \{ w_1 \} \{ w_2 \} &= px \{ \subseteq\text{-trans } w_1 \ w_2 \} \end{aligned}$$

The type of *duplicate* says that we can “weaken” the canon that is implicitly quantified by $\forall [_]$,

Next, we observe that canon union implies value subtyping:

$$\bullet\text{-to-}\subseteq_l : c_1 \bullet c_2 \equiv c \rightarrow (\text{val } c_1) \subseteq (\text{val } c)$$

$$\bullet\text{-to-}\subseteq_r : c_1 \bullet c_2 \equiv c \rightarrow (\text{val } c_2) \subseteq (\text{val } c)$$

Using these ingredients, we can define *fcompose* in terms of *fcompose-eq* as follows:

$$\begin{aligned} \text{fcompose} &: \Box \text{Fragment } c_1 \rightarrow \Box \text{Fragment } c_2 \rightarrow c_1 \bullet c_2 \equiv c \\ &\rightarrow \Box \text{Fragment } c \\ \text{fcompose } \phi_1 \ \phi_2 \ u &= \\ \text{fcompose-eq} & \\ &(\text{duplicate } \phi_1 \{ \bullet\text{-to-}\subseteq_l \ u \}) \\ &(\text{duplicate } \phi_2 \{ \bullet\text{-to-}\subseteq_r \ u \}) \end{aligned}$$

The *duplicate* function is used to “weaken” the c_1 and c_2 canons of the left and right fragment into the canon union c . The lemmas $\bullet\text{-to-}\subseteq_l$ and $\bullet\text{-to-}\subseteq_r$ witness that this weakening is safe.

2.4.3.4 Fragment Composition Examples

With *fcompose* we compose languages from fragments:

$$\text{ArithBoolFrag} = \text{fcompose ArithFrag BoolFrag } \bullet\text{-disjoint}$$

From this fragment, we can derive the intrinsically typed interpreter discussed in Section 3.1:

$$\text{InterpArithBool} = \text{toInterp } (\text{extract ArithBoolFrag})$$

We can also compose *ArithBoolFrag* with *LeqFrag*, to obtain a larger language:

$\text{LeqArithBoolFrag} = \text{fcompose LeqFrag ArithBoolFrag} \bullet\text{-copy}$

In summary, `fcompose` addresses the three concerns we discussed in Section 2.3.5. It provides a single, uniform composition operation for language fragments with compatible but possibly different canonicity assumptions, such that language designers do not have to manually assemble types, values, expressions, and interpreters separately, using four different composition operators. It ensures that language fragments are closed under composition using `fcompose`. And the interpreters we extract from language fragments cannot have conflicting canonicity assumptions.

2.5 LANGUAGE FRAGMENTS WITH LEXICAL VARIABLES AND EFFECTS

Language fragments presuppose both the notion of typing and semantics. The definition of a fragment in Section 2.4 takes well-typed terms to be indexed families ($\text{Expr} : \text{Type} \rightarrow \text{Set}$), and their semantics to be an interpreter ($\text{interp} : \text{Expr } t \rightarrow \text{Val } t$). These presupposed notions limit the expressive power of fragments. For example, well-typed expressions with lexical binding are usually presented as indexed families that are additionally parameterized by a typing context. Their interpreter correspondingly requires an environment of values for variables. Well-typed expressions for ML-style references, on the other hand, have an intrinsically-typed interpretation in the category of monotone predicates [Bach et al., 2018]. This motivates this section, in which we show that the notion of intrinsically-typed language fragments introduced in the previous section can be transported to more expressive semantic domains.

In this section we generalize language fragments to a class of semantic domains (Section 2.5.1) that we show can be used to define intrinsically-typed interpreters for the simply-typed λ -calculus (Section 2.5.2), exceptions (Section 2.5.3), and ML-style references (Section 2.5.4). Although intrinsically-typed language fragments can be transported to this more general setting, frag-

ment composition only combines fragments that are interpreted into the same semantic domain. Since the examples in Sections 2.5.2 to 2.5.4 are interpreted in different domains, they cannot be combined into the same language using fragment composition. Instead, we can manually lift these fragments into a common “super domain” where they can be composed, as we demonstrate with our case study in Section 2.5.5. Furthermore, additional innovation may be required to modularize richer semantic domains that do not fit into the description we introduce in Section 2.5.1.

2.5.1 Fragments for a Class of Semantic Domains

We generalize the definition of language fragments to permit: (1) expressions that are typed relative to a context for de-Bruijn encoded lexical variables, (2) interpretation functions that accept a lexical environment as parameter, and (3) side effects—in particular ML-style references. The generalization is based on a generalization of the codomain of interpretation functions to some *cartesian category* \mathcal{C} —that is, a category with all products and a terminal object. Informally, the semantic domains presupposed by the definition of fragments from Section 2.4 vs. the definition we consider in this section differ as shown in Figure 4

| | Typing | Type Semantics | Term Semantics |
|-------------|-----------------------|---|---|
| Section 2.4 | $e : t$ | $\text{Val } t : \text{Set}$ | $\text{interp } (e : t) : \text{Val } t$ |
| Section 2.5 | $\Gamma \vdash e : t$ | $\text{Val } t : \text{obj}(\mathcal{C})$ | $\text{interp } (\Gamma \vdash e : t) : \mathcal{C}(\text{Env } \Gamma, \text{Val } t)$ |

Figure 4: Overview of how the domain of intrinsically-typed interpreters changes between Section 2.4 and this section

Here, $\text{obj}(\mathcal{C})$ represents the objects of \mathcal{C} , and $\mathcal{C}(\text{Env } \Gamma, \text{Val } t)$ represents the morphisms from $\text{Env } \Gamma$ to $\text{Val } t$ in \mathcal{C} .

In Agda, we can define a type of categories Category_0 and a generalized canon Canon_0 that maps types to objects of some category \mathcal{C}_0 as follows:

```
record Category0 : Set where
  field obj      : Set
      morph : obj → obj → Set
```

```
record Canon0 : Set where
  field ty : Signature
      val : Algebra ty (obj C0)
```

However, it is not possible to define a canon for ML-style reference values in this style! The problem is that the `Canon0` type assumes that we can compositionally map types onto objects of \mathcal{C}_0 . This is not true for ML-style references. To see the issue, consider the following mock case of a `refCanon : Canon0` that we wish to define, assuming that \mathcal{C}_0 is a category of monotone predicates; i.e., `obj C0 = List T → Set`.

```
val refCanon (ref , V :: []) = λ Σ → ???
```

Following how references are traditionally typed [Harper, 1994, Pierce, 2002], the right hand side is supposed to witness that there exists an location of type `t` in the store type Σ , where `t` is a subterm of `ref t`. However, since `Canon0` defines values as plain algebras, the mock case above does not have access to the subterm `t`; only to the object `V : obj C0` resulting from folding over the sub-term `t`.

It is possible to construct a compositional interpretation of reference types—that is, in terms of the object V —in the context of *semantic typing*, where types are viewed as a *set of values*. In our setting this would correspond to defining store types as a list of `Sets`, and mapping reference types to a proof that V is a member of this list. If we naively attempt to define values this way, however, their definition becomes inconsistent, since the type of these membership proofs is simultaneously bigger than and included in the set of values! Although we can escape this paradox by stratifying the interpretation of types [Ahmed et al., 2002, Ahmed, 2004], we would need to find a way to adapt the intrinsically-typed semantics for ML-style references by Bach et al. [2018] to interpret into such a layered domain.

```

record DomainDesc : Set where
  field obj      : (T : Set) → Set
  field morph    : (V : T → obj T) → (X Y : obj T) → Set

record Canon : Set where
  constructor canon
  field
    ty : Signature
    val : ∀ {T} → [ ty ] (T × obj D T)
           → obj D T

record Fragment (c : Canon) : Set where
  field
    expr : ISignature (List (μ (ty c)) × μ (ty c))
    interp : IAlgebra expr
              (λ (Γ , t) → morph D (fold (val c))
                (Env c Γ)
                (fold (val c) t))

```

Figure 5: Definition of categories, canons, and language fragments. The domain description \mathcal{D} is a module parameter, and $\text{Env } c \Gamma$ is a de Bruijn indexed environment comprising value objects whose canonical forms are given by the canon c .

Instead, we opt to generalize canons and categories to reflect that the set of objects and morphisms comprise components assembled from indexed signatures and algebras. This generalized definition is shown and Figure 5. Although similar, the explicit injection of types and values means that the resulting structure (defined in the `DomainDesc` record) is not quite a category. We will refer to this structure as a *domain description*, using \mathcal{D} to range over it. The move from categories to domain descriptions requires three generalizations:

1. Canons are generalized to use a flavor of algebra that supports *paramorphisms* [Meertens, 1992]—i.e., a recursion scheme that provides access to each sub-term both before and after we have recursively folded over it. The `val` field of `Canon` in Figure 5 shows the generalization, which maps a pair of the original sub-term of type T and its folded counterpart of type $\text{obj } \mathcal{D} T$ to an object of type $\text{obj } \mathcal{D} T$, where T is the final set of object language types.
2. Objects of a domain description may depend on a final type of object language types, as the $T : \text{Set}$ parameter of the `obj` field of `DomainDesc` in Figure 5 indicates. This

generalization is used to define modular paramorphic algebras for which we only learn the final type of object language types after we are done composing all canons and fragments.

3. Morphisms of domain descriptions may depend on the final value typing $T \rightarrow \text{obj } \mathcal{D} T$, as the first parameter of the `morph` field of `Fragment` in Figure 5 indicates. This generalization is used to define domain descriptions whose morphisms depend on value typings. For example, to define ML-style references Section 2.5.4 we use morphisms that implicitly thread stores with values that depend on value typings.

Using these generalizations, the `Fragment` type in Figure 5 ensures that a fragment closure $\square \text{Fragment}$ will use the final canon of the language as the definition of the final set of types and the final set of value typings which objects and morphisms depend on.

By deriving \mathcal{D} from the category of Agda `Sets` (i.e. objects are types in `Set` and morphisms are Agda functions) we regain the framework as defined in Section 2.4, but now extended with the necessary infrastructure for variables and stores. It is straightforward to transport the fragments we developed in Section 2.4 to this more expressive setup.

Before we can define intrinsically-typed fragment instances in this more expressive setting, we also need to transport the definitions of isomorphism, canon union, and canon subtyping that we introduced in Section 2.3.1 and Section 2.4. The essential ingredient of these definitions is the notion of isomorphism. Lifting this notion isomorphism to relate objects of domain descriptions via morphisms of domain descriptions, all of the definitions from before translate straightforwardly. The safe upcasting and downcasting operations for translating between between the value typings of the fragment canon and the final fragment look more involved because of the switch to paramorphisms:

```

↑ : { _ : c₁ ⊆ c }
  → morph D (fold (val c))
    (val c₁ (fmap (λ t → t , fold (val c) t) t))
    (fold (val c) (inject t))
↓ : { _ : c₁ ⊆ c }
  → morph D (fold (val c))
    (fold (val c) (inject t))
    (val c₁ (fmap (λ t → t , fold (val c) t) t))

```

In the remainder of this section we illustrate how this more expressive variant of language fragments allows us to define language fragments for different language features and effects.

2.5.2 *Simply-Typed Lambda Calculus*

As our first example, we consider how to define a fragment for the simply-typed λ -calculus. We instantiate the definitions from Section 2.5.1 with a domain description based on the category of Agda `Sets`:

```

Sets : DomainDesc
obj   Sets T      = Set
morph Sets V A B = A → B

```

We use the following canon, interpreting types in `Set`:

```

data TFunShape : Set where
  fun : TFunShape

TFunΣ = TFunShape ▷ λ where
  fun → 2

funCanon : Canon
ty funCanon = TFunΣ
val funCanon (fun , (s , V) :: (t , W) :: []) = V → W
fun' : { TFunΣ ≤ σ } → (s t : μ σ) → μ σ
fun' s t = inject (fun , (s :: t :: []))

```

Here, we see the impact of using a paramorphism: the arguments to `fun` are not just replaced by their value (V/W), but paired with the original type (s/t) as well. The interpretation of function types is defined solely in terms of V and W , but we will need access to the uninterpreted recursive argument to define a language fragment for ML-style references (Section 2.5.4).

The signature `LamExprΣ` defines the three standard constructs of the λ calculus:

```

data LamExprShape { _ : TFunΣ ≲ σ }
  : List (μ σ) × μ σ → Set where
var : t ∈ Γ → LamExprShape (Γ , t)
abs : LamExprShape (Γ , fun' s t)
app : {s : μ σ} → LamExprShape (Γ , t)

LamExprΣ : { TFunΣ ≲ σ } → ISignature (List (μ σ) × μ σ)
LamExprΣ = LamExprShape ▶ λ where
  (var x) → []
  (abs {Γ}{s}{t} → (s :: Γ , t) :: [])
  (app {Γ}{t}{s} → (Γ , fun' s t) :: (Γ , s) :: [])

```

The `var` and `abs` constructors demonstrate the need for tracking a type context Γ . To reference a variable we must supply a witness $t \in \Gamma$ proving that it is in scope, and the type context of the function body is extended with the argument type s when constructing a λ -abstraction.

We then define a language fragment for the simply-typed λ -calculus as follows:

```

stlc : □ Fragment funCanon
expr stlc = LamExprΣ
interp stlc (var x , [] ) nv = fetch x nv
interp stlc (abs , e :: [] ) nv = ↑ (λ v → e (v , nv))
interp stlc (app , e1 :: e2 :: [] ) nv = ↓ (e1 nv) (e2 nv)

```

Variables are interpreted by invoking the function `fetch` : $t \in \Gamma \rightarrow \text{Env } \Gamma \rightarrow \text{fold } (\text{val } c) t$, which performs a safe lookup in the environment. Since function types are mapped to Agda functions,

we can reuse Agda’s function abstraction and application to interpret the `abs` and `app` constructors.

2.5.3 Exceptions

Next, we consider a language fragment for *safe division*, which raises an exception when the divisor is zero. In general, we can define effectful fragments by choosing a suitable *monad* [Moggi, 1989] that encapsulates the effects, and instantiating with a domain description based on the corresponding *Kleisli category* (which applies the monad to the target object of morphisms). For exceptions we use `Maybe`, which is a monad on the category `Sets`. The corresponding domain description is defined as follows:

```
MSets : DomainDesc
obj   MSets _      = Set
morph MSets _ A B = A → Maybe B
```

With this description, terms are interpreted as a function with type `Env Γ → Maybe (fold (val c) t)`. This allows us to implement an interpreter for `div`, which returns `nothing` if the divisor is `zero`, where the function `_/_` takes an (automatically inferred) proof that the divisor is greater than zero.

```
divide : □ Fragment natCanon
expr   divide = DivExprΣ
interp divide (div , m1 :: m2 :: []) nv = do
  v1 ← m1 nv >>= ↓
  v2 ← m2 nv >>= ↓
  case v2 of λ where
    zero   → nothing
    (suc n) → ↑ (v1 / suc n)
```

³¹ <https://agda.readthedocs.io/en/v2.6.2.2/language/syntactic-sugar.html>

Rather than working with values of the `Maybe` type directly, we use Agda’s `do`-notation³¹ as syntactic sugar for monadic computation, as well as the `>>=` operator which denotes monadic bind [Moggi, 1991]. Note that since the domain description is based on a Kleisli category, the result of up- and down casting

(which are defined as morphisms) is now also wrapped in a `Maybe`.

2.5.4 ML-Style References

Finally, we consider how to define a language fragment for ML-style references, based on the intrinsically-typed semantics by [Bach et al. \[2018\]](#). We interpret into a domain based on a Kleisli category generated from the description `ST`, which has store-type-indexed `Sets` as objects and index-preserving functions as morphisms. We will discuss the relevant monad shortly.

```
ST : DomainDesc
obj  ST T      = List T → Set
morph ST V P Q = { Weakenable V } → ∀[ P ⇒ Q ]
```

In the definition of `ST`, we require access to the syntax of types (T) and their interpretation (V) to define the sets of objects and morphisms. Store types are defined in terms of T , and to interpret ML-style references, we need to express the assumption that V is *weakenable*: every value that is well-typed relative to a given store, can also be typed with a *bigger* store.

Note that the `ST` description defined above admits predicates that are not monotone as objects; that is, objects are not guaranteed to be `Weakenable`. We could rectify this by requiring objects to be weakenable; i.e.:

```
ST' : DomainDesc
obj  ST' T      =
  ∃ λ (P : List T → Set) → Weakenable (const P)
morph ST' V P Q = ∀[ proj1 P ⇒ proj1 Q ]
```

However, that would clutter the resulting interpreter, which relies on dependent pattern matching on objects. For that reason, we use `ST` which has less structure, but which allows us to explicitly assume that predicates are weakenable when we need it.

```

Ref : T → List T → Set
Ref t Σ = t ∈ Σ

record Mem (M : Monad ST) : Set where
  field
    alloc   : ∀[ V t ⇒ M V (Ref t) ]
    retrieve : ∀[ Ref t ⇒ M V (V t) ]
    write   : ∀[ Ref t ⇒ V t ⇒ M V U ]

refs : { Mem M } → □ Fragment refCanon
expr refs = RefExprΣ
interp refs (init , m :: []) nv = do
  v ← m nv
  alloc v ≫≧ ↑
interp refs (deref , m :: []) nv = do
  l ← m nv ≫≧ ↓
  retrieve l
interp refs (update , m1 :: m2 :: []) nv = do
  (l , nv) ← ( m1 nv ≫≧ ↓ ) ^ nv < wk-env >
  (v , l ) ← m2 nv          ^ l   < wk-ref   >
  write l v ≫≧ ↑

```

Figure 6: Definition of the `Mem` interface, and a fragment for ML-style references

The canon for references shows why `val` needs to be a paramorphism: the interpretation of the type `ref t` is a proof of the form $t \in \Sigma$, which makes `val` non-compositional.

```

data TRefShape : Set where
  ref unit : TRefShape
TRefΣ = TRefShape ▷ λ where
  ref   → 1
  unit  → 0

refCanon : Canon
ty refCanon = TRefΣ
val refCanon (ref , (t , V) :: []) Σ = t ∈ Σ
val refCanon (unit ,          [] ) Σ = ⊤

```

Interpreting ML-style references has side effects, in the sense that the interpreter can modify a global store. To define the interpreter for ML-style references, we require a monad over store predicates that encapsulates this interaction with the store. Rather than settling on a particular monad, we keep it abstract, and assume that it satisfies the `Mem` interface (Figure 6, left), which provides the operations `alloc`, `retrieve`, and `write`. Following [Bach et al. \[2018\]](#), we also require that the chosen monad

has *tensorial strength* [Moggi, 1991], meaning it supports the following operation, where \sqcap is the usual product type lifted to predicates:

$$_ \wedge _ : \forall [M P \Rightarrow Q \Rightarrow M (P \sqcap Q)]$$

We use this operation whenever we compute a value, but need to perform more computations before we can return that value. This situation occurs, for example, in the `update` case of the interpreter in Figure 6. Since computations may change the store, it is not immediately clear that previously computed values are still typeable relative to the updated store after running these computations. The strength operation allows us to pass these values back into the monad by pairing them with a computation, updating their store typing along the way. The key to implementing strength is to assume that the store only ever increases monotonically during execution (i.e., values can be added or changed, but never deleted), and to require that strength can only be applied to values that are weakenable with respect to this ordering. The interpreter itself (Figure 6, right) is then defined in terms of the operations provided by these interfaces, and defines interpretation for `init`, `deref`, and `update` expressions, which respectively create, read from, and update a reference. Hence, we can use this fragment with any monad that satisfies the `Mem` interface and has tensorial strength.

2.5.5 Case Study

To evaluate our approach we defined a small library of fragments as a case study.³² 7 shows an overview of the `Canons`, monadic interfaces, `Fragments`, and languages that we implemented. Nodes are Agda `modules`, and dashed arrows are `imports`. For each module, we indicate the line count of the corresponding file. In addition to the `Mem` interface from Section 2.5.4, we assume three additional monadic interfaces: `Lambda` (which provides operations for function abstraction and application), `General` (which provides an operation for general recursion), and

³² The case study is available as part of the accompanying artifact [Van der Rest et al., 2022a].

Except (which provides operations for throwing and catching exceptions). It is possible to construct many more languages than shown in the figure, since any unique combination of fragments can be composed into a unique language.

The combination of ML-style references and functions means that, even without assuming the **General** interface, we can encode general recursion using Landin’s Knot. Thus, our interpreter must also be able to assign a semantics to non-terminating programs. A common technique for representing (potentially) non-terminating computations in a total language like Agda, that we also use for our case study, is to use a *fueled interpreter* [Amin and Rompf, 2017, Owens et al., 2016]. A computation that may either return an A or diverge, is represented as a function of type $\mathbb{N} \rightarrow \text{Maybe } A$ that returns a **just** if it finishes computing before running out of fuel, and **nothing** otherwise.

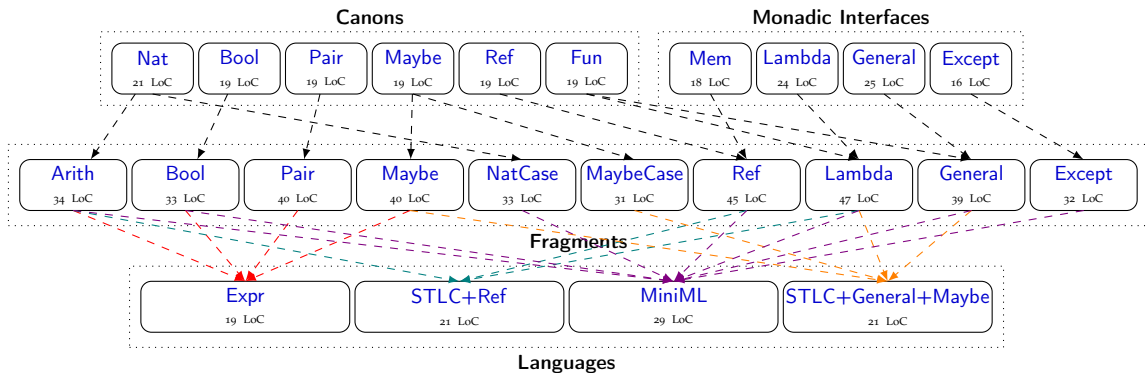


Figure 7: Overview of canons, interfaces, fragments, and languages, together with their dependencies

The type of fragment composition necessitates that we define all the fragments from this case study using the same semantic domain. To support a fragment for ML-style references, this must be the same semantic domain that we developed in Section 2.5.4, hence, we cannot use the exact fragments for the simply-typed lambda calculus and exceptions that we defined in Section 2.5.2 and Section 2.5.3. Instead we must re-define them to interpret into the new semantic domain. To extract an

interpreter for composed languages, we must also provide a monad that instantiates monadic interfaces, such as the `Mem` interface in Figure 6. Our case study achieves this by using a monad that simultaneously instantiates *all* of the interfaces in Figure 7.

The Lambda fragment maps a lambda expression to a monadic operation which accepts a monadic operation (the function body) as input, and produces a closure value as output. Choosing a monad that has *both* this operation *and* satisfies the `Mem` interface discussed in Section 2.5.4 requires some care. We cannot use Agda functions to represent function values, as this results in a *recursive domain equation*: a mutual dependency between values and stores which Agda’s termination checker will (rightly) reject. To solve this problem, we adopt a flavor of effect handler that is similar to the *latent effect handlers* of Van den Berg et al. [2021b].

2.5.6 Discussion

We set out to show that our approach to intrinsically-typed and intrinsically compositional language fragments can be transported from the rather simple semantic domain for which we presented the ingredients in Section 2.4 to a much richer semantic domain from the state-of-the-art in intrinsically-typed interpreters. The key idea of our approach is that the extrinsic proof obligations of both type safety and fragment composition can be made intrinsic to fragment well-typing. We argued that the benefits of this approach are that it helps language component developers to get their semantics right from the start, and that it reduces code size. Our case study indeed demonstrated that these benefits are upheld when we employ the same approach in the relevant semantic domain.

At the same time, the case study also showed that the specifics of particular semantic domains come with their own challenges regarding compositionality. Furthermore, reuse is limited to the chosen semantic domain, which is selected upfront. To

use the presented model for compositional fragments as the underpinning for a meta-language for DSL development, one needs to select a class of domains that is expressive enough to cover a large set of desired DSLs. Making this selection and tackling the compositionality challenges that are specific to that class of domains is not addressed by this paper.

2.6 RELATED WORK

We have presented an approach to constructing type-safe languages from composable, type-safe language fragments. As described in previous sections, our approach builds on *data types à la carte* [Swierstra, 2008]. Here, we describe other related work.

2.6.1 *Meta-Theory à la Carte.*

Our work is closely related to previous work on *Meta-Theory à la Carte* (MTC) [Delaware et al., 2013b], *Modular Monadic Meta-Theory* (3MT) [Delaware et al., 2013c], and *Generic Data Types à la Carte* (GDTC) [Keuchel and Schrijvers, 2013].

Delaware et al. [2013b] and Keuchel and Schrijvers [2013] focus on *pure* language features, and support binders by using *parameterized higher-order abstract syntax* [Chlipala, 2008]. Other than binding, they do not consider effectful language features. Delaware et al. [2013c] extend the MTC approach with effects, using monadic interfaces similar to `Mem` in Figure 6. They also construct their monads modularly using *monad transformers* [Liang et al., 1995b]. We did not build our monad modularly, as that would require composing different domain descriptions, which is an open question.

We need on average 37 LoC (counted with Al Danial’s `cloc` tool) to implement verified language fragments. If we also count lines of code for definitions of canonical forms and monadic interfaces, this number approximately doubles. In comparison, MTC and GDTC report needing respectively on average 1100

LoC and 1050 LoC to define and verify similar language features. In other words, the difference in code size between our intrinsically-typed approach and the extrinsic approach found in previous work is about an order of magnitude.

Our framework code is also more concise: we use 859 LoC, whereas MTC uses 2500 LoC, GDTC uses 3500 LoC, and 3MT uses 4400 LoC. We use the Agda Standard Library for lists, relations over lists, functions for working with functions and predicates, and more. Previous works also seem to use Coq’s standard library, but perhaps to a lesser extent.

Some of the difference in code sizes can be ascribed to code that deals with “wrong” cases, since these cases are absent in intrinsically-typed interpreters. Our language fragment abstraction also saves code when composing languages. This abstraction does not exist in the MTC frameworks, where language features are the sum of their parts (i.e., types, expressions, values, type system, interpreter, and type safety proof). In our framework all of these parts are summed using a single operation: fragment composition.

2.6.2 *Generic Programming and Meta-Theory*

A key technique in our work is to encode data type descriptions as signatures. The [Signature](#) type we used as our data type description is known as a *container* [[Abbott et al., 2005a](#), [Altenkirch et al., 2015](#)]. The *universe of syntaxes* encoding [[Chapman et al., 2010b](#), [Dagand, 2013b](#)] (which we will call just “universe encoding”) encodes data type descriptions in a different way but encodes the same class of data types as containers do. The difference is in how the encoding is defined: a universe encoding is akin to the syntactic representation of a grammar, whereas the container encoding corresponds to a bag of symbols with associated arities. For generic programming applications that operate on the *syntax of data types itself*, the universe encoding is often preferred. For example, *ornamentation* is a technique

for “decorating” data types with additional structure [McBride, 2011, Dagand and McBride, 2014, Ko and Gibbons, 2017, Dagand, 2017].

Allais et al. [2018] use the universe encoding to implement a generic framework for *syntaxes with variable binding*, enabling binding-aware generic programming, with generic correctness guarantees. Whereas our paper focuses on the problem of defining an intrinsically-typed interpreter in a modular way, the work of Allais et al. [2018] suggests a promising direction for meta-theoretical reasoning about object languages and interpreters defined using similar generic programming techniques as our framework. Whether it is possible to apply the generic meta-theoretical reasoning techniques of Allais et al. [2018] to the style of modular intrinsically-typed interpreters we develop in this paper is an interesting question for future work.

Final tag-less interpreters [Carette et al., 2009b] represent a different flavor of generic programming for implementing type safe definitional interpreters. In particular Carette et al. [2009b] showed that it is possible to implement type safe interpreters by modeling object language expression constructors as an abstract interface which, behind the scenes, is constructing a meta-language program that corresponds to interpreting the object language expression. The resulting interpretation of the object language expression is type safe because meta-language programs are type safe. *Object algebras* [d. S. Oliveira and Cook, 2012] are based on a similar idea, but for object oriented meta-languages. Bahr and Hvitved [2012b] extend the final tag-less approach to intrinsically-typed definitional interpreters in Haskell. Their intrinsically-typed interpreter does not treat value types as open, but does support open object language types, by using Haskell types as object language types. A similar result was achieved by Parreaux et al. [2019] in Scala. It is unclear how the final tag-less approach can be used to define type safe semantics of languages with effects that are not built into the meta-language.

2.6.3 *Other Approaches to Modular Semantics and their Proofs*

[Cimini et al. \[2020\]](#) extrinsically verify type safety of language specifications by developing a meta type-system for type-safe language specifications. A (meta) type-correct specification yields a type-safe semantics. Language definitions are monolithic, and cannot be constructed from separately checked fragments. Their approach uses reduction semantics, which has some modularity issues. For example, if we add support for exception handlers, we need to copy-paste the current reduction context definition to express a context up-to the closest handler. On the other hand, their approach make it easy to verify type safety of new extensions, since that is done automatically by their meta-language type checker.

[Schwaab and Siek \[2013\]](#) present an Agda formalization of a modular extrinsic type safety proof for a small-step operational semantics defined using a variant [\[Norell, 2008\]](#) of the universe encoding due to [Chapman et al. \[2010b\]](#), [Dagand \[2013b\]](#) discussed above. The approach is closely related to the extrinsic modular proofs found in MTC [\[Delaware et al., 2013c\]](#), but is based on modular progress/preservation lemmas about a small-step transition relation. A direct comparison with MTC is difficult because of the difference in meta- and object language ([Schwaab and Siek \[2013\]](#) encode a simple object language with just numbers and lists), but we expect that progress/preservation style requires more LOC per feature. On the other hand, progress/preservation lemmas are a time tested paradigm for type safety proofs. As shown by [Wadler et al. \[2020\]](#), it is possible to marry the intrinsically-typed approach with the small-step style. We expect that it is possible to modularize intrinsically-typed small-step semantics using similar techniques as we use in our framework, but leave verification of this expectation for future research.

Modular Structural Operational Semantics (MSOS) [\[Mosses, 2004\]](#) is a framework for modularly specifying small-step operational semantics. [Madlener et al. \[2011\]](#) show how to implement

this framework in the Coq proof assistant, and how to do modular proofs about the small-step transition relation. In a related line of work, [Torrini and Schrijvers \[2015\]](#) describe a different method for modular proofs in Coq, and provide, as case study, a modular extrinsic type safety proof about a modularly-specified small-step transition relation. [Churchill et al. \[2015\]](#) use an extension of MSOS [[Churchill and Mosses, 2013](#)] to modularly specify the static and dynamic semantics of a collection of individual fundamental programming constructs, but do not establish type safety.

Interaction Trees [[Xia et al., 2020](#)] are a general purpose inductive data structure for representing effectful and diverging computations based on the *freer monad* [[Kiselyov et al., 2013](#), [Kiselyov and Ishii, 2015](#)]. As such they support modular reasoning by defining the “visible events” (i.e., effects) of a computation as a functor signature, with the type of interaction trees being closed under the co-product of these signatures. In later work, [Zakowski et al. \[2021\]](#) used interaction trees to give a formal semantics for LLVM IR. The flavor of effect handlers that we used to define a semantics for MiniML in Section 2.5.5 bears some semblance to interaction trees, but to explore their relation in more detail is a subject of further study.

In his popular textbook on *Types and Programming Languages*, [Pierce \[2002\]](#) makes use of *canonical forms* lemmas to make type safety proofs *robust* as new language fragments are gradually introduced. He attributes the idea of using canonical forms lemmas to Bob Harper (no citation given). We have shown how to define interpreters and language fragments in a way that guarantees that the subset of values that an interpreter- or language fragment knows about is guaranteed to be canonical, essentially making language fragments *robust by construction*.

2.7 CONCLUSION

In this paper, we presented a framework for defining composable and safe by construction language fragments that can be

checked in isolation and safely reused to build type safe languages. This makes it easier to develop and reuse interpreters that do not go wrong, and reduces the overhead traditionally associated with modular verification of type safety. This makes mechanized meta-theory available to a wider audience of DSL developers.

POSTSCRIPT

The contributions presented in this chapter indeed give us a handle on understanding how the type safety invariant encoded in intrinsically-typed interpreters can be maintained under composition. While this is an important first step towards employing intrinsically-typed interpreters for defining reusable language components, there is still work ahead. As pointed out in Section 2.5.6, language constructs may impose different requirements on the semantic domain they denote into, and it is still an open question how to mediate these different requirements when composing fragments.

Differences in the semantic domain of fragments arise, for the most part, because their interpreters exhibit different *side effects*. A modular treatment of these side effects is thus a key ingredient for composing fragments with different domains. In the next chapter, we will work towards expanding the set of side effects that can be dealt with modularly.



HEFTY ALGEBRAS: MODULAR ELABORATIONS AND REASONING FOR PROGRAMS WITH HIGHER-ORDER EFFECTS

PREFACE

A key question left open at the end of last chapter was how to mediate language fragments with different semantic domains. By defining the semantics of fragments in terms of operations of some abstract monad, we reduced the problem to finding a structured and modular approach to constructing these monads. Although algebraic effects and handlers [Plotkin and Pretnar, 2009a] provide a principled method for defining the syntax and implementation of abstract monads, the approach is not expressive enough to construct the monad(s) we need to implement the language fragments presented in Section 2.5.

There are two main reasons why algebraic effects and handlers are not sufficient. The first is that they are, by definition, restricted to describing the syntax of monads that only have first-order operations. That is, they cannot describe monadic operations which take computations as an argument, precluding many common operations such as exception catching. Several of the fragments shown in Figure 7, such as `Lambda` and `Except`, are implemented in terms of such higher-order operations. The second reason is that fragments may be denoted into a monad on a category with more structure than Agda's `Set`. The `Ref` fragment in Figure 7, for example, requires a monad in a category

of monotone predicates to enforce well-formedness of references. This chapter contributes a new approach to defining the syntax and semantics of monads with higher-order operations by elaborating them into algebraic effects.

3.1 INTRODUCTION

Defining abstractions for programming with side effects is a research question with a long and rich history. The goal is to define an interface of (possibly) side effecting operations where the interface encapsulates and hides irrelevant operational details about the operations and their side effects. Such encapsulation makes it easy to refactor, optimize, or even change the behavior of a program, by changing the implementation of the interface.

Monads [Moggi, 1989] have long been the preferred solution to this research question. However, *algebraic effects and handlers* [Plotkin and Pretnar, 2009b] are emerging as an attractive alternative solution, due to the modularity benefits that they provide. However, these modularity benefits do not apply to many common operations that take computations as arguments.

3.1.1 Background: Algebraic Effects and Handlers

To understand the benefits of algebraic effects and handlers and the modularity problem with operations that take computations as parameters, we give a brief introduction to algebraic effects, based on the effect handlers tutorial by Pretnar [2015]. Readers familiar with algebraic effects and handlers are encouraged to skim the code examples in this subsection and read its final paragraph.

Consider a simple operation *out* for output which takes a string as argument and returns the unit value. Using algebraic effects and handlers its type is:

$$out : String \rightarrow ()!Output$$

Here *Output* is the *effect* of the operation. In general $A! \Delta$ is a computation type where A is the return type and Δ is a *row* (i.e., unordered sequence) of *effects*, where an *effect* is a label associated with a set of operations. A computation of type $A! \Delta$ may *only* use operations associated with an effect in Δ . An effect can generally be associated with multiple operations (but not the other way around); however, the simple *Output* effect that we consider is only associated with the operation *out*. Thus $()! \text{Output}$ is the type of a computation which may call the *out* operation.

We can think of *Output* as an interface that specifies the parameter and return type of *out*. The implementation of such an interface is given by an *effect handler*. An effect handler defines how to interpret operations in the execution context they occur in. The type of an effect handler is $A! \Delta \Rightarrow B! \Delta'$, where Δ is the row of effects before applying the handler and Δ' is the row after. For example, here is the type of an effect handler for *Output*:

$$hOut : A! \text{Output}, \Delta \Rightarrow (A \times \text{String})! \Delta$$

The *Output* effect is being handled, so it is only present in the effect row on the left.³³ As the type suggests, this handler handles *out* operations by accumulating a string of output. Below is the handler of this type:

$$\begin{aligned} hOut = \mathbf{handler} \{ & (\mathbf{return} \ x) \mapsto \mathbf{return} \ (x, "") \\ & (\mathit{out} \ s; k) \mapsto \mathbf{do} \ (y, s') \leftarrow k \ (); \\ & \mathbf{return} \ (y, s \ ++ \ s') \} \end{aligned}$$

The **return** case of the handler says that, if the computation being handled terminates normally with a value x , then we return a pair of x and the empty string. The case for *out* binds a variable s for the string argument of the operation, but also a variable k representing the *execution context* (or *continuation*). Invoking an operation suspends the program and its execution context up-to the nearest handler of the operation. The handler

³³ *Output* could occur in Δ too. This raises the question: which *Output* effect does a given handler actually handle? We refer to the literature for answers to this question; see, e.g., the row treatment of Morris and McKinna [2019b], the effect lifting of Biernacki et al. [2018], and the effect tunneling of Zhang and Myers [2019].

can choose to re-invoke the suspended execution context (possibly multiple times). The handler case for *out* above always invokes *k* once. Since *k* represents an execution context that includes the current handler, calling *k* gives a pair of a value *y* and a string *s'*, representing the final value and output of the execution context. The result of handling *out s* is then *y* and the current output (*s*) plus the output of the rest of the program (*s'*).

In general, a computation $m : A! \Delta$ can only be run in a context that provides handlers for each effect in Δ . To this end, if $\Delta = \Delta_1, \Delta_2$ and $h : A! \Delta_1 \Rightarrow B! \Delta'_1$, then the expression **(with** *h* **handle** *m*) : $B! \Delta'_1, \Delta_2$ runs *m* in the context of the handler *h*. For example, consider:

```
hello : ()! Output
hello = out "Hello"; out " world!"
```

Using this, we can run *hello* in a scope with the handler *hOut* to compute the following result:

```
(with hOut handle hello)  $\equiv$  ((), "Hello world!")
```

An attractive feature of algebraic effects and handlers is that programs such as *hello* are defined *independently* of how the effectful operations they use are implemented. This makes it possible to refine, refactor, or even change the meaning of operations without having to modify the programs that use them. For example, we can refine the meaning of *out* *without modifying the hello program*, by using a different handler *hOut'* which prints output to the console. However, some operations are challenging to express in a way that provides these modularity benefits.

3.1.2 The Modularity Problem with Higher-Order Operations

Algebraic effects and handlers provide limited support for operations that accept computations as arguments (sometimes called *higher-order operations*). The limitation is subtle but follows from

how handler cases are typed. Following Plotkin and Pretnar [2009b], Pretnar [2015], the left and right hand sides of handler cases are typed as follows:

$$\mathbf{handler} \{ \dots (op \underbrace{v}_{A} ; \underbrace{k}_{B \rightarrow C! \Delta'}) \mapsto \underbrace{c}_{C! \Delta'}, \dots \}$$

Here it is only k whose type is compatible with the right hand side. In theory, the parameter type v would also be compatible if $A = C! \Delta'$. However, encoding computations as parameters in this way is non-modular. The reason is that effect handlers are not applied recursively to parameters of operations [Plotkin and Pretnar, 2009b, Pretnar, 2015]; i.e., if h handles operations other than op , then

$$\begin{aligned} & \mathbf{with\ h\ handle} (\mathbf{do\ } x \leftarrow op\ v; m) \\ \equiv & \mathbf{do\ } x \leftarrow op\ v; (\mathbf{with\ h\ handle\ } m) \end{aligned}$$

This implies that the only way to ensure that v has type $A = C! \Delta'$ whose effects match the context of the operation (e.g., $k : B \rightarrow C! \Delta'$), is to apply handlers of higher-order effect encodings (such as op) before applying other handlers (such as h). In turn, this means that programs can contain at most one higher-order effect encoded in this way (otherwise, which handler do we apply first?). Consequently, encoding computation parameters in terms of the value v carried by an operation does not support modular definition, composition, and handling of higher-order effects.

A consequence of this restriction is that algebraic effects and handlers only support higher-order operations whose computation parameters are *continuation-like*. In particular, for any operation $op : A! \Delta \rightarrow \dots \rightarrow A! \Delta \rightarrow A! \Delta$ and any m_1, \dots, m_n and k ,

$$\begin{aligned} & \mathbf{do\ } x \leftarrow (op\ m_1 \dots m_n); k\ x \\ \equiv & op\ (\mathbf{do\ } x_1 \leftarrow m_1; k\ x_1) \dots (\mathbf{do\ } x_n \leftarrow m_n; k\ x_n) \quad (\dagger) \end{aligned}$$

This property, known as the *algebraicity property* [Plotkin and Power, 2003], says that the computation parameter m_1, \dots, m_n

are only ever run in a way that *directly* passes control to k . Such operations can without loss of generality or modularity be encoded as operations *without computation parameters*; e.g., $op\ m_1 \dots m_n = \mathbf{do}\ x \leftarrow op' \ ();\ select\ x$ where $op' : () \rightarrow D^n ! \Delta$ and $select : D^n \rightarrow A ! \Delta$ is a function that chooses between n different computations using a data type D^n whose constructors are d_1, \dots, d_n such that $select\ d_i = m_i$ for $i = 1..n$. Some higher-order operations obey the algebraicity property; many do not. Examples of operations that do not include:

- Exception handling: let $catch\ m_1\ m_2$ be an operation that handles exceptions thrown during evaluation of computation m_1 by running m_2 instead, and $throw$ be an operation that throws an exception. These operations are not algebraic. For example,

$$\mathbf{do}\ (catch\ m_1\ m_2);\ throw \not\equiv catch\ (\mathbf{do}\ m_1;\ throw)\ (\mathbf{do}\ m_2;\ throw)$$

- Local binding (the *reader monad* [Jones, 1995]): let ask be an operation that reads a local binding, and $local\ r\ m$ be an operation that makes r the current binding in computation m . Observe:

$$\mathbf{do}\ (local\ r\ m);\ ask \not\equiv local\ r\ (\mathbf{do}\ m;\ ask)$$

- Logging with filtering (an extension of the *writer monad* by Jones [1995]): let $out\ s$ be an operation for logging a string, and $cancel\ f\ m$ be an operation for post-processing the output of computation m by applying $f : String \rightarrow String$.³⁴ Observe:

$$\mathbf{do}\ (cancel\ f\ m);\ out\ s \not\equiv cancel\ f\ (\mathbf{do}\ m;\ out\ s)$$

³⁴ The *cancel* operation is a variant of the function by the same name the widely used Haskell `mtl` library.

It is, however, possible to elaborate higher-order operations into more primitive effects and handlers. For example, *sensor* can be elaborated into an inline handler application of *hOut*:

```

sensor : (String → String) → A!Output, Δ → A!Output, Δ
sensor f m = do (x, s) ← (with hOut handle m);
               out (f s);
               return x

```

The other higher-order operations above can be defined in a similar manner.

Elaborating higher-order operations into standard algebraic effects and handlers as illustrated above is a key use case that effect handlers were designed for [Plotkin and Pretnar, 2009b]. However, elaborating operations in this way means the operations are not a part of any effect interface. So, unlike plain algebraic operations, the only way to refactor, optimize, or change the semantics of higher-order operations defined in this way is to modify or copy code. In other words, we forfeit one of the key attractive modularity features of algebraic effects and handlers.

This modularity problem with higher-order effects (i.e., effects with higher-order operations) was first observed by Wu et al. [2014] who proposed *scoped effects and handlers* [Wu et al., 2014, Piróg et al., 2018, Yang et al., 2022] as a solution. Scoped effects and handlers have similar modularity benefits as algebraic effects and handlers, but works for a wider class of effects, including many higher-order effects. However, Van den Berg et al. [2021a] recently observed that operations that defer computation, such as evaluation strategies for λ application or (*multi-*)*staging* [Taha and Sheard, 2000], are beyond the expressiveness of scoped effects. Therefore, Van den Berg et al. [2021a] introduced another flavor of effects and handlers that they call *latent effects and handlers*.

In this paper we present a (surprisingly) simple alternative solution to the modularity problem with higher-order effects, which only uses standard effects and handlers and off-the-shelf

generic programming techniques known from, e.g., *data types à la carte* [Swierstra, 2008].

3.1.3 Solving the Modularity Problem: Elaboration Algebras

We propose to define elaborations such as *sensor* (Section 3.1.2) in a modular way. To this end, we introduce a new type of *computations with higher-order effects* which can be modularly elaborated into computations with only standard algebraic effects:

$$A !! H \xrightarrow{\text{elaborate}} A ! \Delta \xrightarrow{\text{handle}} \text{Result}$$

Here $A !! H$ is a computation type where A is a return type and H is a row comprising both algebraic and higher-order effects. The idea is that the higher-order effects in the row H are modularly elaborated into the row Δ . To achieve this, we define *elaborate* such that it can be modularly composed from separately defined elaboration cases, which we call elaboration *algebras* (for reasons we explain in Section 3.3). Using $A !! H \Rightarrow A ! \Delta$ as the type of elaboration algebras that elaborate the higher-order effects in H to Δ , we can modularly compose any pair of elaboration algebras $e_1 : A !! H_1 \Rightarrow A ! \Delta$ and $e_2 : A !! H_2 \Rightarrow A ! \Delta$ into an algebra $e_{12} : A !! H_1, H_2 \Rightarrow A ! \Delta$.³⁵

³⁵ Readers familiar with *data types à la carte* [Swierstra, 2008] may recognize this as algebra composition.

Elaboration algebras are as simple to define as non-modular elaborations such as *sensor* (Section 3.1.2). For example, here is the elaboration algebra for the higher-order *Censor* effect whose only associated operation is the higher-order operation $\text{censor}_{op} : (\text{String} \rightarrow \text{String}) \rightarrow A !! H \rightarrow A !! H$:

$$\begin{aligned} eCensor : A !! Censor &\Rightarrow A ! \text{Output}, \Delta \\ eCensor (\text{censor}_{op} f m; k) &= \mathbf{do} (x, s) \leftarrow (\mathbf{with} \text{hOut} \mathbf{handle} m); \\ &\quad \text{out} (f s); k x \end{aligned}$$

The implementation of $eCensor$ is essentially the same as *sensor*. There are two main differences. First, elaboration happens in-context, so the value yielded by the elaboration is passed to

the context (or continuation) k . Second, and most importantly, programs that use the $sensor_{op}$ operation are now programmed against the interface given by $Censor$, meaning programs do not (and *cannot*) make assumptions about how $sensor_{op}$ is elaborated. As a consequence, we can modularly refine the elaboration of higher-order operations such as $sensor_{op}$, without modifying the programs that use the operations. For example, the following program censors and replaces “Hello” with “Goodbye”:³⁶

```

sensorHello : () !! Censor, Output
sensorHello = sensorop
  (λs. if (s ≡ “Hello”) then “Goodbye” else s)
  hello

```

³⁶ This program relies on the fact that it is generally possible to lift computation $A! \Delta$ to $A! H$ when $\Delta \subseteq H$.

Say we have a handler $hOut' : (String \rightarrow String) \rightarrow A! Output, \Delta \Rightarrow (A \times String)! \Delta$ which handles each operation $out\ s$ by applying a censor function $(String \rightarrow String)$ to s before emitting it. Using this handler, we can give an alternative elaboration of $sensor_{op}$ which post-processes output strings *individually*:

```

eCensor' : A !! Censor  $\Rightarrow$  A ! Output,  $\Delta$ 
eCensor' (sensorop f m; k) = do x  $\leftarrow$  (with hOut' f handle m);
  out s; k x

```

In contrast, $eCensor$ applies the censoring function $(String \rightarrow String)$ to the batch output of the computation argument of a $sensor_{op}$ operation. The batch output of $hello$ is “Hello world!” which is unequal to “Hello”, so $eCensor$ leaves the string unchanged. On the other hand, $eCensor'$ censors the individually output “Hello”:

```

with hOut handle (with eCensor elaborate sensorHello)
   $\equiv$  ((), “Hello world!”)

```

```

with hOut handle (with eCensor' elaborate sensorHello)
   $\equiv$  ((), “Goodbye world!”)

```

Higher-order operations now have the same modularity benefits as algebraic operations.

3.1.4 Contributions

This paper formalizes the ideas sketched in this introduction by shallowly embedding them in Agda. However, the ideas transcend Agda. Similar shallow embeddings can be implemented in other dependently typed languages, such as Idris [Brady, 2013b]; but also in less dependently typed languages like Haskell, OCaml, or Scala.³⁷ By working in a dependently typed language we can state algebraic laws about interfaces of effectful operations, and prove that implementations of the interfaces respect the laws. We make the following technical contributions:

³⁷ The artifact accompanying this paper [Van der Rest and Bach, 2024] contains a shallow embedding of elaboration algebras in Haskell.

- Section 5.2.3 describes how to encode algebraic effects in Agda, revisits the modularity problem with higher-order operations, and summarizes how scoped effects and handlers address the modularity problem, for some (*scoped* operations) but not all higher-order operations.
- Section 3.3 presents our solution to the modularity problem with higher-order operations. Our solution is to (1) type programs as *higher-order effect trees* (which we dub *hefty trees*), and (2) build modular elaboration algebras for folding hefty trees into algebraic effect trees and handlers. The computations of type $A !! H$ discussed in Section 3.1.3 correspond to hefty trees, and the elaborations of type $A !! H \Rightarrow A ! \Delta$ correspond to hefty algebras.
- Section 3.4 presents examples of how to define hefty algebras for common higher-order effects from the literature on effect handlers.
- Section 3.5 shows that hefty algebras support formal and modular reasoning on a par with algebraic effects and handlers, by developing reasoning infrastructure that supports verification of equational laws for higher-order effects such as exception catching. Crucially, proofs of correctness of elaborations are compositional. When composing two proven correct elaboration, correctness of the

combined elaboration follows immediately without requiring further proof work.

Section 5.6 discusses related work and Section 5.7 concludes. An artifact containing the code of the paper and a Haskell embedding of the same ideas is available online [Van der Rest and Bach, 2024]. A subset of the contributions of this paper were previously published in a conference paper [Bach and Van der Rest, 2023]. While that version of the paper too discusses reasoning about higher-order effects, the correctness proofs were non-modular, in that they make assumptions about the order in which the algebraic effects implementing a higher-order effect are handled. When combining elaborations, these assumptions are often incompatible, meaning that correctness proofs for the individual elaborations do not transfer to the combined elaboration. As a result, one would have to re-prove correctness for every combination of elaborations. For this extended version, we developed reasoning infrastructure to support modular reasoning about higher-order effects in Section 3.5, and proved that correctness of elaborations is preserved under composition of elaborations.

3.2 ALGEBRAIC EFFECTS AND HANDLERS IN AGDA

This section describes how to encode algebraic effects and handlers in Agda. We do not assume familiarity with Agda and explain Agda specific notation in sidenotes. Sections 3.2.1 to 3.2.4 defines algebraic effects and handlers; Section 3.2.5 revisits the problem of defining higher-order effects using algebraic effects and handlers; and Section 3.2.6 discusses how scoped effects [Wu et al., 2014, Piróg et al., 2018, Yang et al., 2022] solves the problem for *scoped* operations but not all higher-order operations.

3.2.1 Algebraic Effects and The Free Monad

³⁸ `Set` is the type of types in Agda. More generally, functors mediate between different categories. For simplicity, this paper only considers endofunctors on `Set`, where an endofunctor is a functor whose domain and codomain coincides; e.g., `Set → Set`.

³⁹ To show that this is truly a functor, we should also prove that `map-sig` satisfies the functor laws. We will not make use of these functor laws in this paper, so we omit them.

⁴⁰ The `_⊕_` function uses copattern matching: <https://agda.readthedocs.io/en/v2.6.2.2/language/copatterns.html>. The `Op` line defines how to compute the `Op` field of the record produced by the function; and similarly for the `Ret` line.

⁴¹ `_⊔_` is a disjoint sum type from the Agda standard library. The `[_,_]` function is its eliminator.

We encode algebraic effects in Agda by representing computations as an abstract syntax tree given by the *free monad* over an *effect signature*. Such effect signatures are traditionally [Awodey, 2010, Swierstra, 2008, Kiselyov and Ishii, 2015, Wu et al., 2014, Kammar et al., 2013] given by a *functor*; i.e., a type of kind `Set → Set` together with a (lawful) mapping function.³⁸ In our Agda implementation, effect signature functors are defined by giving a *container* [Abbott et al., 2003, 2005a]. Each container corresponds to a value of type `Set → Set` that is both *strictly positive* and *universe consistent* [Martin-Löf, 1984], meaning they are a constructive approximation of endofunctors on `Set`. Effect signatures are given by a (dependent) record type:

```
record Effect : Set₁ where
  field Op : Set
       Ret : Op → Set
```

Here, `Op` is the set of operations, and `Ret` defines the *return type* for each operation in the set `Op`. The extension of an effect signature, `[[_]]`, reflects its input of type `Effect` as a value of type `Set → Set`:

```
[[_]] : Effect → Set → Set
[[ Δ ]] X = Σ (Op Δ) λ op → Ret Δ op → X
```

The extension of an effect `Δ` into `Set → Set` is indeed a functor, as witnessed by the following function:³⁹

```
map-sig : (X → Y) → [[ Δ ]] X → [[ Δ ]] Y
map-sig f (op , k) = ( op , f ∘ k )
```

As discussed in the introduction, computations may use multiple different effects. Effect signatures are closed under coproducts:^{40 41}

```
_⊕_ : Effect → Effect → Effect
Op (Δ₁ ⊕ Δ₂) = Op Δ₁ ⊔ Op Δ₂
Ret (Δ₁ ⊕ Δ₂) = [ Ret Δ₁ , Ret Δ₂ ]
```

We compute the co-product of two effect signatures by taking the disjoint sum of their operations and combining the return type mappings pointwise. We co-products to encode effect rows. For example, The effect $\Delta_1 \oplus \Delta_2$ corresponds to the row union denoted as Δ_1, Δ_2 in the introduction.

The syntax of computations with effects Δ is given by the free monad over Δ . We encode the free monad as follows:

```
data Free ( $\Delta$  : Effect) (A : Set) : Set where
  pure   : A                 $\rightarrow$  Free  $\Delta$  A
  impure : [  $\Delta$  ] (Free  $\Delta$  A)  $\rightarrow$  Free  $\Delta$  A
```

Here, `pure` is a computation with no side effects, whereas `impure` is an operation whose syntax is given by the functor `[Δ]`. By applying this functor to `Free Δ A`, we encode an operation whose *continuation* may contain more effectful operations.⁴² To see in what sense, let us consider an example.

EXAMPLE. The following data type defines an operation for outputting a string. Below it is its corresponding effect signature.

```
data OutOp : Set where
  out : String  $\rightarrow$  OutOp

Output : Effect
Op Output      = OutOp
Ret Output (out s) =  $\top$ 
```

The effect signature says that `out` returns a unit value (\top is the unit type). Using this, we can write a simple hello world corresponding to the *hello* program from Section 3.1:

```
hello : Free Output  $\top$ 
hello = impure
  ( out "Hello"
    ,  $\lambda$  _  $\rightarrow$  impure (out " world!" ,  $\lambda$  x  $\rightarrow$  pure x)
```

⁴² By unfolding the definition of `[Δ]` one can see that our definition of the free monad is identical to the I/O trees of Hancock and Setzer [2000], or the so-called freer monad of Kiselyov and Ishii [2015].

Section 3.2.1 shows how to make this program more readable by using monadic `do` notation.

The `hello` program above makes use of just a single effect. Say we want to use another effect, `Throw`, with a single operation, `throw`, which represents throwing an exception (therefore having the empty type \perp as its return type):

```
data ThrowOp : Set where
  throw : ThrowOp

Throw : Effect
Op Throw = ThrowOp
Ret Throw throw =  $\perp$ 
```

Programs that use multiple effects, such as `Output` and `Throw`, are unnecessarily verbose. For example, consider the following program which prints two strings before throwing an exception:⁴³

⁴³ `\perp -elim` is the eliminator for the empty type, encoding the principle of explosion: `\perp -elim : $\perp \rightarrow A$.`

```
hello-throw : Free (Output  $\oplus$  Throw) A
hello-throw = impure (inj1 (out "Hello"),  $\lambda$  _  $\rightarrow$ 
  impure (inj1 (out " world!"),  $\lambda$  _  $\rightarrow$ 
    impure (inj2 throw,  $\perp$ -elim)))
```

To reduce syntactic overhead, we use *row insertions* and *smart constructors* [Swierstra, 2008].

3.2.2 Row Insertions and Smart Constructors

A *smart constructor* constructs an effectful computation comprising a single operation. The type of this computation is polymorphic in what other effects the computation has. For example, the type of a smart constructor for the `out` effect is:

```
`out : { Output  $\lesssim$   $\Delta$  }  $\rightarrow$  String  $\rightarrow$  Free  $\Delta$  T
```

Here, the `{ Output \lesssim Δ }` type declares the row insertion witness as an *instance argument* of ``out`. Instance arguments in Agda are conceptually similar to type class constraints in Haskell: when

we call `\out`, Agda will attempt to automatically find a witness of the right type, and implicitly pass this as an argument [Devriese and Piessens, 2011]. Thus, calling `\out` will automatically inject the `Output` effect into some larger effect row Δ .

We define the \lesssim order on effect rows in terms of a different $\Delta_1 \bullet \Delta_2 \approx \Delta$ which witnesses that any operation of Δ is isomorphic to *either* an operation of Δ_1 *or* an operation of Δ_2 :⁴⁴

```
record _•_≈_ (Δ1 Δ2 Δ : Effect) : Set1 where
  field reorder : ∀ {X} → [ Δ1 ⊕ Δ2 ] X ↔ [ Δ ] X
```

Using this, the \lesssim order is defined as follows:

```
_≲_ : (Δ1 Δ2 : Effect) → Set1
Δ1 ≲ Δ2 = ∃ λ Δ' → Δ1 • Δ' ≈ Δ2
```

It is straightforward to show that \lesssim is a *preorder*; i.e., that it is a *reflexive* and *transitive* relation.

We can also define the following function, which uses a $\Delta_1 \lesssim \Delta_2$ witness to coerce an operation of effect type Δ_1 into an operation of some larger effect type Δ_2 .⁴⁵

```
inj : { Δ1 ≲ Δ2 } → [ Δ1 ] A → [ Δ2 ] A
inj { _ , w } (c , k) = w .reorder .to (inj1 c , k)
```

Furthermore, we can freely coerce the operations of a computation from one effect row type to a different effect row type:⁴⁶

⁴⁷

```
hmap-free : ∀ [ [ Δ1 ] ⇒ [ Δ2 ] ] → ∀ [ Free Δ1 ⇒ Free Δ2 ]
hmap-free θ (pure x)           = pure x
hmap-free θ (impure (c , k)) = impure (θ (c , hmap-free θ ◦ k))
```

Using this infrastructure, we can now implement a generic `inject` function which lets us define smart constructors for operations such as the `out` operation discussed in the previous subsection.

```
inject : { Δ1 ≲ Δ2 } → Free Δ1 A → Free Δ2 A
inject = hmap-free inj

\out : { Output ≲ Δ } → String → Free Δ T
\out s = inject (impure (out s , pure))
```

⁴⁴ Here \leftrightarrow is the type of an isomorphism on `Set` from the Agda Standard Library. It is given by a record with two fields: the `to` field represents the \rightarrow direction of the isomorphism, and `from` field represents the \leftarrow direction of the isomorphism.

⁴⁵ The dot notation `w .reorder` projects the `reorder` field of the record `w`.

⁴⁶ The notation $\forall [_]$ is from the Agda Standard library, and is defined as follows:

```
∀ [ P ] = ∀ x → P x.
```

⁴⁷ We can think of the `hmap-free` function as a “higher-order” map for `Free`: given a natural transformation between (the extension of) signatures, we can transform the signature of a computation. This amounts to the observation that `Free` is a functor over the category of containers and container morphisms; assuming `hmap-free` preserves naturality.

3.2.3 Fold and Monadic Bind for *Free*

Since *Free* Δ is a monad, we can sequence computations using *monadic bind*, which is naturally defined in terms of the fold over *Free*.

$$\begin{aligned} \text{fold} &: (A \rightarrow B) \rightarrow \text{Alg } \Delta B \rightarrow \text{Free } \Delta A \rightarrow B \\ \text{fold } g \ a \ (\text{pure } x) &= g \ x \\ \text{fold } g \ a \ (\text{impure } (op \ , \ k)) &= a \ (op \ , \ \text{fold } g \ a \ \circ \ k) \end{aligned}$$

$$\begin{aligned} \text{Alg} &: (\Delta : \text{Effect}) (A : \text{Set}) \rightarrow \text{Set} \\ \text{Alg } \Delta \ A &= \llbracket \Delta \rrbracket A \rightarrow A \end{aligned}$$

Besides the input computation to be folded (last parameter), the fold is parameterized by a function $A \rightarrow B$ (first parameter) which folds a *pure* computation, and an *algebra* $\text{Alg } \Delta A$ (second parameter) which folds an *impure* computation. We call the latter an algebra because it corresponds to an F-algebra [Arbib and Manes, 1975, Pierce, 1991] over the signature functor of Δ , denoted F_Δ . That is, a tuple (A, α) where A is an object called the *carrier* of the algebra, and α a morphism $F_\Delta(A) \rightarrow A$. Using *fold*, monadic bind for the free monad is defined as follows:

$$\begin{aligned} _ \gg\! = _ &: \text{Free } \Delta A \rightarrow (A \rightarrow \text{Free } \Delta B) \rightarrow \text{Free } \Delta B \\ m \gg\! = g &= \text{fold } g \ \text{impure } m \end{aligned}$$

Intuitively, $m \gg\! = g$ concatenates g to all the leaves in the computation m .

EXAMPLE. The following defines a smart constructor for *throw*:

$$\text{\`throw} : \{ \text{Throw } \lesssim \Delta \} \rightarrow \text{Free } \Delta A$$

Using this and the definition of $\gg\! =$ above, we can use **do**-notation in Agda to make the *hello-throw* program from Section 3.2.1 more readable:

$$\begin{aligned} \text{hello-throw}_1 &: \{ \text{Output } \lesssim \Delta \} \rightarrow \{ \text{Throw } \lesssim \Delta \} \rightarrow \text{Free } \Delta A \\ \text{hello-throw}_1 &= \text{do } \text{\`out "Hello"; } \text{\`out " world!"; } \text{\`throw} \end{aligned}$$

This illustrates how we use the free monad to write effectful programs against an interface given by an effect signature. Next, we define *effect handlers*.

3.2.4 Effect Handlers

An effect handler implements the interface given by an effect signature, interpreting the syntactic operations associated with an effect. Like monadic bind, effect handlers can be defined as a fold over the free monad. The following type of *parameterized handlers* defines how to fold respectively **pure** and **impure** computations:

```
record ⟨ !_ ⇒ _ ⇒ !_ ⟩ (A : Set) (Δ : Effect)
  (P : Set) (B : Set)
  (Δ' : Effect) : Set₁ where
  field ret : A → P → Free Δ' B
        hdl : Alg Δ (P → Free Δ' B)
```

A handler of type $\langle A ! \Delta \Rightarrow P \Rightarrow B ! \Delta' \rangle$ is parameterized in the sense that it turns a computation of type $\text{Free } \Delta A$ into a parameterized computation of type $P \rightarrow \text{Free } \Delta' B$. The following function does so by folding using **ret**, **hdl**, and a **to-front** function:

```
to-front : { Δ₁ • Δ₂ ≈ Δ } → Free Δ A → Free (Δ₁ ⊕ Δ₂) A
to-front { w } = hmap-free (w .reorder .from)

given_handle_ : { w : Δ₁ • Δ₂ ≈ Δ }
  → ⟨ A ! Δ₁ ⇒ P ⇒ B ! Δ₂ ⟩
  → Free Δ A → (P → Free Δ₂ B)

given_handle_ h m = fold
  (ret h)
  ( λ where (inj₁ c , k) p → hdl h (c , k) p
        (inj₂ c , k) p → impure (c , flip k p) )
  (to-front m)
```

Comparing with the syntax we used to explain algebraic effects and handlers in the introduction, the **ret** field corresponds to

the **return** case of the handlers from the introduction, and **hdl** corresponds to the cases that define how operations are handled. The parameterized handler type $\langle A ! \Delta \Rightarrow P \Rightarrow B ! \Delta' \rangle$ corresponds to the type $A ! \Delta, \Delta' \Rightarrow P \rightarrow B ! \Delta'$, and **given h handle m** corresponds to **with h handle m**.

Using this type of handler, the *hOut* handler from the introduction can be defined as follows:

```
hOut : ⟨ A ! Output ⇒ T ⇒ (A × String) ! Δ ⟩
ret hOut x _ = pure (x , "")
hdl hOut (out s , k) p = do (x , s') ← k tt p; pure (x , s ++ s')
```

The handler *hOut* in Section 3.1.1 did not bind any parameters. However, since we are encoding it as a *parameterized* handler, *hOut* now binds a unit-typed parameter. Besides this difference, the handler is the same as in Section 3.1.1. We can use the *hOut* handler to run computations. To this end, we introduce a **Nil** effect with no associated operations which we will use to indicate where an effect row ends:

```
Nil : Effect
Op Nil = ⊥
Ret Nil = ⊥-elim
```

```
un : Free Nil A → A
un (pure x) = x
```

⁴⁸ The *refl* constructor is from the Agda standard library, and witnesses that a propositional equality (\equiv) holds.

Using these, we can run a simple hello world program:⁴⁸

```
hello' : { Output ≲ Δ } → Free Δ T
hello' = do
  `out "Hello"; `out " world!"

test-hello : un (given hOut handle hello' $ tt)
             ≡ (tt , "Hello world!")
test-hello = refl
```

An example of parameterized (as opposed to unparameterized) handlers, is the state effect. Figure 8 declares and illustrates

```

data StateOp : Set where
  get :      StateOp
  put :  $\mathbb{N} \rightarrow$  StateOp

State : Effect
Op State = StateOp
Ret State get    =  $\mathbb{N}$ 
Ret State (put n) =  $\top$ 

hSt :  $\langle A ! \text{State} \Rightarrow \mathbb{N} \Rightarrow (A \times \mathbb{N}) ! \Delta' \rangle$ 
ret hSt x s = pure (x , s)
hdl hSt (put m , k) n = k tt m
hdl hSt (get , k) n = k n    n

`incr : { State  $\lesssim$   $\Delta$  }  $\rightarrow$  Free  $\Delta$   $\top$ 
`incr = do n  $\leftarrow$  `get; `put (n + 1)

incr-test : un ((given hSt handle `incr) 0)  $\equiv$  (tt , 1)
incr-test = refl

```

how to handle such an effect with operations for reading (`get`) and changing (`put`) the state of a memory cell holding a natural number.

3.2.5 The Modularity Problem with Higher-Order Effects, Revisited

Section 3.1.2 described the modularity problem with higher-order effects, using a higher-order operation that interacts with output as an example. In this section we revisit the problem, framing it in terms of the definitions introduced in the previous section. To this end, we use a different effect whose interface is summarized by the `CatchM` record below. The record asserts that the computation type $M : \text{Set} \rightarrow \text{Set}$ has at least a higher-order operation `catch` and a first-order operation `throw`:

```

record CatchM (M : Set  $\rightarrow$  Set) : Set1 where
  field catch : M A  $\rightarrow$  M A  $\rightarrow$  M A
        throw :      M A

```

Figure 8: A state effect (upper), its handler (`hSt` below), and a simple test (`incr-test`, also below) which uses (the elided) smart constructors for `get` and `put`

The idea is that `throw` throws an exception, and `catch` m_1 m_2 handles any exception thrown during evaluation of m_1 by running m_2 instead. The problem is that we cannot give a modular definition of operations such as `catch` using algebraic effects and handlers alone. As discussed in Section 3.1.2, the crux of the problem is that algebraic effects and handlers provide limited support for higher-order operations. However, as also discussed in Section 3.1.2, we can encode `catch` in terms of more primitive effects and handlers, such as the following handler for the `Throw` effect:

```

hThrow : ⟨ A ! Throw ⇒ T ⇒ (Maybe A) ! Δ' ⟩
ret hThrow x _ = pure (just x)
hdl hThrow (throw , k) _ = pure nothing

```

⁴⁹ The `maybe` function is the eliminator for the `Maybe` type. Its first parameter is for eliminating a `just`; the second for `nothing`. Its type is `maybe : (A → B) → B → Maybe A → B`.

⁵⁰ The instance resolution machinery of Agda requires some help to resolve the instance argument of `#` here. We provide a hint to Agda’s instance resolution machinery in an implicit instance argument that we omit for readability in the paper. In the rest of this paper, we will occasionally follow the same convention.

The handler modifies the return type of the computation by decorating it with a `Maybe`. If no exception is thrown, `ret` wraps the yielded value in a `just` constructor. If an exception is thrown, the handler never invokes the continuation k and aborts the computation by returning `nothing` instead. We can elaborate `catch` into an inline application of `hThrow`. To do so we make use of *effect masking* which lets us “weaken” the type of a computation by inserting extra effects in an effect row:

```

#_ : { Δ1 ≲ Δ2 } → Free Δ1 A → Free Δ2 A

```

Using this, the following elaboration defines a semantics for the `catch` operation:⁴⁹ ⁵⁰

```

catch : { Throw ≲ Δ } → Free Δ A → Free Δ A → Free Δ A
catch m1 m2
  = (# (given hThrow handle m1) tt) ≫≅ maybe pure m2

```

If m_1 does not throw an exception, we return the produced value. If it does, m_2 is run.

As observed by Wu et al. [2014], programs that use elaborations such as `catch` are less modular than programs that only use plain algebraic operations. In particular, the effect row type of computations no longer represents the interface of operations

that we use to write programs, since the `catch` elaboration is not represented in the effect type at all. So we have to rely on different machinery if we want to refactor, optimize, or change the semantics of `catch` without having to change programs that use it.

In the next subsection we describe how to define effectful operations such as `catch` modularly using scoped effects and handlers, and discuss how this is not possible for, e.g., operations representing λ -abstraction.

3.2.6 *Scoped Effects and Handlers*

This subsection gives an overview of scoped effects and handlers. While the rest of the paper can be read and understood without a deep understanding of scoped effects and handlers, we include this overview to facilitate comparison with the alternative solution that we introduce in Section 3.3.

Scoped effects extend the expressiveness of algebraic effects to support a class of higher-order operations that [Wu et al. \[2014\]](#), [Piróg et al. \[2018\]](#), [Yang et al. \[2022\]](#) call *scoped operations*. We illustrate how scoped effects work, using a freer monad encoding of the endofunctor algebra approach of [Yang et al. \[2022\]](#). The work of [Yang et al. \[2022\]](#) does not include examples of modular handlers, but the original paper on scoped effects and handlers by [Wu et al. \[2014\]](#) does. We describe an adaptation of the modular handling techniques due to [Wu et al. \[2014\]](#) to the endofunctor algebra approach of [Yang et al. \[2022\]](#).

3.2.6.1 *Scoped Programs*

Scoped effects extend the free monad data type with an additional row for scoped operations. The `return` and `call` constructors of `Prog` below correspond to the `pure` and `impure` constructors of the free monad, whereas `enter` is new:

```
data Prog ( $\Delta \gamma$  : Effect) (A : Set) : Set where
  return : A  $\rightarrow$  Prog  $\Delta \gamma$  A
```

$$\begin{aligned} \text{call} & : \llbracket \Delta \rrbracket (\text{Prog } \Delta \gamma A) \quad \rightarrow \text{Prog } \Delta \gamma A \\ \text{enter} & : \llbracket \gamma \rrbracket (\text{Prog } \Delta \gamma (\text{Prog } \Delta \gamma A)) \rightarrow \text{Prog } \Delta \gamma A \end{aligned}$$

Here, the `enter` constructor represents a higher-order operation with *sub-scopes*; i.e., computations that themselves return computations:

$$\underbrace{\text{Prog } \Delta \gamma}_{\text{outer}} \left(\underbrace{\text{Prog } \Delta \gamma A}_{\text{inner}} \right)$$

This type represents *scoped* computations in the sense that outer computations can be handled independently of inner ones, as we illustrate in Section 3.2.6.2. One way to think of inner computations is as continuations (or join-points) of sub-scopes.

Using `Prog`, the catch operation can be defined as a scoped operation:

```
data CatchOp : Set where
  catch : CatchOp
```

```
Catch : Effect
Op Catch = CatchOp
Ret Catch catch = Bool
```

The effect signature indicates that `Catch` has two scopes since `Bool` has two inhabitants. Following Yang et al. [2022], scoped operations are handled using a structure-preserving fold over `Prog`:

```
hcata : (∀ {X} → X → G X)
        → CallAlg Δ G
        → EnterAlg γ G
        → Prog Δ γ A → G A
```

```
CallAlg : (Δ : Effect) (G : Set → Set) → Set1
CallAlg Δ G =
  {A : Set} → \llbracket Δ \rrbracket (G A) → G A
```

```
EnterAlg : (γ : Effect) (G : Set → Set) → Set1
```

```

EnterAlg  $\gamma$  G =
  {A B : Set}  $\rightarrow$   $\llbracket \gamma \rrbracket$  (G (G A))  $\rightarrow$  G A

```

The first argument represents the case where we are folding a **return** node; the second and third correspond to respectively **call** and **enter**.

3.2.6.2 Scoped Effect Handlers

The following defines a type of parameterized scoped effect handlers:

```

record  $\langle \bullet ! \_ ! \_ \Rightarrow \_ \Rightarrow \_ \bullet ! \_ ! \_ \rangle$  ( $\Delta$   $\gamma$  : Effect) (P : Set)
  (G : Set  $\rightarrow$  Set)
  ( $\Delta'$   $\gamma'$  : Effect) : Set1 where
  field ret    : X  $\rightarrow$  P  $\rightarrow$  Prog  $\Delta'$   $\gamma'$  (G X)
        hcall  : CallAlg  $\Delta$  ( $\lambda$  X  $\rightarrow$  P  $\rightarrow$  Prog  $\Delta'$   $\gamma'$  (G X))
        henter : EnterAlg  $\gamma$  ( $\lambda$  X  $\rightarrow$  P  $\rightarrow$  Prog  $\Delta'$   $\gamma'$  (G X))
        glue   : (k : C  $\rightarrow$  P  $\rightarrow$  Prog  $\Delta'$   $\gamma'$  (G X)) (r : G C)  $\rightarrow$  P
                  $\rightarrow$  Prog  $\Delta'$   $\gamma'$  (G X)

```

A handler of type $\langle \bullet ! \Delta ! \gamma \Rightarrow P \Rightarrow G \bullet ! \Delta' ! \gamma \rangle$ handles operations of Δ and γ *simultaneously* and turns a computation $\text{Prog } \Delta \gamma A$ into a parameterized computation of type $P \rightarrow \text{Prog } \Delta' \gamma' (G A)$. The **ret** and **hcall** cases are similar to the **ret** and **hdl** cases from Section 3.2.4. The crucial addition which adds support for higher-order operations is the **henter** case.

The **henter** field is given by an **EnterAlg** case. This case takes as input a scoped operation whose outer and inner computation have been folded into a parameterized computation of type $P \rightarrow \text{Prog } \Delta' \gamma' (G X)$; and returns as output an interpretation of that operation as a computation of type $P \rightarrow \text{Prog } \Delta' \gamma' (G X)$. The **glue** function is used for modularly *weaving* [Wu et al., 2014] side effects of handlers through sub-scopes of yet-unhandled operations.

3.2.6.3 Weaving

To see why **glue** is needed, it is instructional to look at how the fields in the record type above are used to fold over **Prog**:

```

given_ handle-scoped _ : { w1 : Δ1 • Δ2 ≈ Δ }
                        { w2 : γ1 • γ2 ≈ γ }
                        → ⟨ •! Δ1 ! γ1 ⇒ P ⇒ G •! Δ2 ! γ2 ⟩
                        → Prog Δ γ A → P → Prog Δ2 γ2 (G A)
given h handle-scoped m = hcata (ret h)
  ⊕[ hcall h
    , (λ where (c , k) p → call (c , flip k p)) ]
  ⊕[ (λ {A} → henter h {A})
    , (λ where
        (c , k) p → enter (c ,
          λ x → map-prog (λ y → glue h id y p) (k x p))
      ) ]
  (to-frontΔ (to-frontγ m))

```

The second to last line above shows how **glue** is used. Because **hcata** eagerly folds the current handler over scopes (*sc*), there is a mismatch between the type that the continuation expects (*B*) and the type that the scoped computation returns (*G B*). The **glue** function fixes this mismatch for the particular return type modification $G : \text{Set} \rightarrow \text{Set}$ of a parameterized scoped effect handler.

The scoped effect handler for exception catching is thus:⁵¹

⁵¹ Here,
 $flip : (A \rightarrow B \rightarrow C) \rightarrow$
 $(B \rightarrow A \rightarrow C).$

```

hCatch : ⟨ •! Throw ! Catch ⇒ ⊤ ⇒ Maybe •! Δ ! γ ⟩
ret     hCatch x _ = return (just x)
hcall   hCatch (throw , k) _ = return nothing
henter  hCatch (catch , k) _ = k true tt ≫ λ where
  (just f) → f tt
  nothing → k false tt ≫ maybe ( _ $ tt) (return nothing)
glue hCatch k x _ = maybe (flip k tt) (return nothing) x

```

The **henter** field for the **catch** operation first runs m_1 . If no exception is thrown, the value produced by m_1 is forwarded to

k . Otherwise, m_2 is run and its value is forwarded to k , or its exception is propagated. The `glue` field of `hCatch` says that, if an unhandled exception is thrown during evaluation of a scope, the continuation is discarded and the exception is propagated; and if no exception is thrown the continuation proceeds normally.

3.2.6.4 Discussion and Limitations

As observed by Van den Berg et al. [2021a], some higher-order effects do not correspond to scoped operations. In particular, the `LambdaM` record shown below is not a scoped operation:

```
record LambdaM (V : Set) (M : Set → Set) : Set₁ where
  field lam : (V → M V) → M V
  app : V → M V → M V
```

The `lam` field represents an operation that constructs a λ value. The `app` field represents an operation that will apply the function value in the first parameter position to the argument computation in the second parameter position. The `app` operation has a computation as its second parameter so that it remains compatible with different evaluation strategies.

To see why the operations summarized by the `LambdaM` record above are not scoped operations, let us revisit the `enter` constructor of `Prog`:

$$\text{enter} : \llbracket \gamma \rrbracket \left(\underbrace{\text{Prog } \Delta \gamma}_{\text{outer}} \left(\underbrace{\text{Prog } \Delta \gamma \ A}_{\text{inner}} \right) \right) \rightarrow \text{Prog } \Delta \gamma \ A$$

As summarized earlier in this subsection, `enter` lets us represent higher-order operations (specifically, *scoped operations*), whereas `call` does not (only *algebraic operations*). Just like we defined the computational parameters as scopes (given by the outer `Prog` in the type of `enter`), we might try to define the body of a lambda as a scope in a similar way. However, whereas the `catch` operation always passes control to its continuation (the inner `Prog`), the `lam` effect is supposed to package the body of the lambda into a value and pass this value to the continuation (the inner computation). Because the inner computation is nested

within the outer computation, *the only way to gain access to the inner computation (the continuation) is by first running the outer computation (the body of the lambda)*. This does not give us the right semantics.

It is possible to elaborate the `LambdaM` operations into more primitive effects and handlers, but as discussed in Sections 3.1.2 and 3.2.5, such elaborations are not modular. In the next section we show how to make such elaborations modular.

3.3 HEFTY TREES AND ALGEBRAS

As observed in Section 3.2.5, operations such as `catch` can be elaborated into more primitive effects and handlers. However, these elaborations are not modular. We solve this problem by factoring elaborations into interfaces of their own to make them modular.

To this end, we first introduce a new type of abstract syntax trees (Sections 3.3.1 to 3.3.3) representing computations with higher-order operations, which we dub *hefty trees* (an acronymic pun on *higher-order effects*). We then define elaborations as algebras (*hefty algebras*; Section 3.3.4) over these trees. The following pipeline summarizes the idea, where H is a *higher-order effect signature*:

$$\text{Hefty } H \ A \xrightarrow{\text{elaborate}} \text{Free } \Delta \ A \xrightarrow{\text{handle}} \text{Result}$$

For the categorically inclined reader, `Hefty` conceptually corresponds to the initial algebra of the functor $\text{Hefty}_F \ H \ A \ R = A + H \ R \ (R \ A)$ where $H : (\text{Set} \rightarrow \text{Set}) \rightarrow (\text{Set} \rightarrow \text{Set})$ defines the signature of higher-order operations and is a *higher-order functor*, meaning we have both the usual functorial $\text{map} : (X \rightarrow Y) \rightarrow H \ F \ X \rightarrow H \ F \ Y$ for any functor F as well as a function $\text{hmap} : \text{Nat}(F, G) \rightarrow \text{Nat}(H \ F, H \ G)$ which lifts natural transformations between any F and G to a natural transformation between $H \ F$ and $H \ G$. A hefty algebra is then an F -algebra over a higher-order signature functor H . The notion of elaboration

that we introduce in Section 3.3.4 is an F-algebra whose carrier is a “first-order” effect tree (`Free Δ`).

In this section, we encode this conceptual framework in Agda using a container-inspired approach to represent higher-order signature functors H as a strictly positive type. We discuss and compare our approach with previous work in Section 3.3.5.

3.3.1 Generalizing `Free` to Support Higher-Order Operations

As summarized in Section 3.2.1, `Free Δ A` is the type of abstract syntax trees representing computations over the effect signature Δ . Our objective is to arrive at a more general type of abstract syntax trees representing computations involving (possibly) higher-order operations. To realize this objective, let us consider how to syntactically represent this variant of the *tensor* operation (Section 3.1.2), where M is the type of abstract syntax trees whose type we wish to define:

$$\text{tensor}_{\text{op}} : (\text{String} \rightarrow \text{String}) \rightarrow M \text{ T} \rightarrow M \text{ T}$$

We call the second parameter of this operation a *computation parameter*. Using `Free`, computation parameters can only be encoded as continuations. But the computation parameter of `tensorop` is *not* a continuation, since

$$\text{do } (\text{tensor}_{\text{op}} f m); \backslash \text{out } s \not\equiv \text{tensor}_{\text{op}} f (\text{do } m; \backslash \text{out } s).$$

The crux of the issue is how signature functors $\llbracket \Delta \rrbracket : \text{Set} \rightarrow \text{Set}$ are defined. Since this is an endofunctor on `Set`, the only suitable option in the `impure` constructor is to apply the functor to the type of *continuations*:

$$\text{impure} : \llbracket \Delta \rrbracket (\underbrace{\text{Free } \Delta A}_{\text{continuation}}) \rightarrow \text{Free } \Delta A$$

A more flexible approach would be to allow signature functors to build computation trees with an *arbitrary return type*, including the return type of the continuation. A *higher-order* signature functor of some higher-order signature H , written

$\llbracket H \rrbracket^H : (\text{Set} \rightarrow \text{Set}) \rightarrow \text{Set} \rightarrow \text{Set}$, would fit that bill. Using such a signature functor, we could define the *impure* case as follows:

$$\text{impure} : \llbracket H \rrbracket^H \left(\underbrace{\text{Hefty } H}_{\substack{\text{computation} \\ \text{type}}} \right) \underbrace{A}_{\substack{\text{continuation} \\ \text{return type}}} \rightarrow \text{Hefty } H A$$

Here, *Hefty* is the type of the free monad using higher-order signature functors instead. In the rest of this subsection, we consider how to define higher-order signature functors H , their higher-order functor extensions $\llbracket _ \rrbracket^H$, and the type of *Hefty* trees.

Recall how we defined plain algebraic effects in terms of *containers*:

```
record Effect : Set1 where
  field Op : Set
  Ret : Op → Set
```

Here, *Op* is the type of operations, and *Ret* defines the return type of each operation. In order to allow operations to have sub-computations, we generalize this type to allow each operation to be associated with a number of sub-computations, where each sub-computation can have a different return type. The following record provides this generalization:

```
record EffectH : Set1 where
  field OpH : Set - As before
  RetH : OpH → Set - As before
  Fork : OpH → Set - New
  Ty : {op : OpH} (ψ : Fork op) → Set - New
```

The set of operations is still given by a type field (Op^H), and each operation still has a return type (Ret^H). *Fork* associates each operation with a type that indicates how many sub-computations (or *forks*) an operation has, and *Ty* indicates the return type of

each such fork. For example, say we want to encode an operation op with two sub-computations with different return types, and whose return type is of a unit type. That is, using M as our type of computations:

$$op : M \mathbb{Z} \rightarrow M \mathbb{N} \rightarrow M \top$$

The following signature declares a higher-order effect signature with a single operation with return type \top , and with two forks (we use `Bool` to encode this fact), and where each fork has, respectively \mathbb{Z} and \mathbb{N} as return types.

```
example-op : EffectH
OpH example-op      =  $\top$    - A single operation
RetH example-op tt  =  $\top$    - with return type  $\top$ 
Fork example-op tt   = Bool - with two forks
Ty example-op false  =  $\mathbb{Z}$   - one fork has type  $\mathbb{Z}$ 
Ty example-op true   =  $\mathbb{N}$   - the other has type  $\mathbb{N}$ 
```

The extension of higher-order effect signatures implements the intuition explained above:

$$\begin{aligned} \llbracket _ \rrbracket^H &: \text{Effect}^H \rightarrow (\text{Set} \rightarrow \text{Set}) \rightarrow \text{Set} \rightarrow \text{Set} \\ \llbracket H \rrbracket^H M X &= \Sigma (\text{Op}^H H) \lambda op \rightarrow \\ &(\text{Ret}^H H op \rightarrow M X) \times ((\psi : \text{Fork } H op) \rightarrow M (\text{Ty } H \psi)) \end{aligned}$$

Let us unpack this definition.

$$\underbrace{\Sigma (\text{Op}^H H) \lambda op \rightarrow}_{(1)} \underbrace{(\text{Ret}^H H op \rightarrow M X)}_{(2)} \times \underbrace{((\psi : \text{Fork } H op) \rightarrow M (\text{Ty } H \psi))}_{(3)} \underbrace{\phantom{M (\text{Ty } H \psi)}}_{(4)}$$

The extension of a higher-order signature functor is given by (1) the sum of operations of the signature, where each operation has (2) a continuation (of type $M X$) that expects to be passed a value of the operation's return type, and (3) a set of forks where each fork is (4) a computation that returns the expected type for each fork.

Using the higher-order signature functor and its extension above, our generalized free monad becomes:

```
data Hefty (H : EffectH) (A : Set) : Set where
  pure   : A → Hefty H A
  impure : [[ H ]]H (Hefty H) A → Hefty H A
```

This type of `Hefty` trees can be used to define higher-order operations with an arbitrary number of computation parameters, with arbitrary return types. Using this type, and using a co-product for higher-order effect signatures (`_+_`) which is analogous to the co-product for algebraic effect signatures in Section 3.2.2, Figure 9 represents the syntax of the `ensorop` operation.

```
data CensorOp : Set where
  censor : (String → String)
          → CensorOp

Censor : EffectH
OpH Censor           = CensorOp
RetH Censor (censor f) = T
Fork Censor (censor f) = T
Ty   Censor {censor f} tt = T

ensorop : (String → String) → Hefty (Censor + H) T → Hefty (Censor + H) T
ensorop f m = impure (inj1 (censor f), (λ where tt → m), pure)
```

Figure 9: A higher-order censor effect and operation, with a single computation parameter (declared with `Op = T` in the effect signature top right) with return type `T` (declared with `Ret = λ _ → T` top right)

Just like `Free`, `Hefty` trees can be sequenced using monadic `bind`. Unlike for `Free`, the monadic `bind` of `Hefty` is not expressible in terms of the standard fold over `Hefty` trees. The difference between `Free` and `Hefty` is that `Free` is a regular data type whereas `Hefty` is a *nested datatype* [Bird and Paterson, 1999]. The fold of a nested data type is limited to describe *natural transformations*. As Bird and Paterson [1999] show, this limitation can be overcome by using a *generalized fold*, but for the purpose of this paper it suffices to define monadic `bind` as a recursive function:

```
_>>_ : Hefty H A → (A → Hefty H B) → Hefty H B
pure x           >> g = g x
impure (op , k , s) >> g = impure (op , (_>> g) ∘ k , s)
```

The `bind` behaves similarly to the `bind` for `Free`; i.e., $m \gg= g$ concatenates g to all the leaves in the continuations (but not computation parameters) of m .

In Section 3.3.4 we show how to modularly elaborate higher-order operations into more primitive algebraic effects and handlers (i.e., computations over `Free`), by folding modular elaboration algebras (*hefty algebras*) over `Hefty` trees. First, we show (in Section 3.3.2) how `Hefty` trees support programming against an interface of both algebraic and higher-order operations. We also address (in Section 3.3.3) the question of how to encode effect signatures for higher-order operations whose computation parameters have polymorphic return types, such as the highlighted `A` below:

$$\backslash\text{catch} : \text{Hefty } H \text{ } \boxed{A} \rightarrow \text{Hefty } H \text{ } \boxed{A} \rightarrow \text{Hefty } H \text{ } \boxed{A}$$

3.3.2 Programs with Algebraic and Higher-Order Effects

Any algebraic effect can be lifted to a higher-order effect signature with no fork (i.e., no computation parameters):

$$\begin{aligned} \text{Lift} &: \text{Effect} \rightarrow \text{Effect}^H \\ \text{Op}^H (\text{Lift } \Delta) &= \text{Op } \Delta \\ \text{Ret}^H (\text{Lift } \Delta) &= \text{Ret } \Delta \\ \text{Fork} (\text{Lift } \Delta) &= \lambda _ \rightarrow \perp \\ \text{Ty} (\text{Lift } \Delta) &= \lambda () \end{aligned}$$

Using this effect signature, and using higher-order effect row insertion witnesses analogous to the ones we defined and used in Section 3.2.2, the following smart constructor lets us represent any algebraic operation as a `Hefty` computation:

$$\uparrow_ : \{ w : \text{Lift } \Delta \lesssim^H H \} \rightarrow (op : \text{Op } \Delta) \rightarrow \text{Hefty } H (\text{Ret } \Delta \text{ } op)$$

Using this notion of lifting, `Hefty` trees can be used to program against interfaces of both higher-order and plain algebraic effects.

3.3.3 Higher-Order Operations with Polymorphic Return Types

Let us consider how to define `Catch` as a higher-order effect. Ideally, we would define an operation that is parameterized by a return type of the branches of a particular catch operation, as shown on the left, such that we can define the higher-order effect signature on the right:⁵²

⁵² *d is for dubious.*

```
data CatchOpd : Set1 where
  catchd : Set → CatchOpd
```

```
Catchd : EffectH
OpH Catchd = CatchOpd
RetH Catchd (catchd A) = A
Fork Catchd (catchd A) = Bool
Ty Catchd {catchd A} _ = A
```

The `Fork` field on the right says that the `Catch` operation has two sub-computations (since `Bool` has two constructors), and that each computation parameter has some return type `A`. However, the signature on the right above is not well defined!

The problem is that, because `CatchOpd` has a constructor that quantifies over a type (`Set`), the `CatchOpd` type lives in `Set1`. Consequently it does not fit the definition of `EffectH`, whose operations live in `Set`. There are two potential solutions to this problem: (1) increase the universe level of `EffectH` to allow `OpH` to live in `Set1`; or (2) use a *universe of types* [Martin-Löf, 1984]. Either solution is applicable here. However, for some operations (e.g., λ in Section 3.4.1) it is natural to model types as an interface that we are programming against. For this reason, using a type universe is a natural fit.

A universe of types is a (dependent) pair of a syntax of types (`Ty : Set`) and a semantic function (`[_]ᵀ : Ty → Set`) defining the meaning of the syntax by reflecting it into Agda's `Set`:

```
record Univ : Set1 where
  field Type : Set
       [_]ᵀ : Type → Set
```

Using type universes, we can parameterize the `catch` constructor on the left below by a *syntactic type* `Ty` of some universe u , and use the *meaning of this type* ($\llbracket t \rrbracket^T$) as the return type of the computation parameters in the effect signature on the right below:

```

data CatchOp { u : Univ } : Set where
  catch : Type → CatchOp

Catch : { u : Univ } → EffectH
OpH Catch          = CatchOp
RetH Catch (catch t) =  $\llbracket t \rrbracket^T$ 
Fork Catch (catch t) = Bool
Ty   Catch {catch t} =  $\lambda \_ \rightarrow \llbracket t \rrbracket^T$ 

```

While the universe of types encoding restricts the kind of type that `catch` can have as a return type, the effect signature is parametric in the universe. Thus the implementer of the `Catch` effect signature (or interface) is free to choose a sufficiently expressive universe of types.

3.3.4 Hefty Algebras

As shown in Section 3.2.5, the higher-order catch operation can be encoded as a non-modular elaboration:

```

catch m1 m2
  = (# ((given hThrow handle m1) tt)) >>> (maybe pure m2)

```

We can make this elaboration modular by expressing it as an *algebra* over `Hefty` trees containing operations of the `Catch` signature. To this end, we will use the following notion of hefty algebra (`AlgH`) and fold (or *catamorphism* [Meijer et al., 1991], `cataH`) for `Hefty`:

```

record AlgH (H : EffectH) (F : Set → Set) : Set1 where
  field alg :  $\llbracket H \rrbracket^H F A \rightarrow F A$ 

```

```

cataH : (∀ {A} → A → F A) → AlgH H F → Hefty H A → F A
cataH g a (pure x) = g x

```

$$\begin{aligned} \text{cata}^H g a (\text{impure } (op, k, s)) \\ = \text{alg } a (op, ((\text{cata}^H g a \circ k), (\text{cata}^H g a \circ s))) \end{aligned}$$

Here Alg^H defines how to transform an `impure` node of type `Hefty H A` into a value of type `F A`, assuming we have already folded the computation parameters and continuation into `F` values. A nice property of algebras is that they are closed under higher-order effect signature sums:

$$\begin{aligned} _ \Upsilon _ : \text{Alg}^H H_1 F \rightarrow \text{Alg}^H H_2 F \rightarrow \text{Alg}^H (H_1 \dot{+} H_2) F \\ \text{alg } (A_1 \Upsilon A_2) (\text{inj}_1 op, k, s) = \text{alg } A_1 (op, k, s) \\ \text{alg } (A_1 \Upsilon A_2) (\text{inj}_2 op, k, s) = \text{alg } A_2 (op, k, s) \end{aligned}$$

By defining elaborations as hefty algebras (below) we can compose them using $_ \Upsilon _$.

$$\begin{aligned} \text{Elaboration} : \text{Effect}^H \rightarrow \text{Effect} \rightarrow \text{Set}_1 \\ \text{Elaboration } H \Delta = \text{Alg}^H H (\text{Free } \Delta) \end{aligned}$$

An `Elaboration H Δ` elaborates higher-order operations of signature `H` into algebraic operations of signature `Δ`. Given an elaboration, we can generically transform any hefty tree into more primitive algebraic effects and handlers:

$$\begin{aligned} \text{elaborate} : \text{Elaboration } H \Delta \rightarrow \text{Hefty } H A \rightarrow \text{Free } \Delta A \\ \text{elaborate} = \text{cata}^H \text{pure} \end{aligned}$$

EXAMPLE. The elaboration below is analogous to the non-modular `catch` elaboration discussed in Section 3.2.5 and in the beginning of this subsection:

$$\begin{aligned} \text{eCatch} : \{ u : \text{Univ} \} \{ w : \text{Throw} \lesssim \Delta \} \rightarrow \text{Elaboration Catch } \Delta \\ \text{alg } (\text{eCatch } \{ w = w \}) (\text{catch } t, k, s) = \\ (\# ((\text{given } h\text{Throw } \text{handle } s \text{ true}) \text{ tt})) \ggg \text{maybe } k (s \text{ false} \ggg k) \end{aligned}$$

The elaboration is essentially the same as its non-modular counterpart, except that it now uses the universe of types encoding discussed in Section 3.3.3, and that it now transforms syntactic

representations of higher-order operations instead. Using this elaboration, we can, for example, run the following example program involving the state effect from Figure 8, the throw effect from Section 3.2.1, and the catch effect defined here:

```

transact : { ws : Lift State ≲H H } { wt : Lift Throw ≲H H }
          { w : Catch ≲H H }
          → Hefty H ℕ
transact = do
  ↑ put 1
  \catch (do ↑ (put 2); (↑ throw) ≫= ⊥-elim) (pure tt)
  ↑ get

```

The program first sets the state to 1; then to 2; and then throws an exception. The exception is caught, and control is immediately passed to the final operation in the program which gets the state. By also defining elaborations for `Lift` and `Nil`, we can elaborate and run the program:

```

eTransact : { _ : Throw ≲ Δ } { _ : State ≲ Δ }
           → Elaboration
           (Catch † Lift Throw † Lift State † Lift Nil) Δ
eTransact = eCatch ∨ eLift ∨ eLift ∨ eNil

test-transact : un ( ( given hSt
                    handle ( ( given hThrow
                              handle (
                                elaborate eTransact transact
                              ) tt )
                            )
                  0 ) ≡ (just 2 , 2)
test-transact = refl

```

The program above uses a so-called *global* interpretation of state, where the `put` operation in the “try block” of `\catch` causes the state to be updated globally. In Section 3.4.2.2 we return to this example and show how we can modularly change the elaboration of the higher-order effect `Catch` to yield a so-called *transactional* interpretation of state where the `put` operation in the try block is rolled back when an exception is thrown.

3.3.5 Discussion and Limitations

Which (higher-order) effects can we describe using hefty trees and algebras? Since the core mechanism of our approach is modular elaboration of higher-order operations into more primitive effects and handlers, it is clear that hefty trees and algebras are at least as expressive as standard algebraic effects. The crucial benefit of hefty algebras over algebraic effects is that higher-order operations can be declared and implemented modularly. In this sense, hefty algebras provide a modular abstraction layer over standard algebraic effects that, although it adds an extra layer of indirection by requiring both elaborations and handlers to give a semantics to hefty trees, is comparatively cheap and implemented using only standard techniques such as F-algebras. As we show in Section 3.5, hefty algebras also let us define higher-order effect theories, akin to algebraic effect theories.

Conceptually, we expect that hefty trees can capture any *monadic* higher-order effect whose signature is given by a higher-order functor on $\text{Set} \rightarrow \text{Set}$. Filinski [1999] showed that any monadic effect can be represented using continuations, and given that we can encode the continuation monad using algebraic effects [Schrijvers et al., 2019] in terms of the *sub/jump* operations from Section 3.4.2.2 by Thielecke [1997], Fiore and Staton [2014], it is possible to elaborate any monadic effect into algebraic effects using hefty algebras. The current Agda implementation, though, is slightly more restrictive. The type of effect signatures, Effect^H , approximates the set of higher-order functors by constructively enforcing that all occurrences of the computation type are strictly positive. Hence, there may be higher-order effects that are well-defined semantically, but which cannot be captured in the Agda encoding presented here.

Recent work by Van den Berg and Schrijvers [2023] introduced a higher-order free monad that coincides with our *Hefty* type. Their work shows that hefty trees are, in fact, a free monad. Furthermore, they demonstrate that a range of existing effects

frameworks from the literature can be viewed as instances of hefty trees.

When comparing hefty trees to scoped effects, we observe two important differences. The first difference is that the syntax of programs with higher-order effects is fundamentally more restrictive when using scoped effects. Specifically, as discussed at the end of Section 3.2.6.4, scoped effects impose a restriction on operations that their computation parameters must pass control directly to the continuation of the operation. Hefty trees, on the other hand, do not restrict the control flow of computation parameters, meaning that they can be used to define a broader class of operations. For instance, in Section 3.4.1 we define a higher-order effect for function abstraction, which is an example of an operation where control does not flow from the computation parameter to the continuation.

The second difference is that hefty algebras and scoped effects and handlers are modular in different ways. Scoped effects are modular because we can modularly define, compose, and handle scoped operations, by applying scoped effect handlers in sequence; i.e.:

$$\text{Prog } \Delta_0 \ \gamma_0 \ A_0 \xrightarrow{h_1} \text{Prog } \Delta_1 \ \gamma_1 \ A_1 \xrightarrow{h_2} \dots \xrightarrow{h_n} \text{Prog Nil Nil } A_n$$

As discussed in Section 3.2.6.3, each handler application modularly “weaves” effects through sub computations, using a dedicated `glue` function. Hefty algebras, on the other hand, work by applying an elaboration algebra assembled from modular components in one go. The program resulting from elaboration can then be handled using standard algebraic effect handlers; i.e.:

$$\begin{array}{l} \text{Hefty } (H_0 \dot{+} \dots \dot{+} H_m) \ A \\ \xrightarrow{\text{elaborate } (E_0 \ \gamma \dots \ \gamma \ E_m)} \\ \text{Free } \Delta \ A \\ \xrightarrow{h_1} \dots \xrightarrow{h_k} \\ \text{Free Nil } A_k \end{array}$$

Because hefty algebras eagerly elaborate all higher-order effects in one go, they do not require similar “weaving” as scoped effect handlers. A consequence of this difference is that scoped effect handlers exhibit more effect interaction by default; i.e., different permutations of handlers may give different semantics. In contrast, when using hefty algebras we have to be more explicit about such effect interactions. We discuss this difference in more detail in Section 3.4.2.2.

3.4 EXAMPLES

As discussed in Section 3.2.5, there is a wide range of examples of higher-order effects that cannot be defined as algebraic operations directly, and are typically defined as non-modular elaborations instead. In this section we give examples of such effects and show to define them modularly using hefty algebras. The artifact [Van der Rest and Bach, 2024] contains the full examples.

3.4.1 λ as a Higher-Order Operation

As recently observed by Van den Berg et al. [2021a], the (higher-order) operations for λ abstraction and application are neither algebraic nor scoped effects. We demonstrate how hefty algebras allow us to modularly define and elaborate an interface of higher-order operations for λ abstraction and application, inspired by Levy’s call-by-push-value [Levy, 2004]. The interface we will consider is parametric in a universe of types given by the following record:

```
record LamUniv : Set1 where
  field { u } : Univ
         _  $\mapsto$  _ : Type  $\rightarrow$  Type  $\rightarrow$  Type
         c       : Type  $\rightarrow$  Type
```

The $_ \mapsto _$ field represents a function type, whereas c is the type of *think values*. Distinguishing thinks in this way allows us to

assign either a call-by-value or call-by-name semantics to the interface for λ abstraction summarized by the following smart constructors:

$$\begin{aligned} \backslash\text{lam} &: \{t_1 t_2 : \text{Type}\} \rightarrow (\llbracket c t_1 \rrbracket^T \rightarrow \text{Hefty } H \llbracket t_2 \rrbracket^T) \\ &\rightarrow \text{Hefty } H \llbracket (c t_1) \multimap t_2 \rrbracket^T \\ \backslash\text{var} &: \{t : \text{Type}\} \rightarrow \llbracket c t \rrbracket^T \\ &\rightarrow \text{Hefty } H \llbracket t \rrbracket^T \\ \backslash\text{app} &: \{t_1 t_2 : \text{Type}\} \rightarrow \llbracket (c t_1) \multimap t_2 \rrbracket^T \rightarrow \text{Hefty } H \llbracket t_1 \rrbracket^T \\ &\rightarrow \text{Hefty } H \llbracket t_2 \rrbracket^T \end{aligned}$$

Here $\backslash\text{lam}$ is a higher-order operation with a function typed computation parameter and whose return type is a function value ($\llbracket c t_1 \multimap t_2 \rrbracket^T$). The $\backslash\text{var}$ operation accepts a thunk value as argument and yields a value of a matching type. – The $\backslash\text{app}$ operation is also a higher-order operation: its first parameter is a function value type, whereas its second parameter is a computation parameter whose return type matches the function value parameter type.

The interface above defines a kind of *higher-order abstract syntax* [Pfenning and Elliott, 1988] which piggy-backs on Agda functions for name binding. However, unlike most Agda functions, the constructors above represent functions with side effects. The representation in principle supports functions with arbitrary side effects since it is parametric in what the higher-order effect signature H is. Furthermore, we can assign different operational interpretations to the operations in the interface without having to change the interface or programs written against the interface. To illustrate we give two different implementations of the interface: one that implements a call-by-value evaluation strategy, and one that implements call-by-name.

3.4.1.1 Call-by-Value

We give a call-by-value interpretation $\backslash\text{lam}$ by generically elaborating to algebraic effect trees with any set of effects Δ . Our interpretation is parametric in proof witnesses that the fol-

⁵³ The two sides of an isomorphism $A \leftrightarrow B$ are given by the functions $\text{to} : A \rightarrow B$ and $\text{from} : B \rightarrow A$.

lowing isomorphisms hold for value types (\leftrightarrow is the type of isomorphisms from the Agda standard library):⁵³

$$\begin{aligned} \text{iso}_1 &: \{t_1 \ t_2 : \text{Type}\} \rightarrow \llbracket t_1 \multimap t_2 \rrbracket^T \leftrightarrow (\llbracket t_1 \rrbracket^T \rightarrow \text{Free } \Delta \llbracket t_2 \rrbracket^T) \\ \text{iso}_2 &: \{t : \text{Type}\} \rightarrow \llbracket \text{c } t \rrbracket^T \leftrightarrow \llbracket t \rrbracket^T \end{aligned}$$

The first isomorphism says that a function value type corresponds to a function which accepts a value of type t_1 and produces a computation whose return type matches the function type. The second says that thunk types coincide with value types. Using these isomorphisms, the following defines a call-by-value elaboration of functions:

```
eLamCBV : Elaboration Lam Δ
alg eLamCBV (lam , k , ψ) = k (from ψ)
alg eLamCBV (var x , k , _) = k (to x)
alg eLamCBV (app f , k , ψ) = do
  a ← ψ tt
  v ← to f (from a)
  k v
```

The `lam` case passes the function body given by the sub-tree ψ as a value to the continuation, where the `from` function mediates the sub-tree of type $\llbracket \text{c } t_1 \rrbracket^T \rightarrow \text{Free } \Delta \llbracket t_2 \rrbracket^T$ to a value type $\llbracket (\text{c } t_1) \multimap t_2 \rrbracket^T$, using the isomorphism `iso1`. The `var` case uses the `to` function to mediate a $\llbracket \text{c } t \rrbracket^T$ value to a $\llbracket t \rrbracket^T$ value, using the isomorphism `iso2`. The `app` case first eagerly evaluates the argument expression of the application (in the sub-tree ψ) to an argument value, and then passes the resulting value to the function value of the application. The resulting value is passed to the continuation.

Using the elaboration above, we can evaluate programs such as the following which uses both the higher-order lambda effect, the algebraic state effect, and assumes that our universe has a number type $\llbracket \text{num} \rrbracket^T \leftrightarrow \mathbb{N}$:

```
ex : Hefty (Lam † Lift State † Lift Nil) ℤ
ex = do
```

```

↑ put 1
f ← \lam (λ x → do
  n1 ← \var x
  n2 ← \var x
  pure (from ((to n1) + (to n2))))
v ← \app f incr
pure (to v)
where incr = do s0 ← ↑ get
  ↑ put (s0 + 1)
  s1 ← ↑ get
  pure (from s1)

```

The program first sets the state to 1. Then it constructs a function that binds a variable x , dereferences the variable twice, and adds the two resulting values together. Finally, the application in the second-to-last line applies the function with an argument expression which increments the state by 1 and returns the resulting value. Running the program produces 4 since the state increment expression is eagerly evaluated before the function is applied.

```

elab-cbv : Elaboration (Lam † Lift State † Lift Nil) (State ⊕ Nil)
elab-cbv = eLamCBV √ eLift √ eNil
test-ex-cbv : un ((given hSt handle (elaborate elab-cbv ex)) 0)
              ≡ (4 , 2)
test-ex-cbv = refl

```

3.4.1.2 Call-by-Name

The key difference between the call-by-value and the call-by-name interpretation of our λ operations is that we now assume that thunks are computations. That is, we assume that the following isomorphisms hold for value types:

$$\begin{aligned}
\text{iso}_1 : \{t_1 \ t_2 : \text{Type}\} &\rightarrow \llbracket t_1 \rightsquigarrow t_2 \rrbracket^T \leftrightarrow (\llbracket t_1 \rrbracket^T \rightarrow \text{Free } \Delta \llbracket t_2 \rrbracket^T) \\
\text{iso}_2 : \{t : \text{Type}\} &\rightarrow \llbracket \text{c } t \rrbracket^T \leftrightarrow \text{Free } \Delta \llbracket t \rrbracket^T
\end{aligned}$$

Using these isomorphisms, the following defines a call-by-name elaboration of functions:

```
eLamCBN : Elaboration Lam Δ
alg eLamCBN (lam , k , ψ) = k (from ψ)
alg eLamCBN (var x , k , _) = to x ≫ k
alg eLamCBN (app f , k , ψ) = to f (from (ψ tt)) ≫ k
```

The case for `lam` is the same as the call-by-value elaboration. The case for `var` now needs to force the thunk by running the computation and passing its result to k . The case for `app` passes the argument sub-tree (ψ) as an argument to the function f , runs the computation resulting from doing so, and then passes its result to k . Running the example program `ex` from above now produces `5` as result, since the state increment expression in the argument of `\app` is thunked and run twice during the evaluation of the called function.

```
elab-cbn : Elaboration (Lam † Lift State † Lift Nil) (State ⊕ Nil)
elab-cbn = eLamCBN ∨ eLift ∨ eNil
test-ex-cbn : un ((given hSt handle (elaborate elab-cbn ex)) 0)
              ≡ (5 , 3)
test-ex-cbn = refl
```

3.4.2 Optionally Transactional Exception Catching

A feature of scoped effect handlers [Wu et al., 2014, Piróg et al., 2018, Yang et al., 2022] is that changing the order of handlers makes it possible to obtain different semantics of *effect interaction*. A classical example of effect interaction is the interaction between state and exception catching that we briefly considered at the end of Section 3.3.4 in connection with this `transact` program:

```
transact : { ws : Lift State ≲H H } { wt : Lift Throw ≲H H }
          { w : Catch ≲H H }
          → Hefty H ℕ
```

```

transact = do
  ↑ put 1
  \catch (do ↑ put 2; (↑ throw) >>= ⊥-elim) (pure tt)
  ↑ get

```

The state and exception catching effect can interact to give either of these two semantics:

1. *Global* interpretation of state, where the `transact` program returns `2` since the `put` operation in the “try” block causes the state to be updated globally.
2. *Transactional* interpretation of state, where the `transact` program returns `1` since the state changes of the `put` operation are *rolled back* when the “try” block throws an exception.

With monad transformers [Cenciarelli and Moggi, 1993, Liang et al., 1995b] we can recover either of these semantics by permuting the order of monad transformers. With scoped effect handlers we can also recover either by permuting the order of handlers. However, the `eCatch` elaboration in Section 3.3.4 always gives us a global interpretation of state. In this section we demonstrate how we can recover a transactional interpretation of state by using a different elaboration of the `catch` operation into an algebraically effectful program with the `throw` operation and the off-the-shelf *sub/jump* control effects due to Thielecke [1997], Fiore and Staton [2014]. The different elaboration is modular in the sense that we do not have to change the interface of the `catch` operation nor any programs written against the interface.

3.4.2.1 Sub/Jump

We recall how to define two operations, `sub` and `jump`, due to Thielecke [1997], Fiore and Staton [2014]. We define these operations as algebraic effects following Schrijvers et al. [2019]. The algebraic effects are summarized by the following smart constructors where `CC Ref` is associated with the `sub/jump` operations:

$$\begin{aligned}
& \mathbf{h} (\backslash\text{sub} (\lambda _ \rightarrow p) k) \equiv \mathbf{h} p \\
& \mathbf{h} (\backslash\text{sub} (\lambda r \rightarrow m \gg\gg \backslash\text{jump} r) k) \equiv \mathbf{h} (m \gg\gg k) \\
& \mathbf{h} (\backslash\text{sub} p (\backslash\text{jump} r')) \equiv \mathbf{h} (p r') \\
& \mathbf{h} (\backslash\text{sub} p q \gg\gg k) \equiv \mathbf{h} (\backslash\text{sub} (\lambda x \rightarrow p x \gg\gg k) (\lambda x \rightarrow q x \gg\gg k))
\end{aligned}$$

Figure 10: Laws for the $\backslash\text{sub}$ and $\backslash\text{jump}$ operations

$$\begin{aligned}
\backslash\text{sub} & : \{w : \mathbf{CC} \text{Ref} \lesssim \Delta\} (b : \text{Ref } t \rightarrow \mathbf{Free} \Delta A) \\
& \rightarrow (k : \llbracket t \rrbracket^T \rightarrow \mathbf{Free} \Delta A) \rightarrow \mathbf{Free} \Delta A
\end{aligned}$$

$$\backslash\text{jump} : \{w : \mathbf{CC} \text{Ref} \lesssim \Delta\} (\text{ref} : \text{Ref } t) (x : \llbracket t \rrbracket^T) \rightarrow \mathbf{Free} \Delta B$$

An operation $\backslash\text{sub} f g$ gives a computation f access to the continuation g via a reference value $\text{Ref } t$ which represents a continuation expecting a value of type $\llbracket t \rrbracket^T$. The $\backslash\text{jump}$ operation invokes such continuations. The operations and their handler (abbreviated to \mathbf{h}) satisfy the laws shown in Figure 10. The last law in Figure 10 asserts that $\backslash\text{sub}$ and $\backslash\text{jump}$ are *algebraic* operations, since their computational sub-terms behave as continuations. Thus, we encode $\backslash\text{sub}$ and its handler as an algebraic effect.

3.4.2.2 Optionally Transactional Exception Catching

By using the $\backslash\text{sub}$ and $\backslash\text{jump}$ operations in our elaboration of catch , we get a semantics of exception catching whose interaction with state depends on the order that the state effect and sub/jump effect is handled.

$$\begin{aligned}
\text{eCatchOT} & : \{w_1 : \mathbf{CC} \text{Ref} \lesssim \Delta\} \{w_2 : \mathbf{Throw} \lesssim \Delta\} \\
& \rightarrow \mathbf{Elaboration} \text{Catch } \Delta \\
\text{alg eCatchOT} & (\text{catch } x, k, \psi) = \\
& \text{let } m_1 = \psi \text{ true}; m_2 = \psi \text{ false in} \\
& \backslash\text{sub} (\lambda r \rightarrow (\# ((\text{given } \mathbf{hThrow} \text{ handle } m_1) \text{ tt})) \\
& \quad \gg\gg \text{maybe } k (\backslash\text{jump } r (\text{from tt}))) \\
& (\lambda _ \rightarrow m_2 \gg\gg k)
\end{aligned}$$

The elaboration uses `\sub` to capture the continuation of a higher-order `catch` operation. If no exception is raised, then control flows to the continuation k without invoking the continuation of `\sub`. Otherwise, we jump to the continuation of `\sub`, which runs m_2 before passing control to k . Capturing the continuation in this way interacts with state because the continuation of `\sub` may have been pre-applied to a state handler that only knows about the “old” state. This happens when we invoke the state handler before the handler for `sub/jump`: in this case we get the transactional interpretation of state, so running `transact` gives 1. Otherwise, if we run the `sub/jump` handler before the state handler, we get the global interpretation of state and the result 2.

The `sub/jump` elaboration above is more involved than the scoped effect handler that we considered in Section 3.2.6. However, the complicated encoding does not pollute the higher-order effect interface, and only turns up if we strictly want or need effect interaction.

3.4.3 Logic Programming

Following Schrijvers et al. [2014], Wu et al. [2014], Yang et al. [2022] we can define a non-deterministic choice operation (`\or_`) as an algebraic effect, to provide support for expressing the kind of non-deterministic search for solutions that is common in logic programming. We can also define a `\fail` operation which indicates that the search in the current branch was unsuccessful. The smart constructors below are the lifted higher-order counterparts to the `\or_` and `\fail` operations:

$$\begin{aligned} _ \text{\or_} _ &: \{ \text{Lift Choice } \lesssim^H H \} \rightarrow \text{Hefty } H A \rightarrow \text{Hefty } H A \\ &\rightarrow \text{Hefty } H A \\ _ \text{\fail}^H &: \{ \text{Lift Choice } \lesssim^H H \} \rightarrow \text{Hefty } H A \end{aligned}$$

A useful operator for cutting non-deterministic search short when a solution is found is the `\once` operator. The `\once` operator

is not an algebraic effect, but a scoped (and thus higher-order) effect.

$$\begin{aligned} \text{\code `once} &: \{w : \text{Once} \lesssim^H H\} \{t : \text{Type}\} \\ &\rightarrow \text{Hefty } H \llbracket t \rrbracket^T \rightarrow \text{Hefty } H \llbracket t \rrbracket^T \end{aligned}$$

We can define the meaning of the `once` operator as the following elaboration:

$$\begin{aligned} \text{eOnce} &: \{ \text{Choice} \lesssim \Delta \} \rightarrow \text{Elaboration Once } \Delta \\ \text{alg eOnce} &(\text{once}, k, \psi) = \text{do} \\ & \quad l \leftarrow \# ((\text{given hChoice handle } (\psi \text{ tt})) \text{ tt}) \\ & \quad \text{maybe } k \text{ `fail } (\text{head } l) \end{aligned}$$

The elaboration runs the branch (ψ) of `once` under the `hChoice` handler to compute a list of all solutions of ψ . It then tries to choose the first solution and pass that to the continuation k . If the branch has no solutions, we fail. Under a strict evaluation order, the elaboration computes all possible solutions which is doing more work than needed. Agda 2.6.2.2 does not have a specified evaluation strategy, but does compile to Haskell which is lazy. In Haskell, the solutions would be lazily computed, such that the `once` operator cuts search short as intended.

3.4.4 Concurrency

Finally, we consider how to define higher-order operations for concurrency, inspired by Yang et al.'s [2022] *resumption monad* (see also [Claessen, 1999, Schmidt, 1986, Piróg and Gibbons, 2014]) definition using scoped effects. We summarize our encoding and compare it with the resumption monad. The goal is to define the following operations:

$$\begin{aligned} \text{\code `spawn} &: \{t : \text{Type}\} \rightarrow (m_1 \ m_2 : \text{Hefty } H \llbracket t \rrbracket^T) \rightarrow \text{Hefty } H \llbracket t \rrbracket^T \\ \text{\code `atomic} &: \{t : \text{Type}\} \rightarrow \text{Hefty } H \llbracket t \rrbracket^T \rightarrow \text{Hefty } H \llbracket t \rrbracket^T \end{aligned}$$

The operation ``spawn` $m_1 \ m_2$ spawns two threads that run concurrently, and returns the value produced by m_1 after both

have finished. The operation `\atomic m` represents a block to be executed atomically; i.e., no other threads run before the block finishes executing.

We elaborate `\spawn` by interleaving the sub-trees of its computations. To this end, we use a dedicated function which interleaves the operations in two trees and yields as output the value of the left input tree (the first computation parameter):

$$\begin{aligned} \text{interleave}_1 &: \{ \text{Ref} : \text{Type} \rightarrow \text{Set} \} \rightarrow \text{Free} (\text{CC Ref} \oplus \Delta) A \\ &\rightarrow \text{Free} (\text{CC Ref} \oplus \Delta) B \rightarrow \text{Free} (\text{CC Ref} \oplus \Delta) A \end{aligned}$$

Here, the `CC` effect is the sub/jump effect that we also used in Section 3.4.2.2. The `interleave1` function ensures atomic execution by only interleaving code that is not wrapped in a `\sub` operation. We elaborate `Concur` into `CC` as follows, where the `to-front` and `from-front` functions use the row insertion witness w_a to move the `CC` effect to the front of the row and back again:

$$\begin{aligned} \text{eConcur} &: \{ w : \text{CC Ref} \lesssim \Delta \} \rightarrow \text{Elaboration Concur } \Delta \\ \text{alg eConcur} (\text{spawn } t, k, \psi) &= \\ &\quad \text{from-front} (\text{interleave}_1 (\text{to-front} (\psi \text{ true})) (\text{to-front} (\psi \text{ false}))) \\ &\quad \ggg k \\ \text{alg eConcur} (\text{atomic } t, k, \psi) &= \\ &\quad \text{\sub} (\lambda \text{ ref} \rightarrow \psi \text{ tt} \ggg \text{\jump ref}) k \end{aligned}$$

The elaboration uses `\sub` as a delimiter for blocks that should not be interleaved, such that the `interleave1` function only interleaves code that does not reside in atomic blocks. At the end of an `atomic` block, we `\jump` to the (possibly interleaved) computation context, k . By using `\sub` to explicitly delimit blocks that should not be interleaved, we have encoded what Wu et al. [2014, § 7] call *scoped syntax*.

EXAMPLE. Below is an example program that spawns two threads that use the `Output` effect. The first thread prints 0, 1, and 2; the second prints 3 and 4.

$$\begin{aligned} \text{ex-01234} &: \text{Hefty} (\text{Lift Output} \dot{+} \text{Concur} \dot{+} \text{Lift Nil}) \text{IN} \\ \text{ex-01234} &= \text{\spawn} (\text{do } \uparrow \text{out "0"}; \uparrow \text{out "1"}; \uparrow \text{out "2"}; \text{pure } 0) \\ &\quad (\text{do } \uparrow \text{out "3"}; \uparrow \text{out "4"}; \text{pure } 0) \end{aligned}$$

Since the `Concur` effect is elaborated to interleave the effects of the two threads, the printed output appears in interleaved order:

```
test-ex-01234 : un ( ( given hOut
                      handle
                        ( ( given hCC
                          handle
                            ( elaborate
                              concur-elab ex-01234 )
                            ) tt ) ) tt ) ≡ ( 0 , "03142" )
test-ex-01234 = refl
```

The following program spawns an additional thread with an `\atomic` block

```
ex-01234567 : Hefty (Lift Output † Concur † Lift Nil) ℕ
ex-01234567 =
  \spawn ex-01234
    (\atomic (do ↑ out "5"; ↑ out "6"; ↑ out "7"; pure 0))
```

Inspecting the output, we see that the additional thread indeed computes atomically:

```
test-ex-01234567 : un ( ( given hOut
                          handle
                            ( ( given hCC
                              handle
                                ( elaborate
                                  concur-elab ex-01234567 )
                                ) tt ) ) tt ) ≡ ( 0 , "05673142" )
test-ex-01234567 = refl
```

The example above is inspired by the resumption monad, and in particular by the scoped effects definition of concurrency due to [Yang et al. \[2022\]](#). [Yang et al.](#) do not (explicitly) consider how to define the concurrency operations in a modular style. Instead, they give a direct semantics that translates to the resumption

monad which we can encode as follows in Agda (assuming resumptions are given by the free monad):

```
data Resumption Δ A : Set where
  done : A → Resumption Δ A
  more : Free Δ (Resumption Δ A) → Resumption Δ A
```

We could elaborate into this type using a hefty algebra $\text{Alg}^{\text{H}} \text{Concur} (\text{Resumption } \Delta)$ but that would be incompatible with our other elaborations which use the free monad. For that reason, we emulate the resumption monad using the free monad instead of using the `Resumption` type directly.

3.5 MODULAR REASONING FOR HIGHER-ORDER EFFECTS

A key aspect of algebraic effects and handlers is the ability to state and prove *equational laws* that characterize correct implementations of effectful operations. Usually, an effect comes equipped with multiple laws that govern its intended behavior. An effect and its laws is generally known as an *effect theory* [Hyland et al., 2006, Plotkin and Power, 2002, 2003, Yang and Wu, 2021]. This concept of effect theory extends to *higher-order effect theories*, which describe the intended behavior of higher-order effects. In this section, we first discuss how to define theories for algebraic effects in Agda by adapting the exposition of Yang and Wu [2021], and show how correctness of implementations with respect to a given theory can be stated and proved. We then extend this reasoning infrastructure to higher-order effects, allowing for modular reasoning about the correctness of elaborations of higher-order effects.

Let us consider the state effect as an example, which comprises the `get` and `put` operations. With the state effect, we typically associate a set of equations (or laws) that specify how its implementations ought to behave. One such law is the *get-get* law, which captures the intuition that the state returned by two

subsequent `get` operation does not change if we do not use the `put` operation in between:

$$\text{\code{get}} \gg= \lambda s \rightarrow \text{\code{get}} \gg= \lambda s' \rightarrow k\ s\ s' \equiv \text{\code{get}} \gg= \lambda s \rightarrow k\ s\ s$$

We can define equational laws for higher-order effects in a similar fashion. For example, the following *catch-return* law for the `\catch` operation of the `Catch` effect, stating that catching exceptions in a computation that only returns a value does nothing.

$$\text{\code{\catch}} (\text{\code{pure}}\ x)\ m \equiv \text{\code{pure}}\ x$$

Correctness of an implementation of an algebraic effect with respect to a given theory is defined by comparing the implementations of programs that are equal under that theory. That is, if we can show that two programs are equal using the equations of a theory for its effects, handling the effects should produce equal results. For instance, a way to implement the state effect is by mapping programs to functions of the form $S \rightarrow S \times A$. Such an implementation would be correct if programs that are equal with respect to a theory of the state effect are mapped to functions that give the same value and output state for every input state.

For higher-order effects, correctness is defined in a similar manner. However, since we define higher-order effects by elaborating them into algebraic effects, correctness of elaborations with respect to a higher-order effect theory is defined by comparing the elaborated programs. Crucially, the elaborated programs do not have to be syntactically equal, but rather we should be able to prove them equal using a theory of the algebraic effects used to implement a higher-order effect.

Effect theories are known to be closed under the co-product of effects, by combining the equations into a new theory that contains all equations for both effects [Hyland et al., 2006]. Similarly, theories of higher-order effects are closed under sums of higher-order effect signatures. In Section 3.5.8, we show that composing two elaborations preserves their correctness, with respect to the sum of their respective theories.

3.5.1 Theories of Algebraic Effects

Theories of effects are collections of equations, so we start defining the type of equations in Agda. At its core, an equation for an effect Δ is given by a pair of effect trees of type `Free Δ A`, that define the left- and right-hand side of the equation. However, looking at the *get-get* law above, we see that this equation contains a *term metavariable*; i.e., k . Furthermore, when considering the type of k , which is $S \rightarrow S \rightarrow \text{Free } \Delta A$, we see that it refers to a *type metavariable*; i.e., A . Generally speaking, an equation may refer to any number of term metavariables, which, in turn, may depend on any number of type metavariables. Moreover, the type of the value returned by the left hand side and right hand side of an equation may depend on these type metavariables as well, as is the case for the *get-get* law above. This motivates the following definition of equations in Agda.

```
record Equation ( $\Delta$  : Effect) : Set1 where
  field
    V      : ℕ
    Γ      : Vec Set V → Set
    R      : Vec Set V → Set
    lhs rhs : (vs : Vec Set V) → Γ vs → Free Δ (R vs)
```

An equation consists of five components. The field `V` defines the number of type metavariables used in the equation. Then, the fields `Γ` and `R` define the term metavariables respectively return type of the equation. Both may depend on the type metavariables of the equation, hence they depend on a vector of length `V` containing unapplied substitutions for all type metavariables. Finally, the left-hand side (`lhs`) and right-hand side (`rhs`) of the equation are then defined as functions of type `Free Δ (R vs)`, which depend on unapplied substitutions for both the type and term level metavariables that the equation can refer to.

EXAMPLE . To illustrate how the `Equation` record captures equational laws of effects, we consider how to define the *get-get* as a value of type `Equation State`. Recall that the equation

depends on one type metavariable, and one term metavariable. Furthermore, the return type of the programs on the left and right hand sides should be equal to this type metavariable.

```

get-get : Equation State
V  get-get = 1
Γ  get-get = λ where (A :: []) → ℕ → ℕ → Free State A
R  get-get = λ where (A :: []) → A

```

Since there is exactly one term metavariable, we equate Γ to the type of this metavariable. For equations with more than one metavariable, we would define Γ as a product of the types of all term metavariables. The fields `lhs` and `rhs` for the *get-get* law are then defined as follows:

```

lhs get-get (A :: []) k = `get >>= λ s → `get >>= λ s' → k s s'
rhs get-get (A :: []) k = `get >>= λ s → k s s

```

3.5.2 Modal Necessity

The current definition of equations is too weak, in the sense that it does not apply in many situations where it should. The issue is that it fixes the set of effects that can be used in the left- and right-hand side. To illustrate why this is problematic, consider the following equality:

$$\text{get} \gg= \lambda s \rightarrow \text{get} \gg= \lambda s' \rightarrow \text{throw} \equiv \text{get} \gg= \lambda s \rightarrow \text{throw} \quad (1)$$

We might expect to be able to prove this equality using the *get-get* law, but using the embedding of the law defined above—i.e., *get-get*—this is not possible. The reason for this is that we cannot pick an appropriate instantiation for the term metavariable `k`: it ranges over values of type $S \rightarrow S \rightarrow \text{Free State } A$, inhibiting all references to effectful operation that are not part of the state effect, such as `throw`.

Given an equation for the effect Δ , the solution to this problem is to view Δ as a *lower bound* on the effects that might occur in

the left-hand and right-hand side of the equation, rather than an exact specification. Effectively, this means that we close over all possible contexts of effects in which the equation can occur. This pattern of closing over all possible extensions of a type index is well-known [Allais et al., 2021, Van der Rest et al., 2022b], and corresponds to a shallow embedding of the Kripke semantics of the necessity modality from modal logic. We can define it in Agda as follows.

```
record  $\Box$  (P : Effect  $\rightarrow$  Set1) ( $\Delta$  : Effect) : Set1 where
  constructor necessary
  field
     $\Box$ (_) :  $\forall \{\Delta'\} \rightarrow \{\Delta \lesssim \Delta'\} \rightarrow P \Delta'$ 
```

Intuitively, the \Box modality transforms, for any effect-indexed type ($P : \text{Effect} \rightarrow \text{Set}_1$), an *exact* specification of the set of effects to a *lower bound* on the set of effects. For equations, the difference between terms of type Equation Δ and \Box Equation Δ amounts to the former defining an equation relating programs that have exactly effects Δ , while the latter defines an equation relating programs that have at least the effects Δ but potentially more. The \Box modality is a *comonad*: the counit (`extract` below) witnesses that we can always transform a lower bound on effects to an exact specification, by instantiating the extension witness with a proof of reflexivity.

```
extract : {P : Effect  $\rightarrow$  Set1}  $\rightarrow$   $\Box$  P  $\Delta$   $\rightarrow$  P  $\Delta$ 
extract px =  $\Box$ (px) { $\lesssim$ -refl}
```

We can now redefine the *get-get* law such that it applies to all programs that have the `State` effect, but potentially other effects too.

```
get-get :  $\Box$  Equation State
 $\forall$   $\Box$ (get-get) = 1
 $\Gamma$   $\Box$ (get-get) (A :: []) =  $\mathbb{N} \rightarrow \mathbb{N} \rightarrow \text{Free } \_ A$ 
 $R$   $\Box$ (get-get) (A :: []) = A
lhs  $\Box$ (get-get) (A :: []) k = 'get  $\ggg$   $\lambda s \rightarrow$  'get  $\ggg$   $\lambda s' \rightarrow k s s'$ 
rhs  $\Box$ (get-get) (A :: []) k = 'get  $\ggg$   $\lambda s \rightarrow k s s$ 
```

The above definition of the *get-get* law now lets us prove the equality in Equation (1); the term metavariable k ranges over all continuations that return a tree of type $\text{Free } \Delta' A$, for all Δ' such that $\text{State} \lesssim \Delta'$. This way, we can instantiate Δ' with an effect signature that subsumes both the State and the Throw , which in turn allows us to instantiate k with throw .

3.5.3 Effect Theories

Equations for an effect Δ can be combined into a *theory* for Δ . A theory for the effect Δ is simply a collection of equations, transformed using the \square modality to ensure that term metavariables can range over programs that include more effects than just Δ .

```
record Theory ( $\Delta$  : Effect) : Set1 where
  field
    arity      : Set
    equations  : arity  $\rightarrow$   $\square$  Equation  $\Delta$ 
```

An effect theory consists of an *arity*, that defines the number of equations in the theory, and a function that maps arities to equations. We can think of effect theories as defining a specification for how implementations of an effect ought to behave. Although implementations may vary, depending for example on whether they are tailored to readability or efficiency, they should at least respect the equations of the theory of the effect they implement. We will make precise what it means for an implementation to respect an equation in Section 3.5.5.

Effect theories are closed under several composition operations that allow us to combine the equations of different theories into single theory. The most basic way of combining effect theories is by summing their arities.

```
 $\_ \langle + \rangle \_ : \text{Theory } \Delta \rightarrow \text{Theory } \Delta \rightarrow \text{Theory } \Delta$ 
arity      ( $T_1 \langle + \rangle T_2$ ) = arity  $T_1 \uplus$  arity  $T_2$ 
equations ( $T_1 \langle + \rangle T_2$ ) ( $\text{inj}_1 a$ ) = equations  $T_1 a$ 
equations ( $T_1 \langle + \rangle T_2$ ) ( $\text{inj}_2 a$ ) = equations  $T_2 a$ 
```

This way of combining effects is somewhat limiting, as it imposes that the theories we are combining are theories for the exact same effect. It is more likely, however, that we would want to combine theories for different effects. This requires that we can *weaken* effect theories with respect to the \lesssim relation.

$$\begin{aligned}
\text{weaken-}\square &: \{P : \text{Effect} \rightarrow \text{Set}_1\} \rightarrow \{\Delta_1 \lesssim \Delta_2\} \\
&\rightarrow \square P \Delta_1 \rightarrow \square P \Delta_2 \\
\square \langle \text{weaken-}\square \{w\} px \rangle \{w'\} &= \square \langle px \rangle \{ \lesssim\text{-trans } w w' \} \\
\text{weaken-theory} &: \{\Delta_1 \lesssim \Delta_2\} \rightarrow \text{Theory } \Delta_1 \rightarrow \text{Theory } \Delta_2 \\
\text{arity} &(\text{weaken-theory } T) = \text{arity } T \\
\text{equations} &(\text{weaken-theory } T) = \lambda a \rightarrow \text{weaken-}\square (T.\text{equations } a)
\end{aligned}$$

Categorically speaking, the observation that for a given effect-indexed type P we can transform a value of type $P \Delta_1$ to a value of type $P \Delta_2$ if we know that $\Delta_1 \lesssim \Delta_2$ is equivalent to saying that P is a functor from the category of containers and container morphisms to the category of sets. From this perspective, the existence of weakening for free `Free`, as witnessed by the $\#$ operation discussed in Section 3.3 implies that it too is a such a functor.

With weakening for theories at our disposal, we can combine effect theories for different effects into a theory of the coproduct of their respective effects. This requires us to first define appropriate witnesses relating coproducts to effect inclusion.

$$\begin{aligned}
\lesssim\text{-}\oplus\text{-left} &: \Delta_1 \lesssim (\Delta_1 \oplus \Delta_2) \\
\lesssim\text{-}\oplus\text{-right} &: \Delta_2 \lesssim (\Delta_1 \oplus \Delta_2)
\end{aligned}$$

It is now straightforward to show that effect theories are closed under the coproduct of effect signatures, by summing the weakened theories.

$$\begin{aligned}
_ [+] _ &: \text{Theory } \Delta_1 \rightarrow \text{Theory } \Delta_2 \rightarrow \text{Theory } (\Delta_1 \oplus \Delta_2) \\
T_1 [+] T_2 &= \\
&\text{weaken-theory } \{ \lesssim\text{-}\oplus\text{-left} \} T_1 \\
\langle + \rangle &\text{weaken-theory } \{ \lesssim\text{-}\oplus\text{-right} \} T_2
\end{aligned}$$

While this operation is in principle sufficient for our purposes, it forces a specific order on the effects of the combined theories. We can further generalize the operation above to allow for the effects of the combined theory to appear in any order. This requires the following instances.

$$\begin{aligned} \lesssim\text{-}\bullet\text{-left} & : \{ \Delta_1 \bullet \Delta_2 \approx \Delta \} \rightarrow \Delta_1 \lesssim \Delta \\ \lesssim\text{-}\bullet\text{-right} & : \{ \Delta_1 \bullet \Delta_2 \approx \Delta \} \rightarrow \Delta_2 \lesssim \Delta \end{aligned}$$

We show that effect theories are closed under coproducts up to reordering by, again, summing the weakened theories.

$$\begin{aligned} \text{compose-theory} & : \{ \Delta_1 \bullet \Delta_2 \approx \Delta \} \\ & \rightarrow \text{Theory } \Delta_1 \rightarrow \text{Theory } \Delta_2 \rightarrow \text{Theory } \Delta \\ \text{compose-theory } T_1 T_2 & = \\ & \text{weaken-theory } \{ \lesssim\text{-}\bullet\text{-left} \} T_1 \\ \langle + \rangle \text{weaken-theory } & \{ \lesssim\text{-}\bullet\text{-right} \} T_2 \end{aligned}$$

Since equations are defined by storing the syntax trees that define their left-hand and right-hand side, and effect trees are weakenable, we would expect equations to be weakenable too. Indeed, we can define the following function witnessing weakenability of equations.

$$\text{weaken-eq} : \{ \Delta_1 \lesssim \Delta_2 \} \rightarrow \text{Equation } \Delta_1 \rightarrow \text{Equation } \Delta_2$$

This begs the question: why would we opt to use weakenability of the \square modality (or, bother with the \square modality at all) to show that theories are weakenable, rather than using `weaken-eq` directly? Although the latter approach would indeed allow us to define the composition operations for effect theories defined above, the possible ways in which we can instantiate term metavariables remains too restrictive. That is, we still would not be able to prove the equality in Equation (1), despite the fact that we can weaken the `get-get` law so that it applies to programs that use the `Throw` effect as well. Instantiations of the term metavariable `k` will be limited to weakened effect trees, precluding any instantiation that use operations of effects other than `State`, such as `throw`.

Finally, we define the following predicate to witness that an equation is part of a theory.

$$\begin{aligned} _ \blacktriangleleft _ &: \square \text{Equation } \Delta \rightarrow \text{Theory } \Delta \rightarrow \text{Set}_1 \\ eq \blacktriangleleft T &= \exists \lambda a \rightarrow T . \text{equations } a \equiv eq \end{aligned}$$

We construct a proof $eq \blacktriangleleft T$ that an equation eq is part of a theory T by providing an arity together with a proof that T maps to eq for that arity.

3.5.4 Syntactic Equivalence of Effectful Programs

As discussed, propositional equality of effectful programs is too strict, as it precludes us from proving equalities that rely on a semantic understanding of the effects involved, such as the equality in Equation (1). The solution is to define an inductive relation that captures syntactic equivalence modulo some effect theory. We base our definition of syntactic equality of effectful programs on the relation defining equivalent computations by Yang and Wu [2021], Definition 3.1, adapting their definition where necessary to account for the use of modal necessity in the definition of [Theory](#).

$$\begin{aligned} \text{data } _ \approx \langle _ \rangle _ \{ \Delta \Delta' \} \{ _ : \Delta \lesssim \Delta' \} &: (m_1 : \text{Free } \Delta' A) \\ &\rightarrow \text{Theory } \Delta \\ &\rightarrow (m_2 : \text{Free } \Delta' A) \\ &\rightarrow \text{Set}_1 \text{ where} \end{aligned}$$

A value of type $m_1 \approx \langle T \rangle m_2$ witnesses that programs m_1 and m_2 are equal modulo the equations of theory T . The first three constructors ensure that it is an equivalence relation.

$$\begin{aligned} \approx\text{-refl} &: m \approx \langle T \rangle m \\ \approx\text{-sym} &: m_1 \approx \langle T \rangle m_2 \rightarrow m_2 \approx \langle T \rangle m_1 \\ \approx\text{-trans} &: m_1 \approx \langle T \rangle m_2 \rightarrow m_2 \approx \langle T \rangle m_3 \rightarrow m_1 \approx \langle T \rangle m_3 \end{aligned}$$

Then, we add the following congruence rule, that establish that we can prove equality of two programs starting with the

same operation by proving that the continuations yield equal programs for every possible value.

$$\begin{aligned}
\approx\text{-cong} : & \quad (op : \mathbf{Op} \Delta') \\
& \rightarrow (k_1 k_2 : \mathbf{Ret} \Delta' op \rightarrow \mathbf{Free} \Delta' A) \\
& \rightarrow (\forall x \rightarrow k_1 x \approx \langle T \rangle k_2 x) \\
& \rightarrow \mathbf{impure} (op, k_1) \approx \langle T \rangle \mathbf{impure} (op, k_2)
\end{aligned}$$

The final constructor allows to prove equality of programs by reifying equations of an effect theory.

$$\begin{aligned}
\approx\text{-eq} : & \quad (eq : \square \mathbf{Equation} \Delta) \\
& \rightarrow (px : eq \blacktriangleleft T) \\
& \rightarrow (vs : \mathbf{Vec Set} (\mathbf{V} (\square \langle eq \rangle))) \\
& \rightarrow (\gamma : \Gamma (\square \langle eq \rangle) vs) \\
& \rightarrow (k : \mathbf{R} (\square \langle eq \rangle) vs \rightarrow \mathbf{Free} \Delta' A) \\
& \rightarrow \quad (\mathbf{lhs} (\square \langle eq \rangle) vs \gamma \gg\gg k) \\
& \quad \approx \langle T \rangle (\mathbf{rhs} (\square \langle eq \rangle) vs \gamma \gg\gg k)
\end{aligned}$$

Since the equations of a theory are wrapped in the \square modality, we cannot refer to its components directly, but we must first provide a suitable extension witness.

With the $\approx\text{-eq}$ constructor, we can prove equivalence between the left-hand and right-hand side of an equation, sequenced with an arbitrary continuation k . For convenience, we define the following lemma that allows us to apply an equation where the sides of the equation do not have a continuation.

$$\begin{aligned}
\mathbf{use\text{-}equation} : & \quad \{ _ : \Delta \lesssim \Delta' \} \\
& \rightarrow \{ T : \mathbf{Theory} \Delta \} \\
& \rightarrow (eq : \square \mathbf{Equation} \Delta) \\
& \rightarrow eq \blacktriangleleft T \\
& \rightarrow (vs : \mathbf{Vec Set} (\mathbf{V} \square \langle eq \rangle)) \\
& \rightarrow \{ \gamma : \Gamma (\square \langle eq \rangle) vs \} \\
& \rightarrow \mathbf{lhs} (\square \langle eq \rangle) vs \gamma \approx \langle T \rangle \mathbf{rhs} (\square \langle eq \rangle) vs \gamma
\end{aligned}$$

The definition of $\mathbf{use\text{-}equation}$ follows readily from the right-identity law for monads, i.e., $m \gg\gg \mathbf{pure} \equiv m$, which allows us to instantiate $\approx\text{-eq}$ with \mathbf{pure} .

To construct proofs of equality it is convenient to use the following set of combinators to write proof terms in an equational style. They are completely analogous to the combinators commonly used to construct proofs of Agda’s propositional equality, for example, as found in PLFA [Wadler et al., 2020].

```

module  $\approx$ -Reasoning (T : Theory  $\Delta$ ) { _ :  $\Delta \lesssim \Delta'$  } where

begin _ : {m1 m2 : Free  $\Delta'$  A}
  → m1  $\approx$ < T > m2 → m1  $\approx$ < T > m2
begin eq = eq

_■ : (m : Free  $\Delta'$  A) → m  $\approx$ < T > m
m ■ =  $\approx$ -refl

_ $\approx$ <<>>_ : (m1 : Free  $\Delta'$  A) {m2 : Free  $\Delta'$  A}
  → m1  $\approx$ < T > m2 → m1  $\approx$ < T > m2
m1  $\approx$ <<>> eq = eq

_ $\approx$ <<_>>_ : (m1 {m2 m3} : Free  $\Delta'$  A)
  → m1  $\approx$ < T > m2 → m2  $\approx$ < T > m3 → m1  $\approx$ < T > m3
m1  $\approx$ << eq1 >> eq2 =  $\approx$ -trans eq1 eq2

```

We now have all the necessary tools to prove syntactic equality of programs modulo a theory of their effect. To illustrate, we consider how to prove the equation in Equation (1). First, we define a theory for the `State` effect containing the `get-get` law. While this is not the only law typically associated with `State`, for this example it is enough to only have the `get-get` law.

```

StateTheory : Theory State
arity StateTheory =  $\top$ 
equations StateTheory tt = get-get

```

Now to prove the equality in Equation (1) is simply a matter of invoking the `get-get` law.

```

get-get-throw :
  { _ : Throw  $\lesssim \Delta$  } { _ : State  $\lesssim \Delta$  }
  → (\get  $\gg$   $\lambda$  s → \get  $\gg$   $\lambda$  s' → \throw {A = A})
   $\approx$ < StateTheory > (\get  $\gg$   $\lambda$  s → \throw)

```

```

get-get-throw {A = A} = begin
  \get >>= (\s → \get >>= (\s' → \throw))
  ≈⟨⟨ use-equation get-get (tt , refl) (A :: []) ⟩⟩
  \get >>= (\s → \throw)
  ■
where open ≈-Reasoning StateTheory

```

3.5.5 Handler Correctness

A handler is correct with respect to a given theory if handling syntactically equal programs yields equal results. Since handlers are defined as algebras over effect signatures, we start by defining what it means for an algebra of an effect Δ to respect an equation of the same effect, adapting Definition 2.1 from the exposition of Yang and Wu [2021].

```

Respects : Alg Δ A → Equation Δ → Set1
Respects alg eq = ∀ {vs γ k} →
  fold k alg (lhs eq vs γ) ≡ fold k alg (rhs eq vs γ)

```

An algebra alg respects an equation eq if folding with that algebra produces propositionally equal results for the left- and right-hand side of the equation, for all possible instantiations of its type and term metavariables, and continuations k .

A handler H is correct with respect to a given theory T if its algebra respects all equations of T [Yang and Wu, 2021, Definition 4.3].

```

Correct : {P : Set} → Theory Δ → ⟨ A ! Δ ⇒ P ⇒ B ! Δ' ⟩ → Set1
Correct T H = ∀ {eq} → eq ◀ T → Respects (H .hdl) (extract eq)

```

We can now show that the handler for the `State` effect defined in Figure 8 is correct with respect to `StateTheory`. The proof follows immediately by reflexivity.

```

hStCorrect : Correct {A = A} {Δ' = Δ} StateTheory hSt
hStCorrect (tt , refl) {_ :: []} {γ = k} = refl

```

3.5.6 Theories of Higher-Order Effects

For the most part, equations and theories for higher-order effects are defined in the same way as for first-order effects and support many of the same operations. Indeed, the definition of equations ranging over higher-order effects is exactly the same as its first-order counterpart, the most major difference being that the left-hand and right-hand side are now defined as Hefty trees. To ensure compatibility with the use of type universes to avoid size-issues, we must also allow type metavariables to range over the types in a universe in addition to `Set`. For this reason, the set of type metavariables is no longer described by a natural number, but rather by a list of kinds, which stores for each type metavariable whether it ranges over a types in a universe, or an Agda `Set`.

```
data Kind : Set where set type : Kind
```

A `TypeContext` carries unapplied substitutions for a given set of type metavariables, and is defined by induction over a list of kinds.

```
TypeContext : List Kind → Set1
TypeContext []           = Level.Lift _ T
TypeContext (set :: vs) = Set × TypeContext vs
TypeContext (type :: vs) =
  Level.Lift (sl 0l) Type × TypeContext vs
```

Equations of higher-order effects are then defined as follows.

```
record EquationH (H : EffectH) : Set1 where
  field
    V      : List Kind
    Γ      : TypeContext V → Set
    R      : TypeContext V → Set
    lhs rhs : (vs : TypeContext V) → Γ vs → Hefty H (R vs)
```

This definition of equations suffers the same problem when it comes to term metavariables, which here too can only range

over programs that exhibit the exact effect that the equation is defined for. Again, we address the issue using an embedding of modal necessity to close over all possible extensions of this effect. The definition is analogous to the one in Section 3.5.2, but this time we use higher-order effect subtyping as the modal accessibility relation:

```
record  $\Box$  (P : EffectH → Set1) (H : EffectH) : Set1 where
  constructor necessary
  field  $\Box(\_)$  :  $\forall \{H'\} \rightarrow \{H \lesssim^H H'\} \rightarrow P H'$ 
```

To illustrate: we can define the *catch-return* law from the introduction of this section as a value of type \Box Equation^H Catch a follows. Since the `\catch` operation relies on a type universe to avoid size issues, the sole type metavariable of this equation must range over the types in this universe as well.

```
catch-return :  $\Box$  EquationH Catch
 $\forall$   $\Box$ ⟨ catch-return ⟩ = type :: []
 $\Gamma$   $\Box$ ⟨ catch-return ⟩ (lift t , _) =  $\llbracket t \rrbracket^T \times \text{Hefty } \_ \llbracket t \rrbracket^T$ 
 $R$   $\Box$ ⟨ catch-return ⟩ (lift t , _) =  $\llbracket t \rrbracket^T$ 
lhs  $\Box$ ⟨ catch-return ⟩ _ (x , m) = \catch (pure x) m
rhs  $\Box$ ⟨ catch-return ⟩ _ (x , m) = pure x
```

Theories of higher-order effects bundle extensible equations. The setup is the same as for theories of first-order effects.

```
record TheoryH (H : EffectH) : Set1 where
  field
    arity : Set
    equations : arity →  $\Box$  EquationH H
```

The following predicate establishes that an equation is part of a theory. We prove this fact by providing an arity whose corresponding equation is equal to *eq*.

```
_  $\triangleleft^H$  _ :  $\Box$  EquationH H → TheoryH H → Set1
eq  $\triangleleft^H$  Th =  $\exists \lambda a \rightarrow eq \equiv \text{equations } Th a$ 
```

Weakenability of theories of higher-order effects then follows from weakenability of its equations.

$$\begin{aligned} \text{weaken-}\Box &: \forall \{P\} \rightarrow \{H_1 \lesssim^H H_2\} \rightarrow \Box P H_1 \rightarrow \Box P H_2 \\ \Box \langle \text{weaken-}\Box \{w\} px \rangle \{w'\} &= \Box \langle px \rangle \{ \lesssim^H\text{-trans } w w' \} \end{aligned}$$

$$\begin{aligned} \text{weaken-theory}^H &: \{H_1 \lesssim^H H_2\} \rightarrow \text{Theory}^H H_1 \rightarrow \text{Theory}^H H_2 \\ \text{arity} \quad (\text{weaken-theory}^H Th) &= Th .\text{arity} \\ \text{equations} (\text{weaken-theory}^H Th) a &= \text{weaken-}\Box (Th .\text{equations } a) \end{aligned}$$

Theories of higher-order effects can be combined using the following sum operation. The resulting theory contains all equations of both argument theories.

$$\begin{aligned} _ \langle + \rangle^H _ &: \forall [\text{Theory}^H \Rightarrow \text{Theory}^H \Rightarrow \text{Theory}^H] \\ \text{arity} \quad (Th_1 \langle + \rangle^H Th_2) &= \text{arity } Th_1 \uplus \text{arity } Th_2 \\ \text{equations} (Th_1 \langle + \rangle^H Th_2) (\text{inj}_1 a) &= \text{equations } Th_1 a \\ \text{equations} (Th_1 \langle + \rangle^H Th_2) (\text{inj}_2 a) &= \text{equations } Th_2 a \end{aligned}$$

Theories of higher-order effects are closed under sums of higher-order effect theories as well. This operation is defined by appropriately weakening the respective theories, for which we need the following lemmas witnessing that higher-order effect signatures can be injected in a sum of signatures.

$$\begin{aligned} \lesssim\text{-}\dot{+}\text{-left} &: H_1 \lesssim^H (H_1 \dot{+} H_2) \\ \lesssim\text{-}\dot{+}\text{-right} &: H_2 \lesssim^H (H_1 \dot{+} H_2) \end{aligned}$$

The operation that combines theories under signature sums is then defined like so.

$$\begin{aligned} _ \langle + \rangle^H _ &: \text{Theory}^H H_1 \rightarrow \text{Theory}^H H_2 \rightarrow \text{Theory}^H (H_1 \dot{+} H_2) \\ Th_1 \langle + \rangle^H Th_2 &= \\ &\quad \text{weaken-theory}^H \{ \lesssim\text{-}\dot{+}\text{-left} \} Th_1 \\ &\quad \langle + \rangle^H \text{weaken-theory}^H \{ \lesssim\text{-}\dot{+}\text{-right} \} Th_2 \end{aligned}$$

3.5.7 Equivalence of Programs with Higher-Order Effects

We define the following inductive relation to capture equivalence of programs with higher-order effects modulo the equations of a given theory.

data $\cong \langle _ \rangle _ \{ _ : H_1 \lesssim^H H_2 \} :$ $(m_1 : \text{Hefty } H_2 \ A)$
 $\rightarrow \text{Theory}^H H_1$
 $\rightarrow (m_2 : \text{Hefty } H_2 \ A)$
 $\rightarrow \text{Set}_1 \text{ where}$

To ensure that it is indeed an equivalence relation, we include constructors for reflexivity, symmetry, and transitivity.

$\cong\text{-refl}$: $\forall \{m : \text{Hefty } H_2 \ A\}$
 $\rightarrow m \cong \langle Th \rangle m$

$\cong\text{-sym}$: $\forall \{m_1 : \text{Hefty } H_2 \ A\} \{m_2\}$
 $\rightarrow m_1 \cong \langle Th \rangle m_2$
 $\rightarrow m_2 \cong \langle Th \rangle m_1$

$\cong\text{-trans}$: $\forall \{m_1 : \text{Hefty } H_2 \ A\} \{m_2 \ m_3\}$
 $\rightarrow m_1 \cong \langle Th \rangle m_2 \rightarrow m_2 \cong \langle Th \rangle m_3$
 $\rightarrow m_1 \cong \langle Th \rangle m_3$

Furthermore, we include the following congruence rule that equates two program trees that have the same operation at the root, if their continuations are equivalent for all inputs.

$\cong\text{-cong}$: $(op : \text{Op}^H H_2)$
 $\rightarrow (k_1 \ k_2 : \text{Ret}^H H_2 \ op \rightarrow \text{Hefty } H_2 \ A)$
 $\rightarrow (s_1 \ s_2 : (\psi : \text{Fork } H_2 \ op) \rightarrow \text{Hefty } H_2 \ (\text{Ty } H_2 \ \psi))$
 $\rightarrow (\forall \{x\} \rightarrow k_1 \ x \cong \langle Th \rangle k_2 \ x)$
 $\rightarrow (\forall \{\psi\} \rightarrow s_1 \ \psi \cong \langle Th \rangle s_2 \ \psi)$
 $\rightarrow \text{impure } (op , k_1 , s_1) \cong \langle Th \rangle \text{impure } (op , k_2 , s_2)$

Finally, we include a constructor that equates two programs using an equation of the theory.

$$\begin{aligned}
\cong\text{-eq} & : (eq : \Box \text{Equation}^H H_1) \\
& \rightarrow eq \triangleleft^H Th \\
& \rightarrow (vs : \text{TypeContext } (\bigvee \Box \langle eq \rangle)) \\
& \rightarrow (\gamma : \Gamma \Box \langle eq \rangle vs) \\
& \rightarrow (k : R \Box \langle eq \rangle vs \rightarrow \text{Hefty } H_2 A) \\
& \rightarrow (\text{lhs } \Box \langle eq \rangle vs \gamma \ggg k) \\
& \quad \cong \langle Th \rangle (\text{rhs } \Box \langle eq \rangle vs \gamma \ggg k)
\end{aligned}$$

We can define the same reasoning combinators to construct proofs of equivalence for programs with higher-order effects.

`module $\cong\text{-Reasoning}$ { _ : $H_1 \lesssim^H H_2$ } (Th : $\text{Theory}^H H_1$) where`

$$\begin{aligned}
\text{begin } _ & : \{m_1 m_2 : \text{Hefty } H_2 A\} \\
& \rightarrow m_1 \cong \langle Th \rangle m_2 \rightarrow m_1 \cong \langle Th \rangle m_2 \\
\text{begin } eq = eq &
\end{aligned}$$

$$\begin{aligned}
_\blacksquare & : (c : \text{Hefty } H_2 A) \rightarrow c \cong \langle Th \rangle c \\
c \blacksquare = \cong\text{-refl} &
\end{aligned}$$

$$\begin{aligned}
\cong\langle\langle\rangle\rangle_ & : (m_1 : \text{Hefty } H_2 A) \{m_2 : \text{Hefty } H_2 A\} \\
& \rightarrow m_1 \cong \langle Th \rangle m_2 \rightarrow m_1 \cong \langle Th \rangle m_2 \\
c_1 \cong\langle\langle_\rangle\rangle eq = eq &
\end{aligned}$$

$$\begin{aligned}
\cong\langle\langle\rangle\rangle_ & : (c_1 \{c_2 c_3\} : \text{Hefty } H_2 A) \\
& \rightarrow c_1 \cong \langle Th \rangle c_2 \rightarrow c_2 \cong \langle Th \rangle c_3 \rightarrow c_1 \cong \langle Th \rangle c_3 \\
c_1 \cong\langle\langle eq_1 \rangle\rangle eq_2 = \cong\text{-trans } eq_1 eq_2 &
\end{aligned}$$

To illustrate, we can prove that the programs `catch throw` (`catch f m`) and `catch f m` are equal under a theory for the `afCatch` effect that contains the *catch-return* law.

$$\begin{aligned}
\text{catch-return-censor} & : \forall \{t : \text{Type}\} \{f\} \{x : \llbracket t \rrbracket^T\} \\
& \quad \{m : \text{Hefty } H \llbracket t \rrbracket^T\} \\
& \rightarrow \{ _ : \text{Catch } \lesssim^H H \} \\
& \rightarrow \{ _ : \text{Censor } \lesssim^H H \} \\
& \rightarrow \backslash \text{catch } (\text{pure } x) (\backslash \text{censor } f m) \\
& \quad \cong \langle \text{CatchTheory} \rangle \text{pure } x
\end{aligned}$$

```

catch-return-censor {f = f} {x = x} {m = m} =
begin
  \catch (pure x) (\censor f m)
  ≅⟨⟨ use-equationH catch-return (tt , refl) _ ⟩⟩
  pure x
  ■
  where open ≅-Reasoning _

```

The equivalence proof above makes, again, essential use of modal necessity. That is, by closing over all possible extensions of the `Catch` effect, the term metavariable in the `catch-return` law to range over programs that have higher-order effects other than `Catch`, which is needed to apply the law if the second branch of the `catch` operation contains the `censor` operation.

3.5.8 Correctness of Elaborations

As the first step towards defining correctness of elaborations, we must specify what it means for an algebra over a higher-order effect signature H to respect an equation. The definition is broadly similar to its counterpart for first-order effects in Section 3.5.5, with the crucial difference that the definition of “being equation respecting” for algebras over higher-order effect signatures is parameterized over a binary relation \approx between first-order effect trees. In practice, this binary relation will be instantiated with the inductive equivalence relation defined in Section 3.5.4; propositional equality would be too restrictive, since that does not allow us prove equivalence of programs using equations of the first-order effect(s) that we elaborate into.

$$\begin{aligned}
\text{Respects}^H : (& _ \approx _ : \forall \{A\} \rightarrow \text{Free } \Delta A \rightarrow \text{Free } \Delta A \rightarrow \text{Set}_1) \\
& \rightarrow \text{Alg}^H H (\text{Free } \Delta) \rightarrow \text{Equation}^H H \rightarrow \text{Set}_1 \\
\text{Respects}^H _ \approx _ \text{ alg eq} = & \\
\forall \{vs \ \gamma\} \rightarrow & \text{cata}^H \text{ pure alg (lhs eq vs } \gamma) \\
& \approx \text{cata}^H \text{ pure alg (rhs eq vs } \gamma)
\end{aligned}$$

Since elaborations are composed in parallel, the use of necessity in the definition of equations has additional consequences

for the definition of elaboration correctness. That is, correctness of an elaboration is defined with respect to a theory whose equations have left-hand and right-hand sides that may contain term metavariables that range over programs with more higher-order effects than those the elaboration is defined for. Therefore, to state correctness, we must also close over all possible ways these additional effects are elaborated. For this, we define the following binary relation on extensible elaborations.

```

record  $\sqsubseteq$  (e1 :  $\square$  (Elaboration H1)  $\Delta_1$ )
      (e2 :  $\square$  (Elaboration H2)  $\Delta_2$ ) : Set1 where
  field
    {  $\lesssim$ -eff } :  $\Delta_1 \lesssim \Delta_2$ 
    {  $\lesssim^H$ -eff } : H1  $\lesssim^H$  H2
  preserves-cases
    :  $\forall \{M\} (m : \llbracket H_1 \rrbracket^H M A)$ 
       $\rightarrow (e' : \forall [M \Rightarrow \text{Free } \Delta_2 ])$ 
       $\rightarrow \square \langle e_1 \rangle .\text{alg} (\text{map-sig}^H (\lambda \{x\} \rightarrow e' \{x\}) m)$ 
       $\equiv \text{extract } e_2 .\text{alg}$ 
         $(\text{map-sig}^H (\lambda \{x\} \rightarrow e' \{x\}) (\text{inj}^H \{X = A\} m))$ 

```

A proof of the form $e_1 \sqsubseteq e_2$ witnesses that the elaboration e_1 is included in e_2 . Informally, this means that e_2 may elaborate a bigger set of higher-order effects, for which it may need to refer to a bigger set of first-order effects, but for those higher-order effects that both e_1 and e_2 know how to elaborate, they should agree on how those effects are elaborated.

We then define correctness of elaborations as follows.

```

CorrectH : TheoryH H  $\rightarrow$  Theory  $\Delta$ 
            $\rightarrow \square$  (Elaboration H)  $\Delta \rightarrow$  Set1
CorrectH Th T e =
   $\forall \{\Delta' H'\}$ 
   $\rightarrow (e' : \square$  (Elaboration H')  $\Delta')$ 
   $\rightarrow \{ \_ : e \sqsubseteq e' \}$ 
   $\rightarrow \{ eq : \square$  EquationH  $\_ \}$ 
   $\rightarrow eq \blacktriangleleft^H Th$ 
   $\rightarrow$  RespectsH ( $\_ \approx \langle T \rangle \_$ ) (extract e')  $\square \langle eq \rangle$ 

```

Which is to say that an elaboration is correct with respect to a theory of the higher-order effects it elaborates (Th) and a theory of the first-order effects it elaborates into (T), if all possible extensions of said elaboration respect all equations of the higher-order theory, modulo the equations of the first-order theory.

Crucially, correctness of elaborations is preserved under composition of elaborations. Figure 11 shows the type of the corresponding correctness theorem in Agda; for the full details of the proof we refer to the Agda formalization accompanying this paper [Van der Rest and Bach, 2024]. We remark that correctness of a composed elaboration is defined with respect to the composition of the theories of the first-order effects that the respective elaborations use. Constructing a handler that is correct with respect to this composed first-order effect theory is a separate concern; a solution based on *fusion* is detailed in the work by Yang and Wu [2021].

```

compose-elab-correct : { _ : Δ1 • Δ2 ≈ Δ }
  → (e1 : □ (Elaboration H1) Δ1)
  → (e2 : □ (Elaboration H2) Δ2)
  → (T1 : Theory Δ1)
  → (T2 : Theory Δ2)
  → (Th1 : TheoryH H1)
  → (Th2 : TheoryH H2)
  → CorrectH Th1 T1 e1
  → CorrectH Th2 T2 e2
  → CorrectH (Th1 [+]H Th2) (compose-theory T1 T2)
    (compose-elab e1 e2)

```

Figure 11: The central correctness theorem establishing that correctness of elaborations is preserved under composition

3.5.9 Proving Correctness of Elaborations

To illustrate how the reasoning infrastructure build up in this section can be applied to verify correctness of elaborations,

we show how to verify the *catch-return* law for the elaboration `eCatch` defined in Section 3.3.4. First, we define the following syntax for invoking a known elaboration.

```

module Elab (e : □ (Elaboration H) Δ) where
  ℰ[ ] : Hefty H A → Free Δ A
  ℰ[ m ] = elaborate (extract e) m

```

When opening the module `Elab`, we can use the syntax `ℰ[m]` for elaborating `m` with some known elaboration, which helps to simplify and improve readability of equational proofs for higher-order effects.

Now, to prove that `eCatch` is correct with respect to a higher-order theory for the `Catch` effect containing the *catch-return* law, we must produce a proof that the programs `ℰ `catch (return x) m]` and `ℰ[return]` are equal (in the sense of the inductive equivalence relation defined in Section 3.5.4) with respect to some first-order theory for the `Throw` effect. In this instance, we do not need any equations from this underlying theory to prove the equality, but sometimes it is necessary to invoke equations of the underlying first-order effects to prove correctness of an elaboration.

```

eCatchCorrect : {T : Theory Throw}
  → CorrectH CatchTheory T eCatch
eCatchCorrect {Δ' = Δ'} e' {ζ} (tt , refl) {γ = x , m} =
  begin
    ℰ[ `catch (pure x) m ]
  ≈⟨⟨ from-≡ (sym $ ζ .preserves-cases _ ℰ[ ]) ⟩⟩
    ( # (given hThrow handle (pure x) $ tt)
      ≫≡ maybe' pure (ℰ[ m ] )
    )
  ≈⟨⟨ {- By definition of hThrow -}
    (pure (just x) ≫≡ maybe' pure ((ℰ[ m ] ≫≡ pure)))
  ⟩⟩
  ≈⟨⟨ {- By definition of ≫≡ -}
    ℰ[ pure x ]
  ⟩⟩
  ■
where

```

open \approx -Reasoning _
 open Elab e'

In the Agda formalization accompanying this paper [Van der Rest and Bach, 2024], we verify correctness of elaborations for the higher-order operations that are part of the 3MT library by Delaware et al. [2013c]. Figure 12 shows an overview of first-order and higher-order effects included in the development, and the laws which we prove about their handlers respectively elaborations.

| Effect | Laws | |
|-------------|--|--------------------------------|
| Throw | $\text{'throw} \gg k \equiv k$ | <i>bind-throw</i> |
| State | $\text{'get} \gg \lambda s \rightarrow \text{'get} \gg k s \equiv \text{'get} \gg k s s$ | <i>get-get</i> |
| | $\text{'get} \gg \text{'put} \equiv \text{pure } x$ | <i>get-put</i> |
| | $\text{'put } s \gg \text{'get} \equiv \text{'put } s \gg \text{pure } s$ | <i>put-get</i> |
| | $\text{'put } s \gg \text{'put } s' \equiv \text{'put } s'$ | <i>put-put</i> |
| Reader | $\text{'ask} \gg m \equiv m$ | <i>ask-query</i> |
| | $\text{'ask} \gg \lambda r \rightarrow \text{'ask} \gg k r \equiv \text{'ask} \gg \lambda r \rightarrow k r r$ | <i>ask-ask</i> |
| | $m \gg \lambda x \rightarrow \text{'ask} \gg \lambda r \rightarrow k x r \equiv \text{'ask} \gg \lambda r \rightarrow m \gg \lambda x \rightarrow k x r$ | <i>ask-bind</i> |
| LocalReader | $\text{'local } f (\text{pure } x) \equiv \text{pure } x$ | <i>local-pure</i> |
| | $\text{'local } f (m \gg k) \equiv \text{'local } f m \gg \text{'local } f \circ k$ | <i>local-bind</i> |
| | $\text{'local } f \text{'ask} \equiv \text{pure} \circ f$ | <i>local-ask</i> |
| | $\text{'local } (f \circ g) m \equiv \text{'local } g (\text{'local } f m)$ | <i>local-local</i> |
| Catch | $\text{'catch } (\text{pure } x) m \equiv \text{pure } x$ | <i>catch-pure</i> |
| | $\text{'catch } \text{'throw } m \equiv m$ | <i>catch-throw₁</i> |
| | $\text{'catch } m \text{'throw} \equiv m$ | <i>catch-throw₂</i> |
| Lambda | $\text{'abs } f \gg \lambda f' \rightarrow \text{'app } f' m \equiv m \gg f$ | <i>beta</i> |
| | $\text{pure } f \equiv \text{'abs } (\lambda x \rightarrow \text{'app } f (\text{'var } x))$ | <i>eta</i> |

Figure 12: Overview of effects, their operations, and verified laws in the Agda code

3.6 RELATED WORK

As stated in the introduction of this paper, defining abstractions for programming constructs with side effects is a research ques-

tion with a long and rich history, which we briefly summarize here. Moggi [1990] introduced monads as a means of modeling side effects and structuring programs with side effects; an idea which Wadler [1992] helped popularize. A problem with monads is that they do not naturally compose. A range of different solutions have been developed to address this issue [Jr., 1994, Jones and Duponcheel, 1993, Filinski, 1999, Cenciarelli and Moggi, 1993]. Of these solutions, monad transformers [Cenciarelli and Moggi, 1993, Liang et al., 1995b, Jaskelioff, 2008] is the more widely adopted solution. However, more recently, algebraic effects [Plotkin and Power, 2002] was proposed as an alternative solution which offers some modularity benefits over monads and monad transformers. In particular, whereas monads and monad transformers may “leak” information about the implementation of operations, algebraic effects enforce a strict separation between the interface and implementation of operations. Furthermore, monad transformers commonly require glue code to “lift” operations between layers of monad transformer stacks. While the latter problem is addressed by the Monatron framework of Jaskelioff [2008], algebraic effects have a simple composition semantics that does not require intricate liftings.

However, some effects, such as exception catching, did not fit into the framework of algebraic effects. *Effect handlers* [Plotkin and Pretnar, 2009b] were introduced to address this problem. Algebraic effects and handlers has since been gaining traction as a framework for modeling and structuring programs with side effects in a modular way. Several libraries have been developed based on the idea such as *Handlers in Action* [Kammar et al., 2013], the freer monad [Kiselyov and Ishii, 2015], or Idris’ Effects DSL [Brady, 2013a]; but also standalone languages such as Eff [Bauer and Pretnar, 2015], Koka [Leijen, 2017], Frank [Convent et al., 2020], and Effekt [Brachthäuser et al., 2020].⁵⁴

As discussed in Section 3.1.2 and Section 3.2.5, some modularity benefits of algebraic effects and handlers do not carry over to higher-order effects. Scoped effects and handlers [Wu et al.,

⁵⁴ A more extensive list of applications and frameworks can be found in Jeremy Yallop’s *Effects Bibliography*: <https://github.com/yallop/effects-bibliography>

2014, Piróg et al., 2018, Yang et al., 2022] address this shortcoming for *scoped operations*, as we summarized in Section 3.2.6. This paper provides a different solution to the modularity problem with higher-order effects. Our solution is to provide modular elaborations of higher-order effects into more primitive effects and handlers. We can, in theory, encode any effect in terms of algebraic effects and handlers. However, for some effects, the encodings may be complicated. While the complicated encodings are hidden behind a higher-order effect interface, complicated encodings may hinder understanding the operational semantics of higher-order effects, and may make it hard to verify algebraic laws about implementations of the interface. Our framework would also support elaborating higher-order effects into scoped effects and handlers, which might provide benefits for verification. We leave this as a question to explore in future work.

Although not explicitly advertised, some standalone languages, such as Frank [Convent et al., 2020] and Koka [Leijen, 2017] do have some support for higher-order effects. The denotational semantics of these features of these languages is unclear. A question for future work is whether the modular elaborations we introduce could provide a denotational model.

A recent paper by Van den Berg et al. [2021a] introduced a generalization of scoped effects that they call *latent effects* which supports a broader class of effects, including λ abstraction. While the framework appears powerful, it currently lacks a denotational model, and seems to require similar weaving glue code as scoped effects. The solution we present in this paper does not require weaving glue code, and is given by a modular but simple mapping onto algebraic effects and handlers.

Looking beyond purely functional models of semantics and effects, there are also lines of work on modular support for side effects in operational semantics [Plotkin, 2004]. Mosses' Modular Structural Operational Semantics [Mosses, 2004] (MSOS) defines small-step rules that implicitly propagate an open-ended set of *auxiliary entities* which encode common classes of effects, such as reading or emitting data, stateful mutation, and even control

effects [Sculthorpe et al., 2015]. The K Framework [Rosu and Serbanuta, 2010] takes a different approach but provides many of the same benefits. These frameworks do not encapsulate operational details but instead make it notationally convenient to program (or specify semantics) with side effects.

3.7 CONCLUSION

We have presented a new solution to the modularity problem with modeling and programming with higher-order effects. Our solution allows programming against an interface of higher-order effects in a way that provides effect encapsulation, meaning we can modularly change the implementation of effects without changing programs written against the interface and without changing the definition of any interface implementations. Furthermore, the solution requires a minimal amount of glue code to compose language definitions.

We have shown that the framework supports modular reasoning on a par with algebraic effects and handlers, albeit with some administrative overhead. While we have made use of Agda and dependent types throughout this paper, the framework should be portable to less dependently-typed functional languages, such as Haskell, OCaml, or Scala. An interesting direction for future work is to explore whether the framework could provide a denotational model for handling higher-order effects in standalone languages with support for effect handlers.

POSTSCRIPT

While this chapter describes an approach for modularly defining higher-order effects that is expressive enough for describing the side effects of most of the interpreters discussed in Section 2.5, the language fragment implementing polymorphic references remains beyond its capabilities. The reason is that it requires world-indexing to maintain well-formedness of references. While extending the approach to include support for

that and other intrinsically-typed effects is an interesting subject for future work, in the remainder of this thesis we will address a different problem. That is, by working in Agda we induce significant syntactic overhead into our definitions resulting from modularity. One way to deal with this overhead is to design a dedicated meta-language that includes modularity of definitions as part of its design. The second part of this thesis explores this approach.

Part II

META LANGUAGE DESIGN



TOWARDS A LANGUAGE FOR DEFINING REUSABLE PROGRAMMING LANGUAGE COMPONENTS

PREFACE

We begin our exploration of meta-language design for defining reusable programming language components with the language CS, which aims to provide a convenient surface syntax for working with extensible data types. The goal of designing this language is to gain insight into the programming abstractions that a meta-language could feature for defining reusable programming language components. While the language lacks a formal definition of its type system and semantics, it does have a prototype implementation.

4.1 INTRODUCTION

Developing programming languages is a difficult and time consuming task, which requires a lot of expertise from the language designer. One way to reduce the cost of language development is to build languages from *reusable programming components*, allowing language designers to grab off-the-shelf components for common language features. This approach has the potential to make language development both cheaper and more accessible, while producing specifications that allow us to understand the semantics of language features independent from the language

they are a part of. Modern functional programming languages, however, lack support for reuse of definitions, and as a result, language components built from algebraic data types and pattern matching functions cannot be reused without modifying or copying existing code. To illustrate the kind of language components we would like to define modularly, and where current functional programming languages fall short for this purpose, we consider the implementation of a tiny expression language and its interpreter in Haskell:

```
data Expr = Lit Int | Div Expr Expr
eval :: MonadFail m => Expr -> m Int
eval (Lit x)      = return x
eval (Div e1 e2) = do
  v1 <- eval e1
  v2 <- eval e2
  if v2 /= 0 then return (v1 'div' v2 else fail
```

The *Expr* data type declares an abstract syntax type with constructors for literals and division, and the function *eval* implements an interpreter for *Expr*. Importantly, *eval* needs to account for possible divisions by zero: evaluating *Div (Lit 10) (Lit 0)*, for example, should safely evaluate to a result that indicates a failure, without crashing. For this reason *eval* does not produce an *Int* directly, but rather wraps its result in an abstract monad *m* that encapsulates the side effects of interpretation. In this case, we only assume that *m* is a member of the *MonadFail* typeclass. The *MonadFail* class has one function, *fail*:

```
class MonadFail m where
  fail :: m a
```

We refer to functions, such as *fail*, that allow us to interact with an abstract monad as *operations*. We choose to factor language definitions this way, because it allows us to both define a completely new interpretation such as pretty printing or compilation for *Expr* by writing new functions *pretty* :: *Expr* -> *String* or *compile* :: *Expr* -> *m* [*Instr*], while also having the option to

change the implementation of existing semantics, by supplying alternative implementations for the `fail` operation. We can summarize this approach to defining language components with the following pipeline:

$$\text{Syntax} \xrightarrow{\text{denotation}} \text{Operations} \xrightarrow{\text{implementation}} \text{Result}$$

That is, a *denotation* maps *syntax* to an appropriate domain. In the definition of this domain, we distinguish between the type of the resulting value, and the side effects of computing this result, which are encapsulated in an abstract monad. We interact with this abstract monad using *operations*, and thus to extract a result we must supply a monad that *implements* all required operations.

What if we want to extend this language? To add new constructors to the abstract syntax tree, we must extend the definition of *Expr*, and modify *all* functions that match on *Expr* accordingly. Furthermore, the new clauses for these constructors may impose additional requirements on *m* for which we would need to add more typeclass constraints, and any existing instantiations of *m* would need to be updated to ensure that they are still a member of all required typeclasses.

Clearly, for these reasons *Expr* and *eval* in their current form do not work very well as a reusable language component. We introduce CS, a functional meta-language for defining reusable programming language components. The goal of CS is to provide a language in which one can define language components by defining data types and pattern matching functions, like *Expr* and *eval*, in such a way that we compose the syntax, interpretations, and effects of a language component without affecting existing definition. Importantly, we should also retain the possibility to add completely new interpretations for existing syntax by writing a new pattern matching function. In other words, CS should solve the *expression problem* [Wadler, 1998].

We can summarize this with following concrete design goals. In CS, one should be able to

- extend existing abstract syntax types with new constructors without having to modify existing definitions,
- extend existing denotations with clauses for new constructors, and define new semantics for existing syntax by defining new denotations,
- define abstract effect operations, and use these operations to implement denotation clauses without having to worry about the operations needed by other clauses, and
- define implementations for effect operations that are independent from the implementations of other operations.

There exist abstractions, such as *Data Types à la Carte* [Swierstra, 2008] and *Algebraic Effects and Handlers* [Plotkin and Pretnar, 2009a], that achieve the same goals. These provide the well-understood formalism on which CS is built. CS then provides a convenient surface syntax for working with these abstractions that avoids the overhead that occurs when encoding them in a host language like Haskell.

CS is work in progress. There is a prototype implementation of an interpreter and interactive programming environment which we can use to define and run the examples from this abstract. We are, however, still in the process of developing and implementing a type system. In particular, we should statically detect errors resulting from missing implementations of function clauses.

The name CS is an abbreviation of “CompositionalSemantics”. It is also the initials of Christopher Strachey, whose pioneering work [Strachey, 1966] initiated the development of denotational semantics. In *Fundamental Concepts in Programming Languages*, Strachey [2000] wrote that “the urgent task in programming languages is to explore the field of semantic possibilities”, and that we need to “recognize and isolate the central concepts” of programming languages. Today, five decades later, the words still ring true. The CS language aims to address this urgent task in programming languages, by supporting the definition of

reusable (central) programming language concepts, via compositional denotation functions that map the syntax of programming languages to their meaning.

4.2 CS BY EXAMPLE

In this section, we give an example-driven introduction to CS's features.

4.2.1 Data Types and Functions

CS is a functional programming language, and comes equipped with algebraic data types and pattern matching functions. We declare a new inductive data type for natural numbers as follows:

```
data Nat = Zero | Suc Nat
```

We can write functions over inductive data types by pattern matching, using a “pipe” (`|`) symbol to separate clauses:

```
fun double : Nat → Nat where
  | Zero   ↦ Zero
  | (Suc n) ↦ Suc (Suc (double n))
```

Not all types are user-declared: CS also offers built-in types and syntax for integers, lists, tuples, and strings.

```
fun length : List a → Int where
  | []       ↦ 0
  | (_ :: xs) ↦ 1 + length xs

fun zip : List a → List b → List (a * b) where
  | []       -       ↦ []
  | (x :: xs) (y :: ys) ↦ (x, y) :: zip xs ys
```

Both `length` and `zip` are polymorphic in the type of elements stored in the input list(s). Functions implicitly generalize over any free variables in their type signature.

4.2.2 *Effects and Handlers*

CS supports effectful programs by means of effects and handlers in the spirit of Plotkin and Pretnar [2009a], adapted to support higher-order operations. The key idea of the effects-and-handlers approach is to declare the syntax of effectful *operations*, and assign a semantics to these operations in a separate *handler*. Programs compute values and have side effects, and operations act as the interface through which these side effects are triggered.

We declare a new effect *Fail* with a single operation **fail** in CS as follows:

```
effect Fail where
  fail : {[Fail] a}
```

Effects in CS are declared with the **effect** keyword, and we declare its operations by giving a list of GADT-style signatures. In this case, the **fail** operation is declared to have type {[*Fail*] a}. We enclose the type of **fail** in braces ({-}) to indicate that the name **fail** refers to a *suspended computation* (Section 4.2.3). Suspended computations are annotated with an *effect row*, enclosed in square brackets ([-]), denoting the side effects of running the computation. Invoking the **fail** operation has *Fail* as a side effect.

We can use the *Fail* effect to implement a safe division function that invokes **fail** if the divisor is zero.

```
fun safeDiv : Int → Int → [Fail] Int where
  | x 0 ↦ fail!
  | x y ↦ ...
```

The postfix exclamation mark to **fail** is necessary to force the suspended computation. Here, we want to refer to the “action” of failing, rather than the computation itself, following Frank’s [Convent et al., 2020] separation between “being and doing”. We elaborate on this distinction in Section 4.2.3.

A function’s type signature must explicitly indicate its side effects. In this case, we annotate the return type of *safeDiv* with

the *Fail* effect to indicate that its implementation uses operations of this effect. Removing the annotation would make the above invocation of `fail` ill-typed. For functions that have no side effects, we may omit its row annotation: $a \rightarrow b$ is synonymous to $a \rightarrow [] b$

Handlers discharge an effect from annotations by assigning a semantics to its operations. For the *Fail* effect, we can do this by encoding exceptions in the *Maybe* type.

```
data Maybe a = Just a | Nothing
handler hFail : {[Fail|e] a} → {[e] (Maybe a)} where
  | fail    k ↦ {Nothing}
  | return x ↦ {Just x}
```

The handler *hFail* takes a value annotated with the *Fail* effect, and produces a *Maybe* value annotated with the remaining effects *e*. All free effect row variables in a signature, like *e*, are implicitly generalized over. When defining a handler we must provide a clause for each operation of the handled effect. Additionally, we must write a **return** clause that lifts pure values into the domain that encodes the effect's semantics. Operation clauses have a *continuation parameter*, *k*, with type $b \rightarrow [e] (Maybe a)$, which captures the remainder of the program starting from the current operation. Handlers may use the continuation parameter to decide how execution should resume after the current operation is handled. For example, when handling the `fail` operation we terminate execution of the program by ignoring this continuation.

We use the continuation parameter in a different way when defining a handler for a *State* effect, where $s : Set$ is a parameter of the module in which we define the effect.

```
effect State where
  | get : [State] s
  | put : s → [State] ()
```

```

handler  $hState : \{[State|e] a\} \rightarrow s \rightarrow \{[e] (a * s)\}$  where
| get       $st\ k \mapsto k\ st\ st$ 
| (put  $st'$ )  $st\ k \mapsto k\ ()\ st'$ 
| return  $x\ st \mapsto \{(x, st)\}$ 

```

For both the **get** and **put** operations, we use the continuation parameter k to implement the corresponding branch in $hState$. The continuation expects a value whose type corresponds to the return type of the current operation, and produces a computation with the same type as the return type of the handler. For the **put** operation, for example, this means that k is of type $() \rightarrow s \rightarrow \{[e] (a * s)\}$. The implementation of $hState$ for **get** and **put** then simply invokes k , using the current state as both the value and input state (**get**), or giving a unit value and using the given state st' as the input state (**put**). Effectively, this means that after handling **get** or **put**, execution of the program resumes, respectively with the same state or an updated state st' .

Handlers in CS are so-called *deep handlers*, meaning that they are automatically distributed over continuations. For the example above, this means that the *State* effect is already handled in the computation returned by k . The alternative is *shallow handlers*, in which case k would return a computation of type $\{[State|e] a\}$. When using shallow handlers, the programmer is responsible for recursively applying the handler to the result of continuations. While shallow handlers are more expressive, unlike deep handlers they are not guaranteed to correspond to a fold over the effect tree, meaning that they are potentially harder to reason about.

4.2.3 Order of Evaluation, Suspension, and Enactment

Inspired by Frank [Convent et al., 2020], CS allows effectful computations to be used as if they were pure values, without having to sequence them. Sub-expressions in CS are evaluated from left to right, and the side effects of computational sub-

expressions are evaluated eagerly in that order. For example, consider the following program:

```
fun f : Int → [Fail] Int where
  | n ↦ fail! + n
```

Here, we use the expression **fail!** (whose type is instantiated to $[Fail] Int$) as the first argument to $+$, where a value of type Int is expected. This is fine, because side effects that occur during evaluation of sub-terms are discharged to the surrounding context. That is, the side effects of evaluating computational sub-terms in the definition of f become side effects of f itself.

In practice, this means that function application in CS is not unlike programming in an *applicative style* in Haskell. For instance, when using the previously-defined handler $hFail$, which maps the *Fail* effect to a *Maybe*, we can informally understand the semantics of the CS program above as equivalent to the following Haskell program:

```
f :: Int → Maybe Int
f n = (+) <$> Nothing <*> pure n
```

Equivalently, we could write the following *monadic* program in Haskell, which makes the evaluation order explicit.

```
f :: Int → Maybe Int
f n = do x ← Nothing
      y ← pure n
      return (x + y)
```

CS's eager treatment of side effects means that effectful computations are not first-class values, in the sense that we cannot refer to an effectful computation without triggering its side effects. To treat computations as first-class, we must explicitly *suspend* their effects using *braces*:

```
fun f' : Int → {[Fail] Int} where
  | n ↦ {f n}
```

The function f' is no longer a function on Int that may fail, but instead a function from Int to a computation that returns and Int , but that could also fail. We indicate a suspended computation in types using braces ($\{\}/\}$), and construct a suspension at the term level using the same notation.

To *enact* the side effects of a suspended computation, we postfix it with an exclamation mark (!). For example, the expression $(f' 0)!$ has type $[Fail] Int$, whereas the expression $(f' 0)$ has type $\{[Fail] Int\}$. We see the same distinction with operations declared using the **effect** keyword. When we write **fail**, we refer to the operation in a descriptive sense, and we can treat it like any other value without having to worry about its side effects. When writing **fail!**, on the other hand, we are really performing the action of abruptly terminating: **fail** *is* and **fail!** *does*.

4.2.4 Modules and Imports

CS programs are organized using modules. Modules are delimited using the **module** and **end** keywords, and their definitions can be brought into scope elsewhere using the **import** keyword. All declarations—i.e., data types, functions, effects, and handlers—must occur inside a module.

```

module A where
  fun  $f : Int \rightarrow Int$  where
    |  $n \mapsto n + n$ 
end

module B where
  import A
  fun  $g : Int \rightarrow Int$  where
    |  $n \mapsto f\ n$ 
end

```

In addition to being an organizational tool, modules play a key role in defining and composing modular data types and functions.

4.2.5 Composable Data Types and Functions

In addition to plain algebraic data types and pattern matching functions, declared using the **data** and **fun** keywords, CS also supports case-by-case definitions of *extensible* data types and functions. In effect, CS provides a convenient surface syntax for working with DTC-style [Swierstra, 2008] definitions, which relies on an embedding of the initial algebra semantics [Goguen, 1976, Johann and Ghani, 2007] of data types to give a semantics to extensible and composable algebraic data types and functions, meaning that extensible functions have to correspond to a *fold* [Meijer et al., 1991]. In CS, one can program with extensible data types and functions in the same familiar way as with their plain, non-extensible counterparts.

The module system plays an essential role in the definition of composable data types and functions. That is, modules can inhabit a *signature* that declares the extensible types and functions for which that module can give a partial definition. In a signature declaration, we use the keyword **sort** to declare an extensible data type, and the **alg** keyword to declare an extensible function, or *algebra*. By requiring extensible functions to be defined as algebras over the functor semantics of extensible data types, we enforce *by construction* that they correspond to a fold.

As an example, consider the following signature that declares an extensible data type *Expr*, which can be evaluated to an integer using *eval*.

```
signature Eval where
  sort Expr : Set
  alg eval : Expr → Int
end
```

To give cases for *Expr* and *eval*, we define modules that inhabit the *Eval* signature.

```

module Lit : Eval where
  cons Lit : Int → Expr
  case eval (Lit x) ↦ x
end

module Add : Eval where
  cons Add : Expr → Expr → Expr
  case eval (Add x y) ↦ x + y
end

```

The **cons** keyword declares a new constructor for an extensible data type, where we declare any arguments by giving a GADT-style type signature. We declare clauses for functions that match on an extensible type using the **case** keyword. For every newly declared constructor of an extensible data type, we have an obligation to supply exactly one corresponding clause *for every extensible function that matches on that type*. CS has a coverage checker that checks whether modules indeed contain all necessary definitions, in order to rule out partiality resulting from missing patterns. For example, omitting the *eval* case from either the module *Lit* or *Add* above will result in a static error. Coverage is checked locally in modules, and preserved when composing signature instances.

In the definition of *eval* in the module *Add*, we see the implications of defining function clauses as algebras. We do not have direct control over recursive calls to *eval*. Instead, in **case** declarations, any recursive arguments to the matched constructor are replaced with the result of recursively invoking *eval* on them. In this case, this implies that *x* and *y* do not refer to expressions. Rather, if we invoke *eval* on the expression *Add e1 e2*, in the corresponding **case** declaration, *x* and *y* are bound to *eval e1* and *eval e2* respectively. We could encode the same example in Haskell as follows, but to use *eval* on concrete expressions additionally requires explicit definitions of a type level fixpoint and fold operation. In CS, this encoding layer is hidden by the language.

```
data Add e = Add e e
eval :: Add Int → Int
eval (Add x y) = x + y
```

To compose signature instances we merely have to import them from the same location.

```
module Program where
  import Lit, Add
  -- Evaluates to 3
  fun test : Int = eval (Add (Lit 1) (Lit 2))
end
```

By importing both the *Lit* and *Add* modules, the names *Expr* and *eval* will refer to the composition of the constructors/clauses defined in the imported signature instances. Here, this means that we can construct and evaluate expressions that consist of both literals and addition. Furthermore, to add a new constructor into the mix, we can simply define a new module that instantiates the *Eval* signature, and add it to the import list.

To define an alternative interpretation for *Expr*, we declare a new signature. In order to reference the **sort** declaration for *Expr*, we must import the *Eval* signature.

```
signature Pretty where
  import Eval -- brings 'Expr' into scope
  alg pretty : Expr → String
end
```

We declare cases for *pretty* by instantiating the newly defined signature, adding **import** declarations to bring relevant **cons** declaration into scope.

```
module PrettyAdd : Pretty where
  import Add -- brings 'Add' into scope
  case pretty (Add s1 s2) = s1 ++ " + " ++ s2
end
```

It is possible for two modules to be *conflicting*, in the sense that they both define an algebra case for the same constructor. This would happen, for example, if we were to define another module *PrettyAdd2* that also implements *pretty* for the constructor *Add*. Importing two conflicting modules should result in a type error, since the semantics of their composition is unclear.

4.3 DEFINING REUSABLE LANGUAGE COMPONENTS IN CS

In this section, we demonstrate how to use the features of CS introduced in the previous section to define reusable language components. We work towards defining a reusable component for function abstraction, which can be composed with other constructs, and for which we can define alternative implementations. As an example, we will show that we can use the same component defining functions with both a call-by-value and call-by-name strategy.

4.3.1 A Signature for Reusable Components

The first step is to define an appropriate module signature. We follow the same setup as for the *Eval* signature in Section 4.2.5. That is, we declare an extensible sort *Expr*, together with an algebra *eval* that consumes values of type *Expr*. The result of evaluation is a *Value*, with potential side effects *e*. The side effects are still abstract in the signature definition. The effect variable *e* is universally quantified, but instantiations of the signature *Eval* can impose additional constraints depending on which effects they need to implement *eval*. As a result, when invoking *eval*, *e* can be instantiated with any effect row that satisfies all the constraints imposed by the instances of *Eval* that are in scope.

```

signature Eval where
  sort Expr : Set
  alg eval : Expr → {[e] Value}
end

```

We will consider the precise definition of *Value* later in Section 4.3.3. For now, it is enough to know that it has a constructor $Num : Int \rightarrow Value$ that constructs a value from an integer literal.

4.3.2 A Language Component for Arithmetic Expressions

Let us start by defining instances of the *Eval* signature for the expression language from the introduction. First, we define a module for integer literals.

```

module Lit : Eval where
  cons Lit : Int → Expr
  case eval (Lit n) = {Num n}
end

```

The corresponding clause for *eval* simply returns the value n stored inside the constructor. Because the interpreter expects that we return a suspended computation, we must wrap n in a suspension, even though it is a pure value. Enacting this suspension, however, does not trigger any side effects, and as such importing *Lit* imposes no constraints on the effect row e .

Next, we define a module *Div* that implements integer division.

```

module Div : Eval where
  cons Div : Expr → Expr → Expr
  case eval (Div m1 m2) = {safeDiv m1! m2! }

```

Looking at the implementation of *eval* in the module *Div* we notice two things. First, the recursive arguments to *Div* have been replaced by the result of calling *eval* on them, meaning that m_1 and m_2 are now *computations* with type {[e] *Int*}, and hence we must use enactment before we can pass the result to

safeDiv. Enacting these computations may trigger side effects, so the order in which sub-expressions are evaluated determines in which order these side effects occur in the case that expressions contain more than one enactment. Sub-expressions in CS are evaluated from left to right. Second, the implementation uses the function *safeDiv*, defined in Section 4.2.2, which guards against errors resulting from division by zero.

The function *safeDiv* is annotated with the *Fail* effect, which supplies the **fail** operation. By invoking *safeDiv* in the definition of *eval*, which from the definition of *Eval* has type $Expr \rightarrow \{[e] Int\}$, we are imposing a constraint on the effect row *e* that it contains at least the *Fail* effect. In other words, whenever we import the module *Div* we have to make sure that we instantiate *e* with a row that has *Fail* in it. Consequently, before we can extract a value from any instantiation of *Eval* that includes *Div*, we must apply a handler for the *Fail* effect.

Since the interpreter now returns a *Value* instead of an *Int*, we must modify *safeDiv* accordingly. In practice this means that we must check if its arguments are constructed using the *Num* constructor before further processing the input. Since *safeDiv* already has *Fail* as a side effect, we can invoke the **fail** operation in case an argument was constructed using a different constructor than *Num*.

4.3.3 Implementing Functions as a Reusable Effect

CS's effect system can describe much more sophisticated effects than *Fail*. The effect system permits fine-grained control over the semantics of operations that affect a program's control flow, even in the presence of other effects. To illustrate its expressiveness, we will now consider how to define function abstraction as a reusable effect, and implement two different handlers for this effect corresponding to a call-by-value and call-by-name semantics. Implementing function abstraction as an effect is especially challenging since execution of the function body is deferred until the function is applied. From a handler's perspective, this

means that the function body and its side effect have to be postponed until a point beyond its own control, a pattern that is very difficult to capture using traditional algebraic effects.

We will see shortly how CS addresses this challenge. A key part of the solution is the ability to define *higher-order operations*: operations with arguments that are themselves effectful computations, leaving it up to the operation's handler to enact the side effects of higher-order arguments. The *Fun* effect, which implements function abstraction, has several higher-order operations.

effect *Fun* where

| | | |
|--------------|--|----------------------------------|
| lam | : <i>String</i> → {[<i>Fun</i>] <i>Value</i> } | → {[<i>Fun</i>] <i>Value</i> } |
| app | : <i>Value</i> → <i>Value</i> | → {[<i>Fun</i>] <i>Value</i> } |
| var | : <i>String</i> | → {[<i>Fun</i>] <i>Value</i> } |
| thunk | : {[<i>Fun</i>] <i>Value</i> } | → {[<i>Fun</i>] <i>Value</i> } |

The *Fun* effect defines four operations, three of which correspond to the usual constructs of the λ -calculus. The **thunk** operation has no counterpart in the λ -calculus, and postpones evaluation of a computation. It is necessary for evaluation to support both a call-by-value and call-by-name evaluation strategy

When looking at the **lam** and **thunk** operations, we find that they both have parameters annotated with the *Fun* effect. This annotation indicates that they are higher-order parameters. By allowing higher-order parameters to operations, effects in CS do not correspond directly to algebraic effects. Instead, to give as semantics to effects in CS, we must use a flavor of effects that permits higher-order syntax, such as *Latent Effects* [Van den Berg et al., 2021a].

As a result, any effects of the computations stored in a closure or **thunk** are postponed, leaving it up to the handler to decide when these take place.

USING THE *Fun* EFFECT To build a language with function abstractions that uses the *Fun* effect, we give an instance of the

Eval signature that defines the constructors *Abs*, *App*, and *Var* for *Expr*. We extend *eval* for these constructors by mapping onto the corresponding operation.

```

module Lambda : Eval where
  cons Abs : String → Expr → Expr
    | App : Expr → Expr → Expr
    | Var : String → Expr
  case eval (Abs x m) = lam x m
    | eval (App m1 m2) = app m1! (thunk m2!)!
    | eval (Var x) = var x
end

```

Crucially, in the case for *Abs* we pass the effect-annotated body *m*, with type $\{[e] \textit{Value}\}$, to the **lam** operation directly without extracting a pure value first. This prevents any effects in the body of a lambda from being enacted at the definition site, and instead leaves the decision of when these effects should take place to the used handler for the *Fun* effect. Similarly, in the case for *App*, we pass the function argument *m2* to the **thunk** operation directly, postponing any side effects until we force the constructed thunk. The precise moment at which we will force the thunk constructed for function arguments will depend on whether we employ a call-by-value or call-by-name strategy. We must, however, enact the side effects of evaluating the function itself (i.e., *m1*), because the **app** operation expects its arguments to be a pure value.

We implement call-by-value and call-by-name handlers for *Fun* in a new module, which also defines the type of values, *Value*, for our language. To keep the exposition simple, *Value* is not an extensible **sort**, but it is possible to do so in CS.

Values in this language are either numbers (*Num*), function closures (*Clo*), or thunked computations (*Thunk*). We define the type of values together in the same module as the handler(s) for the *Fun* effect. This module is parameterized over an effect row *e*, that denotes the *remaining effects* that are left after handling the *Fun* effect. In this case, *e* is a module parameter to express

```

handler hCBV : {[Fun|e] Value}
                → Env → {[Fail|e] Value} where
| (lam x f)                nv k ↦ k (Clo x f nv) nv
| (app (Clo x f nv') (Thunk t)) nv k ↦ k (f ((x, t!) :: nv'))! nv
| (app _ _)                _ _ ↦ {fail!}
| (var x)                  nv k ↦ k (lookup nv x)! nv
| (thunk f)                 nv k ↦ k (Thunk {f nv}) nv
| return v                  nv  ↦ {v}

```

that the remaining effects in the handlers that we will define coincide with the effect annotations of the computations stored in the *Clo* and *Thunk* constructors, allowing us to run these computations in the handler.

```

module HLambda (e : Effects) where
  import Fun
  type Env = List (String * Value)
  data Value = Num Int
            | Clo String (Env → {[Fail|e] Value}) Env
            | Thunk ({[Fail|e] Value})
  -- ... (handlers for the Fun effect) ...
end

```

CALL-BY-VALUE We are now ready to define a handler for the *Fun* effect that implements a call-by-value evaluation strategy. Figure 13 shows its implementation.

The **return** case is unremarkable: we simply ignore the environment *nv* and return the value *v*. The cases for **lam** and **thunk** are similar, as in both cases we do not enact the side effects associated with the stored computation *f*, but instead wrap this computation in a *Closure* or *Thunk* which is passed to the continuation *k*. For variables, we resolve the identifier *x* in the environment and pass the result to the continuation.

Figure 13: A Handler for the *Fun* effect, implementing a call-by-value semantics for function arguments. The gray highlights indicate where thunks constructed for function arguments are forced.

```

handler  $hCBN : \{[Fun|e] Value\}$ 
   $\rightarrow Env \rightarrow \{[Fail|e] Value\}$  where
| (lam  $x f$ )            $nv k \mapsto k (Clo\ x\ f\ nv)\ nv$ 
| (app  $(Clo\ x\ f\ nv')\ v$ )  $nv k \mapsto k (f\ ((x, v) :: nv'))!\ nv$ 
| (app  $- -$ )            $- - \mapsto \{fail!\}$ 
| (var  $x$ )            $nv k \mapsto$  match  $(lookup\ x\ nv)!$  with
  | (Thunk  $t$ )  $\mapsto k\ t!\ nv$ 
  |  $v$             $\mapsto k\ v\ nv$ 
  end
| (thunk  $f$ )          $nv k \mapsto k (Thunk\ \{f\ nv\})\ nv$ 
| return  $v$          $nv \mapsto \{v\}$ 

```

Figure 14: A Handler for the *Fun* effect, implementing a call-by-name semantics for function arguments. The gray highlights indicate where thunks constructed for function arguments are forced.

A call-by-value semantics arises from the implementation of the **app** case. The highlights (e.g., **t!**) indicate where the thunk we constructed for the function argument in *eval* is forced. In this case, we force this argument thunk immediately when encountering a function application, meaning that any side effects of the argument take place *before* we evaluate the function body.

CALL-BY-NAME The handler in Figure 14 shows an implementation of a call-by-name semantics for the *Fun* effect. The only case that differ from the call-by-value handler in Figure 13 are the **app** and **var** cases.

In the case for **app**, we now put the argument thunk in the environment immediately, without forcing it first. Instead, in the case for **var**, we check if the variable we look up in the environment is a *Thunk*. If so, we force it and pass the resulting value to the continuation. In effect, this means that for a variable that binds an effectful computation, the associated side effects take place every time we use that variable, but not until we reference it for the first time.

4.3.4 Example Usage

To illustrate how to use the reusable components defined in this section, and the difference between the semantics implemented by *hCBV* (Figure 13) and *hCBN* (Figure 14), we combine the *Lambda* module with the modules for *Div* and *Lit*. Figure 15 shows the example.

When importing the modules *HLambdaCBV* and *HLambdaCBN*, we pass an explicit effect row that corresponds to the effects that remain after handling the *Fun* effect. Because we handle *Fun* after handling the *Fail* effect introduced by *Div*, we pass the empty row. To evaluate expressions, we have to invoke *hFail* twice: first to handle the instance of the *Fail* effect introduced by *eval* for the *Div* constructor, and later to handle the *Fail* instance introduced by applying *hCBV/hCBN*. Consequently, the result of evaluating is a nested *Maybe*, where the inner instance indicates errors resulting from division by zero, and the outer instance errors thrown by the handler. Evaluating *result1* and *result2* shows the difference between using the call-by-value and call-by-name semantics for functions.

4.4 RELATED WORK

4.4.1 Effect Semantics

Monads, originally introduced by Moggi [1991] have long been the dominant approach to modelling programs with side effects. They are, however, famously hard to compose, leading to the development of *monad transformers* [Liang et al., 1995a] as a technique for building monads from individual definitions of effects. *Algebraic effects* [Plotkin and Power, 2003] provide a more structured approach towards this goal, where an effect is specified in terms of the *operations* that we can use to interact with it. The behaviour of these operations is governed by a set of equations that specify its well-behavedness. Later, Plotkin and Pretnar [Plotkin and Pretnar, 2009a] extended the

```

module Test where
  import Prelude
    ,Fun
    ,Fail
    ,HLambdaCBV []
    ,HLambdaCBN []
    ,Lambda
    ,Lit
    ,Div

  fun execCBV : Expr → Maybe (Maybe Value) where
    | e ↦ (hFail (hCBV (hFail (eval e)) []))!

  fun execCBN : Expr → Maybe (Maybe Value) where
    | e ↦ (hFail (hCBN (hFail (eval e)) []))!

  fun expr : Expr = App (Abs "x" (Lit 10)) (Div (Lit 5) (Lit 0))
    -- evaluates to Just Nothing

  fun result1 : Maybe (Maybe Value) = execCBV expr
    -- evaluates to Just (Just (Num 10))

  fun result2 : Maybe (Maybe Value) = execCBN expr
end

```

Figure 15: Examples of different outcomes when using a call-by-value or call-by-name evaluation strategy

approach with *handlers*, which define interpretations of effectful operations by defining a homomorphism from a *free model* that trivially inhabits the equational theory (i.e., syntax) to a programmer-defined domain, making the approach attractive for implementing effects as well. Perhaps the most well-known implementation of algebraic effects and handlers is the *free monad* [Kammar et al., 2013], and this implementation is often taken as the semantic foundation of languages with support for effect handlers. Schrijvers et al. [2019] showed that algebraic effects implemented using the free monad correspond to a subclass of monad-transformers. The algebraic effects and handlers approach provides a solid formal framework for understanding

effectful programs in which we intend to ground CS' semantics of effects and handlers.

A crucial difference between CS' effects and handlers, and the original formulation by Plotkin and Pretnar [2009a], is the support for *higher-order operations*. Although it is possible to implement such operations in algebraic effects by defining them as handlers, this breaches the separation between the syntax and semantics of effects that underpins CS' design. *Scoped Effects* [Wu et al., 2014] were proposed as an alternative flavor of algebraic effects that supports higher-order syntax, recovering a separation between the syntax semantics of effects for higher-order operations. In subsequent work, Piróg et al. [2018] adapted the categorical formulation of algebraic effects to give Scoped Effects a similar formal underpinning. Unfortunately, Scoped Effects is not suitable out-of-the-box as a model for effects and handlers in CS, because it cannot readily capture operations that arbitrarily postpone the execution of their higher-order arguments, such as `lam`. *Latent effects* were developed by Van den Berg et al. [2021a] as a refinement of scoped effects that solves this issue. Key to their approach is a *latent effect functor*, which explicitly tracks semantic residue of previously-installed handlers, allowing for a more fine-grained specification of the types of the computational arguments of operations. With Latent Effects, it is possible to capture function abstraction as a higher-order operation. It remains future work to formulate a precise model of effectful computation for CS, and to establish if and how CS' effect handlers correspond to Latent Effects.

4.4.2 Implementations of Algebraic Effects and Handlers

There are many languages with support for algebraic effects and handlers. Perhaps the most mature is Koka [Leijen, 2017], which features a Hindley/Milner-style row-polymorphic type system. While we borrow from Frank [Convent et al., 2020] a CBPV-inspired [Levy, 2004] distinction between computations and values, Koka is purely call-by-value, and only functions

can be effectful. Frank [Convent et al., 2020], on the other hand, does maintain this distinction between values and computations. Its type system relies on an *ambient ability* and implicit row polymorphism to approximate effects. Handlers are not first-class constructs in Frank. Instead, functions may adjust the ambient ability of their arguments by specifying the behaviour of operations. This provides some additional flexibility over built-in handlers, for example by permitting *multihandlers* that handle multiple effects at once. Both Koka and Frank lack native support for higher order effects, thus higher-order operations must be encoded in terms of handlers. This means that it is not possible to define higher order operations while maintaining the aforementioned distinction between the syntax and semantics of effects.

Eff [Bauer and Pretnar, 2015] is a functional language with support for algebraic effects and handlers, with the possibility to dynamically generate new operations and effects. In later work, Bauer and Pretnar [2014] developed a type-and-effect system for Eff, together with an inference algorithm [Pretnar, 2014]. The language Links [Lindley and Cheney, 2012] employs row-typed algebraic effects in the context of database programming. Their system is based on System F extended with effect rows and row polymorphism, and limits effectful computations to functions similar to Koka. Importantly, their system tracks effects using Rémy-style rows [Rémy, 1989], maintaining so-called *presence types* that can additionally express an effect’s absence from a computation. Brachthäuser et al. [2020] presented *Effekt* as a more practical implementation of effects and handlers, using *capability based* type system where effect types express a set of capabilities that a computation requires from its context.

4.4.3 Semantics of Composable Data Types and Functions

We give a semantics to extensible data types and functions in CS using the initial algebra semantics [Goguen, 1976, Johann and Ghani, 2007] of an underlying signature functor. *Data Types*

à la Carte (DTC) [Swierstra, 2008] solves the expression problem in Haskell by embedding this semantics into the host language. In later work, Bahr [2014] and Bahr and Hvitved [2011, 2012a] extended the approach to improve its expressiveness and flexibility.

DTC, like any approach that relies on initial algebra semantics, limits the modular definition of functions to functions that correspond to a *fold* over the input data. While this may seem restrictive, in practice more complicated traversals can often be encoded as a fold, such as paramorphisms [Meijer et al., 1991] or some classes of attribute grammars [Johnsson, 1987]. While CS currently only has syntax for plain algebras and folds, we plan to extend the syntax for working with extensible data types and functions to accommodate a wider range of traversals in the future.

Carette et al. [2009a] showed how to define interpreters by encoding object language expressions as a term that represents their traversal. These traversals correspond to a fold, but abstract over the algebra that defines the computation, meaning that alternative semantics can be assigned to terms by picking a suitable algebra. Semantics are defined as type class instances in the host language, thus one can build object language terms in a modular way by defining multiple different type classes that correspond to different syntactical constructs.

4.4.4 Row Types

While a concrete design for CS' type system is still emerging, we anticipate that it will make heavy use of *row types*, both for tracking effects and typing extensible types and functions. While to the best of our knowledge no type system exists with this combination of features, all the ingredients are there in the literature. Originally, row types were incepted as a means to model inheritance in object-oriented languages [Wand, 1989, Rémy, 1989], and later extensible records [Blume et al., 2006, Gaster and Jones, 1996]. More recently, they also gained popularity in

the form of row-based effect systems with the development of languages such as Koka [Leijen, 2017] and Links [Lindley and Cheney, 2012]. Their use for typing extensible algebraic data types and pattern matching functions is less well-studied. For the most part, row types in this context exist implicitly as part of encoding techniques such as DTC [Swierstra, 2008], where we can view the use of signature functors and functor co-products as an embedding of row-typed extensible variants in the host language’s type system. Various refinements of DTC [Morris, 2015, d. S. Oliveira et al., 2015, Bahr, 2014] make this connection more explicit by using type-level lists to track the composition of extensible data. A notable exception is the ROSE [Morris and McKinna, 2019a] language, which has a row-based type system with built-in support for extensible variants and pattern matching function.

4.5 FUTURE WORK

CS is an ongoing research project. Here, we briefly summarize the current state, and some of the challenges that still remain.

While we can implement the examples from this paper in the current prototype implementation of CS, the language still lacks a complete specification. As such, the immediate next steps are to develop specifications of the type system and operational semantics. This requires us to address several open research questions, such as giving a semantics to the flavor of higher-order effects used by CS, and applying row types to type CS’ extensible data types and functions. While the ROSE language supports extensible variants, this support is limited to non-recursive types. For CS, we would need to adapt their type system to support recursive extensible data types as well. Designing a small core calculus into which CS can be translated could be potential way to explore these questions, making a formalization of the language in a proof assistant more attainable, by formalizing the core language. Further down the line, we also intend to explore a denotational model for effect handlers

in CS, giving the language a more solid formal foundation, similar to existing programming languages based on algebraic effects and handlers.

In the future, we also hope to enforce stronger properties about specifications defined in CS through the language's type system. The prime example are *intrinsically-typed definitional interpreters* [Augustsson and Carlsson, 1999], which specify a language's operational semantics such that it is *type sound by construction*.

4.6 CONCLUSION

Reusable programming language components have the potential to significantly reduce the amount of time, effort, and expertise needed for developing programming languages. In this paper, we presented CS, a functional meta-language for defining reusable programming language components. CS enables the definition of reusable language components using algebraic data types and pattern matching functions, by supporting *extensible data types and functions*, which are defined on a case-by-case basis. Additionally, CS features built-in support for *effects and handlers* for defining the side effects of a language. The flavor of effects and handlers implemented by CS supports *higher-order* operations, and can be used to define features that affect a program's control flow, such as function abstraction, as a reusable effect. We illustrated how these features can be used for developing reusable programming language components by defining a component for function abstraction, which can be composed with other language components and evaluated using both a call-by-value and call-by-name strategy.

POSTSCRIPT

The language design we discussed in this section provides us for a programming model for working with extensible data types. What is clearly missing is a formal specification of a typing

discipline. Crucially, such a typing discipline should guarantee that all extensions of existing data are safe, in the sense that the addition of new constructors to a data type should not result in a scenario where functions over an extensible data type can be applied to inputs on which they are not defined. When embedding modular data types in functional languages using the Data Types á la Carte approach, type safe modularity is enforced by the host languages type system using a de facto shallow embedding of *row-typed algebraic data types*. This already provides us with some hints for how a suitable type system could be designed. In the next chapter, we design such a type system for a small calculus featuring primitives for programming with modular data types.



TYPES AND SEMANTICS FOR EXTENSIBLE DATA TYPES

PREFACE

In this chapter, we work towards providing a formal foundation for the language design proposed in Chapter 5. We do this by developing a core calculus that features primitives for working with extensible data types, together with a type system and (categorical) semantics. A potential way to connect the work from this and the previous chapter would be to define a desugaring from programs in CS to the calculus presented in this chapter, although that is at this point still future work.

5.1 INTRODUCTION

A common litmus test for a programming language's capability for modularity is whether a programmer is able to extend existing data with new ways to construct it as well as to add new functionality for this data. All in a way that preserves static type safety; a conundrum which Wadler [1998] dubbed the *expression problem*. When working in pure functional programming languages, another modularity question is how to model side effects modularly using, e.g., *monads* [Moggi, 1991]. Ideally, we would keep the specific monad used to model the effects of a program abstract and program against an *interface* of effectful

operations instead, defining the syntax and implementation of such interfaces separately and in a modular fashion.

The traditional approach for tackling these modularity questions in pure functional programming languages is by embedding the *initial algebra semantics* [Goguen, 1976] of inductive data types in the language's type system. By working with such embeddings in favor of the language's built-in data types we gain modularity without sacrificing type safety. This approach was popularized by Swierstra's *Data Types à la Carte* [Swierstra, 2008] as a solution to the expression problem, where it was used to derive modular interpreters for a small expression language. In later work, similar techniques were applied to define the syntax and implementation of a large class of monads using (algebraic) effects and handlers based on different flavors of inductively defined *free monads*. This was shown to be an effective technique for modularizing both first order [Kammar et al., 2013] and higher-order [Wu et al., 2014, Bach and Van der Rest, 2023, Van den Berg et al., 2021a] effectful computations.

The key idea that unifies these techniques is the use of *signature functors*, which act as a de facto syntactic representation of an inductive data type or inductively defined free monad. Effectively, this defines a generic inductive data type or free monad that takes its constructors as a parameter. The crucial benefit of this setup is that we can compose data types and effects by taking the coproduct of signature functors, and we can compose function cases defined over these signature functors in a similarly modular way. Inductive data types and functions in mainstream functional programming languages generally do not support these kinds of composition.

While embedding signature functors has proven itself as a tremendously useful technique for enhancing functional languages with a higher degree of type safe modularity, the approach has some downsides:

- Encodings of a data type's initial algebra semantics lacks the syntactic convenience of native data types, especially when it comes to constructing and pattern matching on

values. Further overhead is introduced by their limited interoperability, which typically relies on user-defined isomorphisms.

- The connection between initial algebra semantics encodings of data types, and the mathematical concepts that motivate them remains implicit. This has two drawbacks: (1) the programmer has to write additional code witnessing that their definitions possess the required structure (e.g., by defining instances of the `Functor` typeclass), and (2) a compiler cannot leverage the properties of this structure, such as by implementing (provably correct) optimizations based on the well-known map and fold fusion laws.

In this paper, we explore an alternative perspective by making type-safe modularity part of the language’s design, by including built-in primitives for the functional programmer’s modularity toolkit—e.g., functors, folds, fixpoints, etc. We believe that this approach has the potential to present the programmer with more convenient syntax for working with extensible data types (see, for example, the language design proposed by [Van der Rest and Bach \[2022\]](#)). Furthermore, by supporting type-safe modularity through dedicated language primitives, we open the door for compilers to benefit from their properties, for example by applying fusion based optimizations.

5.1.1 Contributions

The semantics of (nested) algebraic data types has been studied extensively in the literature (e.g., by [Johann et al. \[2021\]](#), [Johann and Polonsky \[2019\]](#), [Johann and Ghiorzi \[2021\]](#), and [Abel and Matthes \[2002\]](#), [Abel et al. \[2003, 2005\]](#)) resulting in the development of various calculi with the purpose of studying different aspects of the semantics of programming with algebraic data types. In this paper, we build on these works to develop a core calculus that seeks to distill the essential language features needed for developing programming languages

with built-in support for type-safe modularity while retaining the same formal foundations. Although the semantic ideas that we build on to develop our calculus are generally well-known, their application to improving the design of functional programming languages has yet to be explored in depth. It is still future work to leverage the insights gained by developing this calculus in the design of programming language that provide better ergonomics for working with extensible data types, but we believe the development of a core calculus capturing the essentials of programming with extensible data types to be a key step for achieving this goal. To bridge from the calculus presented in this paper to a practical language design, features such as *smart constructors*, *row types*, and (*functor*) *subtyping* (as employed, for example, by [Morris and McKinna \[2019a\]](#) and [Hubers and Morris \[2023\]](#)) would be essential. We make the following technical contributions:

- We show (in Section [5.2](#)) how modular functions over algebraic data types in the style of Data Types à la Carte and modular definitions of first-order and higher-order (algebraic) effects and handlers based on inductively defined free monads can be captured in the calculus.
- We present (in Section [5.3](#)) a formal definition of the syntax and type system.
- We give (in Section [5.4](#)) a categorical semantics for our calculus.
- We present (in Section [5.5](#)) an operational semantics for our calculus, and discuss how it relates to the categorical semantics.

Section [5.6](#) discusses related work, and Section [5.7](#) concludes.

5.2 PROGRAMMING WITH EXTENSIBLE DATA TYPES, BY EXAMPLE

The basis of our calculus is the polymorphic λ -calculus extended with kinds and restricted to rank-1 polymorphism, allowing the definition of many familiar polymorphic functions, such as $(\text{id} : \forall\alpha.\alpha \Rightarrow \alpha) = \lambda x.x$ or $(\text{const} : \forall\alpha.\forall\beta.\alpha \Rightarrow \beta \Rightarrow \alpha) = \lambda x.\lambda y.x$. Types are closed under products and coproducts, with the unit type $(\mathbb{1})$ and empty type (\emptyset) acting as their respective units. Furthermore, we include a type-level fixpoint (μ) , which can be used to encode many well-known algebraic data types. For example, the familiar type of lists is encoded as $\text{List} \triangleq \lambda\alpha.\mu(\lambda X.\mathbb{1} + (\alpha \times X))$. A key feature of the calculus is that all higher-order types (i.e., that have one or more type argument) are, by construction, functorial in all their arguments. While this imposes some restrictions on the types we can define, it also means that the programmer gets access to primitive mapping and folding operations that they would otherwise have to define themselves. For the type List , for example, this means that we get both the usual mapping operation transforming its elements, as well as an operation corresponding to Haskell's `foldr`, for `free`.

Although the mapping and folding primitives for first-order type constructors (i.e., those taking arguments of kind \star and producing a type of kind \star) are already enough to solve the expression problem for regular algebraic data types (Section 5.2.2) and to encode modular algebraic effects (Section 5.2.3), they can readily be generalized to higher-order type constructors. That is, type constructors that construct higher-order types from higher-order types. The benefit of this generalization is that our calculus can also capture the definition of so-called *nested data types* [Bird and Meertens, 1998], which arise as the fixpoint of a *higher-order functor*. We make essential use of the calculus' higher-order capabilities in Section 5.2.4 to define modular handlers for scoped effects [Yang et al., 2022] and modular elaborations for higher-order effects [Bach and Van der Rest, 2023],

as in both cases effect trees that represents monadic programs with higher-order operations is defined as a nested data type.

5.2.1 Notation

All code examples in this section directly correspond to programs in our calculus, but we take some notational liberty to simplify the exposition. Abstraction and application of type variables is left implicit. Similarly, we omit first-order universal quantifications. By convention, we denote type variables bound by type-level λ -abstraction using capital letters (e.g., X), and those bound by universal quantification using Greek letters (e.g., α, β).

5.2.2 Modular Interpreters in the style of Data Types *à la Carte*

We consider how to define a modular interpreter for a small expression language of simple arithmetic operations. For starters, we just include literals and addition. The corresponding BNF equation and signature functor are given below:

$$e ::= \mathbb{N} \mid e + e \qquad \text{Expr} \triangleq \lambda X. \mathbb{N} + (X \times X)$$

Now, we can define an eval that maps expressions—given by the fixpoint of Expr—to their result:

$$\begin{aligned} \text{expr} &: \mathbb{N} + (\mathbb{N} \times \mathbb{N}) \Rightarrow \mathbb{N} \\ \text{expr} &= (\lambda x. x) \blacktriangledown (\lambda x. \pi_1 x + \pi_2 x) \end{aligned}$$

$$\begin{aligned} \text{eval} &: \mu(\text{Expr}) \Rightarrow \mathbb{N} \\ \text{eval} &= \langle \! \langle \text{expr} \rangle \! \rangle^{\text{Expr}} \end{aligned}$$

Terms typeset in **purple** are built-in operations. π_1 and π_2 are the usual projection functions for products, and $- \blacktriangledown -$ is an eliminator for coproducts. Following Meijer et al. [1991], we write $\langle \! \langle \text{alg} \rangle \! \rangle^\tau$ (i.e., “banana brackets”) to denote a fold over the type $\mu(\tau)$ with an *algebra* of type $\text{alg} : \tau \tau' \Rightarrow \tau'$. The calculus

does not include a general term level fixpoint; the only way to write a function that recurses on the substructures of a μ -type is by using the built-in folding operation. While this limits the operations we can define for a given type, it also ensures that all well-typed terms in the calculus have a well-defined semantics.

Now, we can extend this expression language with support for a multiplication operation as follows, where $\text{Mul} \triangleq \lambda X.X \times X$:

$$\begin{aligned} \text{mul} : \mathbb{N} \times \mathbb{N} &\Rightarrow \mathbb{N} & \text{eval} : \mu(\text{Expr} + \text{Mul}) &\Rightarrow \mathbb{N} \\ \text{mul} = \lambda x.\pi_1 \ x * \pi_2 \ x & & \text{eval} = \langle \text{expr} \ \nabla \ \text{mul} \rangle^{\text{Expr} + \text{Mul}} \end{aligned}$$

5.2.3 Modular Algebraic Effects using the Free Monad

As our second example we consider how to define modular algebraic effects and handlers [Plotkin and Pretnar, 2009a] in terms of the free monad following Swierstra [2008]. First, we define the Free type which constructs a free monad for a given signature functor f . We can think of a term with type $\text{Free } f \ \alpha$ as a syntactic representation of a monadic program producing a value of type α with f describing the operations which we can use to interact with the monadic context.

$$\text{Free} : (\star \rightsquigarrow \star) \rightsquigarrow \star \rightsquigarrow \star \quad \triangleq \quad \lambda f.\lambda \alpha.\mu(\lambda X.\alpha + fX)$$

Note that the type Free is actually a functor in both its arguments, and thus there are two ways to “map over” a value of type $\text{Free } f \ \alpha$; we can transform the values at the leaves using a function $\alpha \Rightarrow \beta$, or the shape of the nodes using a natural transformation $\forall \alpha.f \ \alpha \Rightarrow g \ \alpha$. The higher-order map can be used, for example, for defining function that reorders the operations of effect trees with a composite signature.

$$\begin{aligned} \text{reorder} : \text{Free } (f + g) \ \alpha &\Rightarrow \text{Free } (g + f) \ \alpha \\ \text{reorder} = \text{map} \langle \iota_2 \ \nabla \ \iota_1 \rangle^{\text{Free}} \end{aligned}$$

Here, we use higher-order instances at kind $\star \rightsquigarrow \star$ of the co-product eliminator $-\ \nabla \ -$, the coproduct injection functions ι_1 , ι_2 , and the functorial map operation $\text{map} \langle - \rangle^-$.

Effect handlers can straightforwardly be implemented as folds over `Free`. In fact, the behavior of a handler is entirely defined by the algebra that we use to fold over the effect tree, allowing us write a generic handle function:

$$\begin{aligned} \text{handle} &: (\alpha \Rightarrow \beta) \Rightarrow (f (\text{Free } g \beta) \Rightarrow \text{Free } g \beta) \\ &\Rightarrow \text{Free } (f + g) \alpha \Rightarrow \text{Free } g \beta \\ \text{handle} &= \lambda h. \lambda i. \lambda \mathbf{i}. \left((\mathbf{in} \circ \mathbf{i}_1 \circ h) \blacktriangledown i \blacktriangledown (\mathbf{in} \circ \mathbf{i}_2) \right)^{\alpha + (fX) + (gX)} \end{aligned}$$

Here, \mathbf{in} is the constructor of a type-level fixpoint (μ). The fold above distinguishes three cases: (1) pure values, in which case we return it again using the function h ; (2) an operation of the signature f which is handled using the function i ; or (3) an operation of the signature g which is preserved by reconstructing the effect tree and doing nothing.

As an example, we consider how to implement a handler for the `Abort` effect, which has a single operation indicating abrupt termination of a computation. We define its signature functor as follows:

$$\text{Abort} : * \rightsquigarrow * \triangleq \lambda X. \mathbb{1}$$

The definition of `Abort` ignores its argument, X , which is the type of the continuation. After aborting a computation, there is no continuation, thus the `Abort` effect does not need to store one. A handler for `Abort` is then defined like so, invoking the generic handle function defined above:

$$\begin{aligned} \text{hAbort} &: \text{Free } (\text{Abort} + f) \alpha \Rightarrow \text{Free } f (\text{Maybe } \alpha) \\ \text{hAbort} &= \text{handle Just } (\lambda x. \mathbf{in} (\mathbf{i}_1 \text{ Nothing})) \end{aligned}$$

5.2.4 Modular Higher-Order Effects

To describe the syntax of computations that interact with their monadic context through higher-order operations (that is, operations whose arguments can themselves also be monadic computations) we need to generalize the free monad as follows.

$$\text{Prog} : ((* \rightsquigarrow *) \rightsquigarrow * \rightsquigarrow *) \rightsquigarrow * \rightsquigarrow * \triangleq \lambda f. \mu (\lambda X. \lambda \alpha. \alpha + (f X \alpha))$$

Note that, unlike the Free type, Prog is defined as the fixpoint of a higher-order functor. This generalization allows for signature functors to freely choose the return type of continuations. Following [Yang et al. \[2022\]](#), we use this additional expressivity to describe the syntax of higher-order operations by nesting continuations. For example, the following defines the syntax of an effect for exception catching, that we can interact with by either throwing an exception, or by declaring an exception handler that first executes its first argument, and only runs the second computation if an exception was thrown.

$$\text{Catch} : (\star \rightsquigarrow \star) \rightsquigarrow \star \rightsquigarrow \star \triangleq \lambda X. \lambda \alpha. \mathbb{1} + (X(X\alpha) \times (X(X\alpha)))$$

A value of type $\text{Prog Catch } \alpha$ is then a syntactic representation of a monadic program that can both throw and catch exceptions. From this syntactic representation we can proceed in two different ways. The first option is to replace exception catching with an application of the `hAbort` handler, in line with [Plotkin and Pretnar \[2009a\]](#) original strategy for capturing higher-order operations. In recent work, [Bach and Van der Rest \[2023\]](#) demonstrated how such abbreviations can be made modular and reusable by implementing them as algebras over the Prog type. Following their approach, we define the following elaboration of exception catching into a first-order effect tree.

$$\text{eCatch} : \text{Prog Catch } \alpha \Rightarrow \text{Free Abort } \alpha$$

$$\begin{aligned} \text{eCatch} = & \left(\text{in} \circ \iota_1 \right) \\ & \blacktriangledown \left(\text{in} \circ \iota_2 \right) \\ & \blacktriangledown \left(\lambda x. \text{hAbort } (\pi_1 x) \gg \text{maybe } (\text{join } (\pi_2 x)) \text{ id} \right) \\ & \left. \right)_{\alpha + \text{Catch } X \ \alpha} \end{aligned}$$

Here, the applications of monadic bind (\gg) and join refer to the monadic structure of Free. Alternatively, we can define a

handler for exception catching directly by folding over the `Prog` type, following the *scoped effects* approach by [Wu et al. \[2014\]](#):

$$\begin{aligned} \text{hCatch} &: \text{Prog} (\text{Catch} + \text{h}) \alpha \Rightarrow \text{Prog} \text{h} (\text{Maybe} \alpha) \\ \text{hCatch} &= \llbracket (\text{in} \circ \iota_1 \circ \text{Just}) \\ &\quad \blacktriangledown (\lambda x. \text{in} (\iota_1 \text{Nothing})) \\ &\quad \blacktriangledown (\lambda x. \pi_1 x \gg= \text{maybe} (\pi_2 x \gg= \text{fwd}) \text{id}) \\ &\quad \blacktriangledown (\text{in} \circ \iota_2) \rrbracket^{\alpha + (\text{Catch} \times \alpha) + (\text{h} \times \alpha)} \end{aligned}$$

Where the function `fwd` establishes that `Maybe` commutes with the `Prog` type in a suitable way:

$$\text{fwd} : \text{Maybe} (\text{Prog} \text{h} (\text{Maybe} \alpha)) \Rightarrow \text{Prog} \text{h} (\text{Maybe} \alpha)$$

That is, we show that `Prog h` is a *modular carrier* for `Maybe` [[Schrijvers et al., 2019](#)].

As demonstrated, our calculus supports defining higher-order effects and their interpretations. To conveniently sequence higher-order computations we typically also want to use a monadic bind function, such as $\gg= : \text{Prog} \text{h} \alpha \rightarrow (\alpha \rightarrow \text{Prog} \text{h} \beta) \rightarrow \text{Prog} \text{h} \beta$. While it is possible to define monadic bind for *Free* from Section 5.2.3 in terms of a plain fold, defining the monadic bind for *Prog* generally requires a *generalized fold* [[Bird and Paterson, 1999](#), [Yang et al., 2022](#)]. Adding this and other recursion principles [[Meijer et al., 1991](#)] to our calculus is future work.

5.3 THE CALCULUS

The previous section demonstrated how a language with built-in support for functors, folds, and fixpoints provides support for defining and working with state-of-the-art techniques for type safe modular programming. In this section we present a core calculus for such a language. The basis of our calculus is the first-order fragment of System F^ω —i.e., the polymorphic λ -calculus with kinds, where universal quantification is limited to prenex normal form à la Hindley-Milner. Additionally, the syntax of

$$\begin{array}{lcl}
& & \alpha, \beta, \gamma, X, Y \in \text{String} \\
\\
\textit{Kind} & \ni & k ::= \star \mid k \rightsquigarrow k \\
\textit{KindEnv} & \ni & \Delta, \Phi ::= \emptyset \mid \Delta, \alpha : k \\
\\
\textit{Type} & \ni & \tau ::= \alpha \mid X \mid \tau \tau \mid \lambda X. \tau \mid \mu(\tau) \mid \tau \Rightarrow \tau \\
& & \mid \mathbb{1} \mid \tau \times \tau \mid \tau + \tau \\
\textit{Scheme} & \ni & \sigma ::= \forall \alpha. \sigma \mid \tau
\end{array}$$

Figure 16: Type syntax

types, defined in Figure 16, includes primitives for constructing recursive types ($\mu(-)$), products (\times) and coproducts ($+$), as well as a unit type ($\mathbb{1}$) and empty type (\emptyset). In the definition of the syntax of types, the use of \forall -types is restricted by stratifying the syntax into two layers, types and type schemes. Consequently, our calculus is, by design, *predicative*: \forall -types can quantify over types but not type schemes.

The motivation for this predicative design is that it permits a relatively straightforward categorical interpretation of \forall -types in terms of *ends* (see Section 5.4.2.3). Whereas the restriction of universal quantification to prenex normal form is usually imposed to facilitate type inference, our calculus does not support inference in its current form due to the structural treatment of data types. In a structural setting, inference requires the reconstruction of (recursive) data type definitions from values, which is, in general, not possible.

We remark that the current presentation of the type system is *declarative*, meaning certain algorithmic aspects crucial for type checking, such as normalization and equality checking of types, are not covered in the current exposition. Regarding decidability of the type system: our system is a subset of System F_ω , whose Church-style formulation is decidable while its Curry-style formulation is not. As such, we expect our type system to inherit these properties. Since we are restricting ourselves to

| | | |
|--|--|---|
| $\Delta \mid \Phi \vdash \tau : k$ | | |
| $\frac{\text{K-VAR} \quad k : \alpha \in \Delta}{\Delta \mid \Phi \vdash \alpha : k}$ | $\frac{\text{K-FVAR} \quad \Phi(X) \mapsto k}{\Delta \mid \Phi \vdash X : k}$ | $\frac{\text{K-APP} \quad \Delta \mid \Phi \vdash \tau_1 : k_1 \rightsquigarrow k_2 \quad \Delta \mid \Phi \vdash \tau_2 : k_1}{\Delta \mid \Phi \vdash \tau_1 \tau_2 : k_2}$ |
| $\frac{\text{K-ABS} \quad \Delta \mid \Phi, (X \mapsto k_1) \vdash \tau : k_2}{\Delta \mid \Phi \vdash \lambda X. \tau : k_1 \rightsquigarrow k_2}$ | $\frac{\text{K-FIX} \quad \Delta \mid \Phi \vdash \tau : k \rightsquigarrow k}{\Delta \mid \Phi \vdash \mu(\tau) : k}$ | $\frac{\text{K-FUN} \quad \Delta \mid \emptyset \vdash \tau_1 : * \quad \Delta \mid \Phi \vdash \tau_2 : *}{\Delta \mid \Phi \vdash \tau_1 \Rightarrow \tau_2 : *}$ |
| $\frac{\text{K-EMPTY}}{\Delta \mid \Phi \vdash : k}$ | $\frac{\text{K-UNIT}}{\Delta \mid \Phi \vdash \mathbb{1} : k}$ | $\frac{\text{K-PRODUCT} \quad \Delta \mid \Phi \vdash \tau_1 : k \quad \Delta \mid \Phi \vdash \tau_2 : k}{\Delta \mid \Phi \vdash \tau_1 \times \tau_2 : k}$ |
| $\frac{\text{K-SUM} \quad \Delta \mid \Phi \vdash \tau_1 : k \quad \Delta \mid \Phi \vdash \tau_2 : k}{\Delta \mid \Phi \vdash \tau_1 + \tau_2 : k}$ | | |
| $\frac{\text{SC-FORALL} \quad \Delta, (\alpha \mapsto k) \vdash \sigma}{\Delta \vdash \forall \alpha. \sigma}$ | $\frac{\text{SC-TYPE} \quad \Delta \mid \emptyset \vdash \tau : *}{\Delta \vdash \tau}$ | $\Delta \vdash \sigma$ |

Figure 17: Well-formedness rules for types and type schemes

a predicative subset of F_ω , we are optimistic that the Curry-style formulation of our type system will be decidable too, but verifying this expectation is future work.

5.3.1 Well-Formed Types

Types are well-formed with respect to a kind k , describing the arity of a type's parameters, if it has any. Well-formedness of types is defined using the judgment $\Delta \mid \Phi \vdash \tau : k$, stating that the type τ has kind k under contexts Δ and Φ . Similarly, well-formedness of type schemes is defined by the judgment $\Delta \vdash \sigma$, stating that the type scheme σ is well-formed with respect to the context Δ .

Following [Johann et al. \[2021\]](#), well-formedness of types is defined with respect to two contexts, one containing functorial variables (Φ), and one containing variables with mixed variance (Δ). Specifically, the variables in the context Φ are restricted to occur only in *strictly positive* [[Abbott et al., 2005b](#), [Coquand and Paulin, 1988](#)] positions (i.e., they can never appear to the left of a function arrow), while the variables in Δ can have mixed variance. This restriction on the occurrence of the variables in Φ is enforced in the well-formedness rule for function types, K-FUN , which requires that its domain is typed under an empty context of functorial variables, preventing the domain type from dereferencing any functorial variables bound in the surrounding context. While it may seem overly restrictive to require type expressions to be strictly positive—rather than merely positive—in Φ , this is necessary to ensure that μ -types, as well as its introduction and elimination forms, have a well-defined semantics (see Section [5.4.2.1](#)). Variables in Φ are bound by type-level λ -abstraction, meaning that any type former with kind $k_1 \rightsquigarrow k_2$ is functorial in its argument. In contrast, the variables in Δ are bound by \forall -quantification.

Products (\times), coproducts ($+$), units ($\mathbb{1}$) and empty types ($()$) can be constructed at any kind, reflecting the fact that the corresponding categorical (co)limits can be lifted from SET to its functor categories by computing them pointwise. This pointwise lifting of these (co)limits to functor categories is reflected in the β equalities for these type formers (shown in [Figure 20](#)), which allow an instance at kind $k_1 \rightsquigarrow k_2$, when applied with a type argument, to be replaced with an instance at kind k_2 .

The well-formed judgements for types effectively define a (simply typed) type level λ -calculus with base “type” \star . Consequently, the same type has multiple equivalent representations in the presence of β -redexes, raising the question of how we should deal with type normalization. The approach we adopt here is to add a non-syntactic conversion rule to the definition of our type system that permits any well-formed term to be

$$\begin{array}{l}
x, y \in \text{String} \\
\\
Env \ni \Gamma ::= \emptyset \mid \Gamma, x : \sigma \\
Term \ni M, N ::= x \mid M N \mid \lambda x. M \mid \mathbf{let} (x : \sigma) = M \mathbf{in} N \\
\quad \mid \Lambda \alpha. M \mid M @_{\tau} \mid \mathbf{in} \mid \mathbf{unin} \mid \mathbf{map} \langle M \rangle^{\tau} \mid (\mid M \mid)^{\tau} \\
\quad \mid \pi_1 \mid \pi_2 \mid M \blacktriangle N \mid \iota_1 \mid \iota_2 \mid M \blacktriangledown N \mid \mathbf{tt} \mid \mathbf{absurd} \\
\\
\tau_1 \xrightarrow{*} \tau_2 \triangleq \tau_1 \Rightarrow \tau_2 \quad (\text{Arrow Types}) \\
\tau_1 \xrightarrow{(k_1 \rightsquigarrow k_2)} \tau_2 \triangleq \forall \alpha. \tau_1 \alpha \xrightarrow{k_2} \tau_2 \alpha \\
\text{where } \Delta \vdash \tau_1 \xrightarrow{k} \tau_2 \text{ if } \Delta \mid \emptyset \vdash \tau_1, \tau_2 : k
\end{array}$$

Figure 18: Term syntax

typed under an equivalent type scheme. Section 5.3.3 discusses type equivalence in more detail.

5.3.2 Well-Typed Terms

Figure 18 shows the term syntax of our calculus. Along with the standard syntactic forms of the polymorphic λ -calculus we include explicit type abstraction and application, as well as introduction and elimination forms for recursive types (**in/unin**), products ($\pi_1/\pi_2/- \blacktriangle -$), coproducts ($\iota_1/\iota_2/- \blacktriangledown -$), and the unit (**tt**) and empty (**absurd**) types. Furthermore, the calculus includes dedicated primitives for mapping ($\mathbf{map} \langle - \rangle^{\tau}$) and folding ($(\mid - \mid)^{\tau}$) over a type.

Figure 18 also includes the definition of *arrow types*. In spirit of the syntactic notion of natural transformations used by Abel and Matthes [2002], Abel et al. [2003, 2005] to study generalized (Mendler) iteration, an arrow type of the form $\tau_1 \xrightarrow{k} \tau_2$ (where $\tau_1, \tau_2 : k$) defines the type of *morphisms* between the objects that interpret τ_1 and τ_2 . Arrow types are defined by induction over k , since the precise meaning of morphism for any pair of types depends on their kind. If $k = *$, then a morphism between τ_1 and τ_2 is simply a function type. However, if τ_1 and τ_2 have

one or more type argument, they are to be interpreted as objects in a suitable functor category, meaning that their morphisms are natural transformations. This is reflected in the definition of arrow types, by unfolding an arrow $\tau_1 \xrightarrow{k} \tau_2$ to a \forall -type that closes over all type arguments of τ_1 and τ_2 , capturing the intuition that polymorphic functions correspond to natural transformations.⁵⁵ For instance, we would type the inorder traversal of binary trees as $inorder : Tree \xrightarrow{* \rightsquigarrow * *} List (\triangleq \forall \alpha. Tree \alpha \Rightarrow List \alpha)$, describing a natural transformation between the *Tree* and *List* functors.

⁵⁵ This intuition is made formal by Theorem 1 in Section 5.4.4.

The typing rules are shown in shown in Figure 19. The rules rely on arrow types for introduction and elimination forms. For example, Products can be constructed at any kind (following rule K-PRODUCT in Figure 17), so the rules for terms that operate on these (i.e., T-FST, T-SND, and T-FORK) use arrow types at any kind k . Consequently, arrow types should correspond to morphisms in a suitable category, such that the semantics of a product type and its introduction/elimination forms can be expressed as morphisms in this category.

5.3.3 Type Equivalence

In the presence of type level λ -abstraction and application, the same type can have multiple representations. For this reason, the type system defined in Figure 19 includes a non-syntactic conversion rule that allows a well-typed term to be re-typed under any equivalent type scheme. The relevant equational theory for types is defined in Figure 20, and includes the customary β and η equivalences for λ -terms, as well as β rules for product, sum, unit, and empty types. The equations shown in Figure 20 are motivated by the semantic model we discuss in Section 5.4, in the sense that equivalent types are interpreted to naturally isomorphic functors. The relation is also reflexive and transitive, motivated by respectively the identity and composition of natural isomorphisms. Viewing the equalities in Figure 20 left-to-right provides us with a basis for a normalization strategy

$$\boxed{\Gamma \vdash M : \sigma}$$

$$\begin{array}{c}
\text{T-VAR} \\
\frac{x : \sigma \in \Gamma}{\Gamma \vdash x : \sigma} \\
\\
\text{T-APP} \\
\frac{\Gamma \vdash M : \tau_1 \Rightarrow \tau_2 \quad \Gamma \vdash N : \tau_1}{\Gamma \vdash MN : \tau_2} \\
\\
\text{T-ABS} \\
\frac{\Gamma, (x : \tau_1) \vdash M : \tau_2}{\Gamma \vdash \lambda x. M : \tau_1 \Rightarrow \tau_2} \\
\\
\text{T-LET} \\
\frac{\Gamma \vdash M : \sigma_1 \quad \Gamma, x : \sigma_1 \vdash N : \sigma_2}{\Gamma \vdash \mathbf{let} (x : \sigma_1) = M \mathbf{in} N : \sigma_2} \\
\\
\text{T-TYPEABS} \\
\frac{\Gamma \vdash M : \sigma \quad \alpha \notin \text{freevars}(\Gamma)}{\Gamma \vdash \Lambda \alpha. M : \forall \alpha. \sigma} \\
\\
\text{T-TYPEAPP} \\
\frac{\Gamma \vdash M : \forall \alpha. \sigma}{\Gamma \vdash M @\tau : \sigma[\tau/\alpha]} \\
\\
\text{T-IN} \\
\frac{}{\Gamma \vdash \mathbf{in} : \tau \mu(\tau) \xrightarrow{k} \mu(\tau)} \\
\\
\text{T-OUT} \\
\frac{}{\Gamma \vdash \mathbf{unin} : \mu(\tau) \xrightarrow{k} \tau \mu(\tau)} \\
\\
\text{T-MAP} \\
\frac{\Gamma \vdash M : \tau_1 \xrightarrow{k_1} \tau_2}{\Gamma \vdash \mathbf{map} \langle M \rangle^\tau : \tau \tau_1 \xrightarrow{k_2} \tau \tau_2} \\
\\
\text{T-FOLD} \\
\frac{\Gamma \vdash M : \tau_1 \tau_2 \xrightarrow{k} \tau_2}{\Gamma \vdash \langle M \rangle^{\tau_1} : \mu(\tau_1) \xrightarrow{k} \tau_2} \\
\\
\text{T-FST} \\
\frac{}{\Gamma \vdash \boldsymbol{\pi}_1 : \tau_1 \times \tau_2 \xrightarrow{k} \tau_1} \\
\\
\text{T-SND} \\
\frac{}{\Gamma \vdash \boldsymbol{\pi}_2 : \tau_1 \times \tau_2 \xrightarrow{k} \tau_2} \\
\\
\text{T-FORK} \\
\frac{\Gamma \vdash M : \tau \xrightarrow{k} \tau_1 \quad \Gamma \vdash N : \tau \xrightarrow{k} \tau_2}{\Gamma \vdash M \blacktriangle N : \tau \xrightarrow{k} \tau_1 \times \tau_2} \\
\\
\text{T-INL} \\
\frac{}{\Gamma \vdash \boldsymbol{\iota}_1 : \tau_1 \xrightarrow{k} \tau_1 + \tau_2} \\
\\
\text{T-INR} \\
\frac{}{\Gamma \vdash \boldsymbol{\iota}_2 : \tau_2 \xrightarrow{k} \tau_1 + \tau_2} \\
\\
\text{T-JOIN} \\
\frac{\Gamma \vdash M : \tau_1 \xrightarrow{k} \tau \quad \Gamma \vdash M : \tau_2 \xrightarrow{k} \tau}{\Gamma \vdash M \blacktriangledown N : \tau_1 + \tau_2 \xrightarrow{k} \tau} \\
\\
\text{T-UNIT} \\
\frac{}{\Gamma \vdash \mathbf{tt} : \mathbb{1}} \\
\\
\text{T-EMPTY} \\
\frac{}{\Gamma \vdash \mathbf{absurd} : \xrightarrow{k} \tau} \\
\\
\text{T-CONV} \\
\frac{\Gamma \vdash M : \sigma_1 \quad \sigma_1 \equiv \sigma_2}{\Gamma \vdash M : \sigma_2}
\end{array}$$

Figure 19: Well-formed terms

$$\begin{array}{l}
(\lambda X. \tau_1) \tau_2 \equiv \tau_1[\tau_2/X] \\
(\lambda X. \tau X) \equiv \tau \quad T := \Box \mid T \tau \mid \tau T \mid \mu(T) \mid T \Rightarrow \tau \mid \tau \Rightarrow T \\
(\tau_1 \times \tau_2) \tau \equiv (\tau_1 \tau) \times (\tau_2 \tau) \mid T \times \tau \mid \tau \times T \mid T + \tau \mid \tau + T \\
(\tau_1 + \tau_2) \tau \equiv (\tau_1 \tau) + (\tau_2 \tau) \\
\mathbb{1} \tau \equiv \mathbb{1} \\
\tau \equiv \tau
\end{array}
\quad \frac{\tau_1 \equiv \tau_2}{T[\tau_1] \equiv T[\tau_2]}$$

for types, which would be required for implementing the type system.

Figure 20: Equational theory for types

5.4 CATEGORICAL SEMANTICS

In this section, we consider how to define a categorical semantics for our calculus, drawing inspiration from the semantics defined by [Johann and Polonsky \[2019\]](#) and [Johann et al. \[2021\]](#), [Johann and Ghiorzi \[2021\]](#). To define this semantics, we must show that each type in our calculus corresponds to a functor, and that all such functors have initial algebras. In [Section 5.4.3](#) we discuss the requirements for these initial algebras to exist, and argue informally why they should exist for the functors interpreting our types. Although [Johann and Polonsky \[2019\]](#) present a detailed argument for the existence of initial algebras of the functors underlying nested data types, it is still future work to adapt this argument to our setting.

The general setup of our semantics is to interpret types of kind \star as objects in SET (the category of sets), higher-order types as functors on SET , and type schemes as objects in SET_2 (the category of large sets). This size bump is necessary to model the universal quantification over types in type schemes. Crucially, SET is a *full subcategory* of SET_1 , as witnessed by the existence of a fully faithful inclusion functor I :

$$\text{SET} \xrightarrow{I} \text{SET}_1$$

Assuming cumulative universes (i.e., the collection of all large sets also includes all small sets), I is just the identity functor. We remark that both SET and SET_1 are *complete* and *cocomplete* and *cartesian closed*. Importantly, since I is fully faithful, the cartesian closed structure of SET is reflected in SET_1 for those objects that lie in the image of I .

The subcategory relation between SET and SET_1 reflects the syntactic restriction of types to rank-1 polymorphism: all objects in SET can also be found in SET_1 , but SET_1 is sufficiently larger than SET that it also includes objects modelling quantification over objects in SET . This intuition is embodied by fact that every functor $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \text{SET}_1$, where \mathcal{C} is smaller than SET_1 (which includes SET), has an end in SET_1 . This follows from completeness of SET_1 [MacLane, 1971, p. 224, corollary 2]. We discuss the use of ends for modelling universal quantification in more detail in Section 5.4.2.3.

5.4.1 Interpreting Kinds and Kind Environments

We associate with each kind k a category whose objects interpret the types of that kind. The semantics of kinds is defined by induction over k , where we map the base kind \star to SET , and kinds of the form $k_1 \rightsquigarrow k_2$ to the category of functors between their domain and codomain.⁵⁶

⁵⁶ Here, \mathbf{CAT} denotes the (very large) category of large categories. Although SET itself is locally small, its functor categories have a large set of morphisms.

$$\begin{aligned} \llbracket - \rrbracket & : \text{Kind} \rightarrow \mathbf{CAT} \\ \llbracket \star \rrbracket & = \text{SET} \\ \llbracket k_1 \rightsquigarrow k_2 \rrbracket & = [\llbracket k_1 \rrbracket , \llbracket k_2 \rrbracket] \end{aligned}$$

By interpreting types of kind $k_1 \rightsquigarrow k_2$ as objects in a functor category, we formalize the intuition that higher-order types correspond to functors. The semantics of kind contexts is then defined on a per-entry basis, as a chain of products of the categories that interpret their elements.

$$\begin{aligned}
\llbracket - \rrbracket & : \text{Context} \rightarrow \mathbf{CAT} \\
\llbracket \emptyset \rrbracket & = \bullet \\
\llbracket \Delta, \alpha \mapsto k \rrbracket & = \llbracket \Delta \rrbracket \times \llbracket k \rrbracket
\end{aligned}$$

Here, \bullet denotes the *trivial category*, which has a single object, $*$, together with its identity morphism, id_* . It is worth mentioning that \bullet and $- \times -$, together with the operation of constructing a functor category, $[-, -]$, imply that \mathbf{CAT} is a cartesian closed category. We will use this cartesian closed structure to give a semantics to the fragment of well-formed types that corresponds to the simply-typed λ -calculus.

5.4.2 Interpreting Types

Since a well-formed type $\Delta \mid \Phi \vdash \tau : k$ is intended to be functorial in all variables in Φ , it is clear that its semantics should be a functor over the category associated with Φ (i.e., $\llbracket \Phi \rrbracket$). But what about the variables in Δ , which can occur both in covariant and contravariant positions? For example, in the type of the identity function, $\forall \alpha. \alpha \Rightarrow \alpha$, we cannot interpret the sub-derivation for $\alpha \Rightarrow \alpha$ as a functor over the category interpreting its free variables since there would not be a sensible way to define its action on morphisms due to the negative occurrence of α . To account for the mixed variance of universally quantified type variables, we instead adopt a *difunctorial semantics*, interpreting types as a functor on the product category $\llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket$ (similar representations of type expressions with mixed variance appear, for example, when considering Mendler-style inductive types [Uustalu and Vene, 1999], or the object calculus semantics by Glimming and Ghani [2004]). Well-formed types (left) and type schemes (right) are interpreted as a functors over their contexts of the following form:

$$\begin{aligned}
\llbracket \Delta \mid \Phi \vdash \tau : k \rrbracket & : (\llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket) \times \llbracket \Phi \rrbracket \rightarrow \llbracket k \rrbracket \\
\llbracket \Delta \vdash \sigma \rrbracket & : \llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket \rightarrow \mathbf{SET}_1
\end{aligned}$$

$$\begin{aligned}
\llbracket \Delta \mid \Phi \vdash \alpha : \tau \rrbracket &= \mathbf{lookup}_\alpha^\Delta \circ \pi_2 \circ \pi_1 \\
\llbracket \Delta \mid \Phi \vdash X : \tau \rrbracket &= \mathbf{lookup}_X^\Phi \circ \pi_2 \\
\llbracket \Delta \mid \Phi \vdash \tau_1 \tau_2 : k_2 \rrbracket &= \mathbf{eval} \circ \langle \llbracket \Delta \mid \Phi \vdash \tau_1 : k_1 \rightsquigarrow k_2 \rrbracket, \llbracket \Delta \mid \Phi \vdash \tau_2 : k_1 \rrbracket \rangle \\
\llbracket \Delta \mid \Phi \vdash \lambda X. \tau : k_1 \rightsquigarrow k_2 \rrbracket &= \mathbf{curry}(\llbracket \Delta \mid \Phi, X : k_1 \vdash \tau : k_2 \rrbracket) \\
\llbracket \Delta \mid \Phi \vdash \mu(\tau) : k \rrbracket &= \boldsymbol{\mu}(\llbracket \Delta \mid \Phi \vdash \tau : k \rightsquigarrow k \rrbracket) \\
\llbracket \Delta \mid \Phi \vdash \tau_1 \Rightarrow \tau_2 : \star \rrbracket &= \mathbf{exp}(\llbracket \Delta \mid \emptyset \vdash \tau_1 : \star \rrbracket, \llbracket \Delta \mid \Phi \vdash \tau_2 : \star \rrbracket) \\
\llbracket \Delta \mid \Phi \vdash : \star \rrbracket &= \perp \\
\llbracket \Delta \mid \Phi \vdash \mathbb{1} : \star \rrbracket &= \top \\
\llbracket \Delta \mid \Phi \vdash \tau_1 \times \tau_2 : k \rrbracket &= \llbracket \Delta \mid \Phi \vdash \tau_1 : k \rrbracket \times \llbracket \Delta \mid \Phi \vdash \tau_2 : k \rrbracket \\
\llbracket \Delta \mid \Phi \vdash \tau_1 + \tau_2 : k \rrbracket &= \llbracket \Delta \mid \Phi \vdash \tau_1 : k \rrbracket + \llbracket \Delta \mid \Phi \vdash \tau_2 : k \rrbracket \\
\llbracket \Delta \vdash \forall \alpha. \sigma \rrbracket &= \mathbf{end}(\mathbf{curry}(\llbracket \Delta, \alpha : k \vdash \sigma \rrbracket \circ \mathbf{sift})) \\
\llbracket \Delta \vdash \tau \rrbracket &= \mathbf{I} \circ \llbracket \Delta \mid \emptyset \vdash \tau : \star \rrbracket
\end{aligned}$$

Figure 21: Semantics of well-formed types and type schemes

Ultimately, the goal of this setup is to interpret \forall -types as *ends* in \mathbf{SET}_1 , which allows us to formally argue that terms that are well-formed with an arrow type of the form $\tau_1 \xrightarrow{k} \tau_2$ (which unfolds to $\forall \bar{\alpha}. \tau_1 \bar{\alpha} \Rightarrow \tau_2 \bar{\alpha}$) correspond, in a suitable sense, to the natural transformations between the functors interpreting τ_1 and τ_2 . Or, put differently, terms with an arrow type define a morphism between the interpretation of their domain and codomain. We discuss the semantics of universal quantification further in Section 5.4.2.3, and give a more precise account of the relation between arrow types and natural transformations in Section 5.4.4.

Figure 21 defines the semantics of well-formed types and type schemes. The interpretation of the empty type, unit type, and (co)product types follow immediately from (co)completeness of \mathbf{SET} . Since they can be constructed at any kind, the semantics of (co)product types depends crucially on the fact that functor categories preserve all (co)limits of their codomain category,

which implies that $\llbracket k \rrbracket$ is (co)complete for any k . To interpret variables, we utilize the cartesian closed structure of **CAT** to compute an appropriate projection based on the position of the variable in the environment.

$$\begin{aligned} \text{lookup}_\alpha^\Delta & : \llbracket \Delta \rrbracket \rightarrow \llbracket k \rrbracket \\ \text{lookup}_\alpha^{\Delta, \alpha:k} & \mapsto \pi_2 \\ \text{lookup}_\alpha^{\Delta, \beta:k} & \mapsto \text{lookup}_\alpha^\Delta \circ \pi_1 \quad (\text{where } \alpha \neq \beta) \end{aligned}$$

Similarly, the cartesian closed structure of **CAT** also implies the existence of functors $\text{eval} : [\mathcal{C}, \mathcal{D}] \times \mathcal{C} \rightarrow \mathcal{D}$ and $\text{curry}(F) : \mathcal{C} \rightarrow [\mathcal{D}, \mathcal{E}]$, for any $F : \mathcal{C} \times \mathcal{D} \rightarrow \mathcal{E}$, which immediately provide a semantics for type-level application and abstraction respectively. The remaining type and type scheme constructors are interpreted using specifically-defined functors. Although their definitions are typical examples of how (co)limits are lifted to functor categories by computing them pointwise, we discuss the definition of these functors separately and in more detail respectively in Section 5.4.2.1 (recursive types), Section 5.4.2.2 (function types), and Section 5.4.2.3 (\forall -types).

5.4.2.1 Recursive Types

Following the usual categorical interpretation of inductive data types [Goguen, 1976], the semantics of recursive types is given by an *initial algebras*. We summarize the setup here. An *F-algebra* for an endofunctor $F : \mathcal{C} \rightarrow \mathcal{C}$ is defined as a tuple (A, α) of an object $A \in \mathcal{C}$ (called the *carrier*), and a morphism $\alpha : FA \rightarrow A$. An *algebra homomorphism* between *F-algebras* (A, α) and (B, β) is given by a morphism $f : A \rightarrow B$ such that the following diagram commutes.

$$\begin{array}{ccc} FA & \xrightarrow{\alpha} & A \\ Ff \downarrow & & \downarrow f \\ FB & \xrightarrow{\beta} & B \end{array}$$

F-algebras and their homomorphisms form a category. If F is an endofunctor, we denote the initial object of the category of

F-algebras (which, if it exists, we refer to as the initial algebra) as $(\mu F, \text{in})$. Initial algebras give a semantics to inductive data types, with their universal property providing an induction principle. Given an F-algebra (A, α) , we denote unique F-algebra homomorphism that factors through A by $\text{cata}(\alpha) : \mu F \rightarrow A$. Instantiating the diagram above with $\text{cata}(\alpha)$ gives us the familiar universal property of folds, $\text{cata}(\alpha) \circ \text{in} = \alpha \circ F(\text{cata}(\alpha))$, which defines their computational behavior.

To interpret recursive types in our calculus, we construct the functor $\mu(F)$, which sends objects pointwise to the initial algebras of a functor $F : \mathcal{C} \rightarrow [\mathcal{D}, \mathcal{D}]$. For a morphism $f : X \rightarrow Y$, the action of $\mu(F)$ on f is defined by factoring through the algebra defined by precomposing the initial algebra of $F(Y)$ with the action of F on f , which defines a natural transformation $F(X) \rightarrow F(Y)$, at component $\mu(F(Y))$.

$$\begin{aligned} \mu(F)(-) & : \mathcal{C} \rightarrow \mathcal{D} \\ \mu(F)(x) & \mapsto \mu(F(x)) \\ \mu(F)(f) & \mapsto \text{cata}(\text{in} \circ F(f)_{\mu(F(Y))}) \end{aligned}$$

In general, it is not guaranteed that an initial algebra exists for any endofunctor $F : \mathcal{C} \rightarrow \mathcal{C}$. Typically, the existence of an initial algebras is shown by iterating F and showing that it converges, applying the classic theorem by [Adámek \[1974\]](#). This approach imposes some additional requirements on the functor F and underlying category \mathcal{C} , which we discuss in more detail in Section 5.4.3.

5.4.2.2 Function Types

The functor $\text{exp}(-)$ is defined by mapping onto exponential objects in SET . But we have to take some additional care to ensure that we can still define its action on morphism, as the polarity of free variables is reversed in domain of a function type. Indeed, when computed pointwise, exponential objects give rise to a bifunctor of the form $\mathcal{C}^{\text{op}} \rightarrow \mathcal{C} \rightarrow \mathcal{C}$, meaning that functors are not, in general, closed under exponentiation. To some extent we anticipated this situation already in the

design of our type system by defining the well-formedness rule for function types such that the context of functorial variables, Φ , is discarded in its domain. Of course, the variables in Δ can occur both in covariant and contravariant positions, but by adopting a difunctorial semantics we limit ourselves to a specific class of functors that is closed under exponentiation. The key observation is that constructing the opposite category of the product of a category and its opposite is an idempotent (up to isomorphism) operation. That is, we have the following equivalence of categories: $(\mathcal{C}^{\text{op}} \times \mathcal{C})^{\text{op}} \simeq \mathcal{C}^{\text{op}} \times \mathcal{C}$. As a result, a pointwise mapping of difunctors to exponential objects does give rise to a new difunctor. We use this fact to our advantage to define the following functor $\mathbf{exp}(F, G)$ for functors $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{E}$ and $G : (\mathcal{C}^{\text{op}} \times \mathcal{C}) \times \mathcal{D} \rightarrow \mathcal{E}$, of which the interpretation of function types is an instance.

$$\begin{aligned} \mathbf{exp}(F, G)(-) & : (\mathcal{C}^{\text{op}} \times \mathcal{C}) \times \mathcal{D} \rightarrow \mathcal{E} \\ \mathbf{exp}(F, G)((x, y), z) & \mapsto G((x, y), z)^{F(y, x)} \\ \mathbf{exp}(F, G)((f, g), h) & \\ & \mapsto \text{curry}(G((f, g), h) \circ \text{eval} \circ (\text{id}_{\mathbf{exp}(F, G)((x, y), z)} \times F(g, f))) \end{aligned}$$

We remark that $\mathbf{exp}(F, G)$ does not define an exponential object in the functor category $[(\mathcal{C}^{\text{op}} \times \mathcal{C}) \times \mathcal{D}, \mathcal{E}]$. Fortunately, for defining the semantics of term level λ -abstraction or application it is sufficient that the action on objects maps to exponentials in SET .

5.4.2.3 Universal quantification

The semantics of universal quantifications is expressed in terms of ends in the category SET_1 . If $F : \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{D}$ is a functor, then an *end* of F is an object $\int_{x \in \mathcal{C}} F(x, x) \in \mathcal{D}$ equipped with a projection map given by an extranatural transformation $\pi_x : \int_{c \in \mathcal{C}} F(c, c) \rightarrow F(x, x)$. Formally, an end of the F is defined as the universal wedge of the following diagram:

$$F(x, x) \xrightarrow{F(\text{id}_x, f)} F(x, y) \xleftarrow{F(f, \text{id}_y)} F(y, y)$$

For all $x, y \in \mathcal{C}$ and $f : x \rightarrow y$. The universal property of ends then states that any other wedge $W \in \mathcal{D}$ with maps $i : W \rightarrow F(x, x)$ and $j : W \rightarrow F(y, y)$ uniquely factors through $\int_{c \in \mathcal{C}} F(c, c)$.

$$\begin{array}{ccccc}
 W & \xrightarrow{i} & F(x, x) & & \\
 \downarrow & \searrow j & \nearrow & \searrow F(\text{id}_x, f) & \\
 \text{factor}(W) & & & & \\
 \downarrow \sphericalangle & \nearrow \pi_x & & & \\
 \int_{c \in \mathcal{C}} F(c, c) & \xrightarrow{\pi_y} & F(y, y) & \xrightarrow{F(f, \text{id}_y)} & F(x, y)
 \end{array}$$

To model the more general situation where a \forall -quantified type can contain free variables that are bound by another quantifier above it in the lexical hierarchy, we define the semantics of universal quantification in terms of the *end functor*, $\mathbf{end}(-)$, which for a functor $G : \mathcal{C} \rightarrow [\mathcal{D}^{\text{op}} \times \mathcal{D}, \mathcal{E}]$ defines a functor $\mathbf{end}(G) : \mathcal{C} \rightarrow \mathcal{E}$ whose object action is computed pointwise from ends in \mathcal{E} . Its action on morphisms, $\mathbf{end}(f) : \int_{d \in \mathcal{D}} G(X)(d, d) \rightarrow \int_{d \in \mathcal{D}} G(Y)(d, d)$, follows from the universal property stated above. To define the action on morphisms, we observe that the object $\int_{d \in \mathcal{D}} G(X)(d, d)$ is a wedge of the following diagram.

$$\int_{d \in \mathcal{D}} G(X)(d, d) \xrightarrow{G(Y)(\text{id}_x, f)} \int_{d \in \mathcal{D}} G(Y)(d, d) \xleftarrow{G(Y)(f, \text{id}_y)} \int_{d \in \mathcal{D}} G(Y)(d, d)$$

Where the vertices of the cone are constructed by composing the projection map with the action of G on f , i.e., $G(f)(x, x) \circ \pi_x$. By universality, this wedge uniquely factors through the end $\int_{d \in \mathcal{D}} G(Y)(d, d)$. This factorization defines the morphism action $\mathbf{end}(f)$.

$$\begin{aligned}
 \mathbf{end}(G)(-) & : \mathcal{C} \rightarrow \mathcal{E} \\
 \mathbf{end}(G)(x) & \mapsto \int_{d \in \mathcal{D}} G(x)(d, d) \\
 \mathbf{end}(G)(f) & \mapsto \text{factor}\left(\int_{d \in \mathcal{D}} G(x)(d, d)\right)
 \end{aligned}$$

An important subtlety here is that $F(X)$ should have an end in \mathcal{E} for every X . In our case, this is a consequence of completeness of \mathbf{SET}_1 .⁵⁷ To actually use the functor \mathbf{end} to define the semantics of universal quantifications, we need to precompose the semantics

⁵⁷ See Mac Lane [MacLane, 1971] chapter 9.5 corollary 2.

of its body with the **sift** functor to separate the quantified variable from the remainder of the context.

$$\begin{aligned} \mathbf{sift} &: ([\Delta] \times [k])^{\text{op}} \times ([\Delta] \times [k]) \times [\Phi] \\ &\rightarrow (([\Delta]^{\text{op}} \times [\Delta]) \times [\Phi]) \times ([k]^{\text{op}} \times [k]) \end{aligned}$$

We note that **sift** defines an isomorphism in **CAT**.

5.4.3 On the Existence of Initial Algebras

In general, it is not the case that any endofunctor has an initial algebra. For certain classes of endofunctors, it can be shown that an initial algebra exists by means of Adámek's theorem [Adámek, 1974]. Here, we present a condensed argument for why we expect that functors interpreting well-formed types of kind $k \rightsquigarrow k$ (for any k) have initial algebras; a more thorough formal treatment of the construction of initial algebras is a subject of further study.

The intuition behind Adámek's construction is that repeated applications of an endofunctor $F : \mathcal{C} \rightarrow \mathcal{C}$ converge after infinite iterations, reaching a fixpoint. If \mathcal{C} has an initial object and ω -colimits,⁵⁸ we can define the initial algebra of F as the ω -colimit of the following chain:

$$\perp \xrightarrow{!} F\perp \xrightarrow{F!} FF\perp \xrightarrow{FF!} FFF\perp \xrightarrow{FFF!} \dots$$

Where \perp is the initial object in \mathcal{C} and $!_X : \perp \rightarrow X$ the unique map from \perp to X . A crucial stipulation is that F should be ω -cocontinuous, meaning that it preserves ω -colimits.

Thus, for the functors interpreting higher-order types to have an initial algebra, we must argue that all higher-order types are interpreted to a ω -cocontinuous functor. This prompts a refinement of the semantics for kinds discussed in Section 5.4.1, where we impose the additional restriction that the interpretation of a kind of the form $k_1 \rightsquigarrow k_2$ is a ω -cocontinuous functor from $[[k_1]]$ to $[[k_2]]$. Subsequently, we must show that Figure 21 actually inhabits this refined semantics.

⁵⁸ That is, colimits over diagrams defined as a functor on the thin category generated from the poset of natural numbers.

Johann and Polonsky [2019] present an inductive argument showing the existence of initial algebras for a universe of higher-kinded data types is similar to our definition of well-formed terms in Figure 17. While their proof establishes the more general property of λ -cocontinuity (for an arbitrary limit ordinal λ) for the functors interpreting higher-kinded types, we expect that the relevant cases of their inductive proof—specifically the cases for products, coproducts, type application, and the μ functor—can be adapted to our setting. What remains is to show that the semantics of type level λ -abstraction and function types is a ω -cocontinuous functor. For λ -abstraction, we transport along the currying isomorphism, which should preserve ω -cocontinuity. For function types, we require that the functor $(-)^X : \text{SET} \rightarrow \text{SET}$ is ω -cocontinuous for all X , which, as Johann and Polonsky [2019] point out, is indeed the case. Expanding this proof sketch into a full proof of the existence of initial algebras is future work.

5.4.4 Arrow Types Correspond to Morphisms

To define the semantics of well-typed terms, it is crucial that we can relate arrow types—i.e., of the form $\tau_1 \xrightarrow{k} \tau_2$ —to morphisms in the category $\llbracket k \rrbracket$. To make this more precise, consider the typing rule for left projections. To define its semantics, we would like to use the cartesian structure of the category $\llbracket k \rrbracket$, which implies the existence of a *morphism* $\pi_1 : \llbracket k \rrbracket(x \times y, x)$ for $x, y \in \llbracket k \rrbracket$. However, the rule T-FST implies that π_1 should be related to an *object* in SET_1 , i.e., $\llbracket \tau_1 \times \tau_2 \xrightarrow{k} \tau_1 \rrbracket$. To mediate between morphisms in $\llbracket k \rrbracket$ and objects in SET_1 calls for a suitable currying/uncurrying isomorphism for arrow types, though we highlight that the required isomorphism is different from the usual currying isomorphism arising from the existence of right adjoints for the tensor product in closed monoidal categories, in the sense that $\llbracket \tau_1 \xrightarrow{k} \tau_2 \rrbracket$ does not define an internal hom for the objects $\llbracket \tau_1 \rrbracket, \llbracket \tau_2 \rrbracket$ but rather internalizes the morphisms between these objects in a *different* category.

Theorem 1. *Given a kind \mathbb{k} , morphisms of the category $\llbracket \mathbb{k} \rrbracket$ are internalized as objects in SET_1 through the following bijection between hom-sets:*

$$\llbracket \mathbb{k} \rrbracket(F(\delta) \times \llbracket \tau_1 \rrbracket(\delta^\circ), \llbracket \tau_2 \rrbracket(\delta)) \simeq \text{SET}_1(F(\delta), \llbracket \tau_1 \xrightarrow{\mathbb{k}} \tau_2 \rrbracket(\delta)) \quad (2)$$

Where $\delta \in \llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket$ and $\delta^\circ \in (\llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket)^{\text{op}}$ its complement, which is defined by swapping the objects representing contravariant respectively covariant occurrences of the variables in Δ . Let $F : \llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket \rightarrow \text{SET}_1$ be a functor. In a slight abuse of notation, we also write $F(\delta)$ for the “lifting” of F to an object in the (functor) category $\llbracket \mathbb{k} \rrbracket$ that ignores all the additional variables on which $\llbracket \tau_1 \rrbracket$ and $\llbracket \tau_2 \rrbracket$ depend.

$$\begin{aligned} & \llbracket \mathbb{k} \rrbracket(F(\delta) \times \llbracket \tau_1 \rrbracket(\delta^\circ), \llbracket \tau_2 \rrbracket(\delta)) \\ = & \int_{x_1 \in \llbracket k_1 \rrbracket} \cdots \int_{x_n \in \llbracket k_n \rrbracket} \text{SET}_1(F(\delta) \times \llbracket \tau_1 \rrbracket(\delta^\circ)(x_1) \cdots (x_n), \llbracket \tau_2 \rrbracket(\delta)(x_1) \cdots (x_n)) \\ \simeq & \int_{x_1 \in \llbracket k_1 \rrbracket} \cdots \int_{x_n \in \llbracket k_n \rrbracket} \text{SET}_1(F(\delta), \llbracket \tau_2 \rrbracket(\delta)(x_1) \cdots (x_n))^{\llbracket \tau_1 \rrbracket(\delta^\circ)(x_1) \cdots (x_n)} \\ \simeq & \text{SET}_1(F(\delta), \int_{x_1} \cdots \int_{x_n} \llbracket \tau_2 \rrbracket(\delta)(x_1) \cdots (x_n \in \llbracket k_n \rrbracket))^{\llbracket \tau_1 \rrbracket(\delta^\circ)(x_1) \cdots (x_n)} \\ \simeq & \text{SET}_1(F(\delta), \llbracket \tau_1 \xrightarrow{\mathbb{k}} \tau_2 \rrbracket(\delta)) \end{aligned}$$

Proof. Figure 22 shows how the isomorphism in Equation (2) is computed. The first step of the derivation rewrites the left-hand side of the isomorphism to a sequence of zero or more ends in the category of very large sets, allowing us to apply currying for exponentials in SET_1 in the subsequent step. This is justified by cartesian closedness of SET , because the objects $\llbracket \tau_1 \rrbracket(\delta^\circ)(x_1) \cdots (x_n)$ and $\llbracket \tau_2 \rrbracket(\delta)(x_1) \cdots (x_n)$ are included in the image of the fully faithful inclusion functor I . Next, we use the fact that the covariant hom-functor $\text{SET}_1(x, -)$ is continuous and thus preserves ends:⁵⁹

$$\int_{y \in e} \text{SET}_1(x, G(y, y)) \simeq \text{SET}_1(x, \int_{y \in e} G(y, y)) \quad (3)$$

Figure 22: Derivation of the isomorphism in Equation (2), where $\mathbb{k} = k_1 \rightsquigarrow \cdots \rightsquigarrow k_n \rightsquigarrow \star$

⁵⁹ See [MacLane \[1971\]](#), page 225 Equation 4.

By repeatedly applying the identity above, we can distribute the aforementioned sequence of ends over the functor $\text{SET}_1(F(\delta), -)$. Intuitively, this corresponds to distributing universal quantification over logical implication in the scenario that the quantified variable does not occur freely in the antecedent, which is axiomatized in some flavors of first-order logic, though we apply a much more general instance of the same principle here. The final step then follows from the standard definition of η -equivalence implied by cartesian closedness of **CAT**. \square

We write $\uparrow(-)/\downarrow(-)$ for the functions that transport along the isomorphism defined in Equation (2).

5.4.5 Interpreting Terms

Well-typed terms, of the form $\Gamma \vdash M : \sigma$, are interpreted as natural transformations from the interpretation their context, $\llbracket \Gamma \rrbracket$, to the interpretation of their type, $\llbracket \sigma \rrbracket$. At component $\delta \in \llbracket \Delta \rrbracket^{\text{op}} \times \llbracket \Delta \rrbracket$ this transformation is given by a function with the following type:

$$\llbracket \Gamma \vdash M : \sigma \rrbracket_{\delta} : \llbracket \Gamma \rrbracket(\delta) \rightarrow \llbracket \sigma \rrbracket(\delta)$$

Here, $\llbracket \Gamma \rrbracket$ is defined componentwise by mapping contexts to a left-associated product of its elements, analogous to how we defined the interpretation of kind contexts in Section 5.4.1. Figure 23 shows the interpretation of well-typed terms in its entirety.

The interpretation of λ -abstraction and application is defined in terms of the cartesian closed structure of **SET**, which is preserved by its inclusion in SET_1 . For a type abstractions of the form $\Lambda\alpha.M$, its semantics follows from the fact that hom-functors preserves ends (see Equation (3)), which implies a bijection between the set of morphisms that interprets the type abstraction and the set of morphisms into which we interpret its body. We remark that this only works because α does not occur free in Γ , meaning that we know that $\llbracket \Gamma \rrbracket$ does not depend on

$$\begin{aligned}
\llbracket \Gamma \vdash x : \sigma \rrbracket_{\delta} &= \mathbf{lookup}_x^{\Gamma} \\
\llbracket \Gamma \vdash M N : \tau_2 \rrbracket_{\delta} &= \mathbf{eval} \circ \langle \llbracket \Gamma \vdash M : \tau_1 \Rightarrow \tau_2 \rrbracket_{\delta}, \llbracket \Gamma \vdash N : \tau_1 \rrbracket_{\delta} \rangle \\
\llbracket \Gamma \vdash \lambda x. M : \tau_1 \Rightarrow \tau_2 \rrbracket_{\delta} &= \mathbf{curry}(\llbracket \Gamma, x : \tau_1 \vdash M : \tau_2 \rrbracket_{\delta}) \\
\llbracket \Gamma \vdash \mathbf{let} (x : \sigma_1) = M \mathbf{in} N : \sigma_2 \rrbracket_{\delta} &= \mathbf{eval} \circ \langle \mathbf{curry}(\llbracket \Gamma, x : \sigma_1 \vdash N : \sigma_2 \rrbracket_{\delta}), \llbracket \Gamma \vdash M : \sigma_1 \rrbracket_{\delta} \rangle \\
\llbracket \Gamma \vdash \Lambda \alpha. M : \forall \alpha. \sigma \rrbracket_{\delta} &= \llbracket \Gamma \vdash M : \sigma \rrbracket_{\delta} \quad (\text{isomorphic per Equation (3)}) \\
\llbracket \Gamma \vdash M @ \tau : \sigma[\tau/\alpha] \rrbracket_{\delta} &= \pi_{[\tau]} \circ \llbracket \Gamma \vdash M : \forall \alpha. \sigma \rrbracket_{\delta} \\
\llbracket \Gamma \vdash \mathbf{in} : \tau \mu(\tau) \xrightarrow{k} \mu(\tau) \rrbracket_{\delta} &= \uparrow(\mathbf{in} \circ \pi_2) \\
\llbracket \Gamma \vdash \mathbf{unin} : \mu(\tau) \xrightarrow{k} \tau \mu(\tau) \rrbracket_{\delta} &= \uparrow(\mathbf{unin} \circ \pi_2) \\
\llbracket \Gamma \vdash \mathbf{map} \langle M \rangle^{\tau} : \tau \tau_1 \xrightarrow{k_2} \tau \tau_2 \rrbracket_{\delta} &= \uparrow(\lambda(\gamma, x). \llbracket \tau \rrbracket(\delta)(\lambda y. \downarrow(\llbracket \Gamma \vdash M : \tau_1 \xrightarrow{k_1} \tau_2 \rrbracket_{\delta})(\gamma, y))) \\
\llbracket \Gamma \vdash (M) \tau_1 : \mu(\tau_1) \xrightarrow{k} \tau_2 \rrbracket_{\delta} &= \\
&\quad \uparrow(\lambda(\gamma, x). \mathbf{cata}(\lambda y. \downarrow(\llbracket \Gamma \vdash M : \tau_1 \tau_2 \xrightarrow{k} \tau_2 \rrbracket_{\delta})(\gamma, y))) \\
\llbracket \Gamma \vdash \pi_1 : \tau_1 \times \tau_2 \xrightarrow{k} \tau_1 \rrbracket_{\delta} &= \uparrow(\pi_1 \circ \pi_2) \\
\llbracket \Gamma \vdash \pi_2 : \tau_1 \times \tau_2 \xrightarrow{k} \tau_2 \rrbracket_{\delta} &= \uparrow(\pi_2 \circ \pi_2) \\
\llbracket \Gamma \vdash M \blacktriangle N : \tau \xrightarrow{k} \tau_1 \times \tau_2 \rrbracket_{\delta} &= \\
&\quad \uparrow(\langle \downarrow(\llbracket \Gamma \vdash M : \tau \xrightarrow{k} \tau_1 \rrbracket_{\delta}), \downarrow(\llbracket \Gamma \vdash N : \tau \xrightarrow{k} \tau_2 \rrbracket_{\delta}) \rangle) \\
\llbracket \Gamma \vdash \iota_1 : \tau_1 \xrightarrow{k} \tau_1 + \tau_2 \rrbracket_{\delta} &= \uparrow(\iota_1 \circ \pi_2) \\
\llbracket \Gamma \vdash \iota_2 : \tau_2 \xrightarrow{k} \tau_1 + \tau_2 \rrbracket_{\delta} &= \uparrow(\iota_2 \circ \pi_2) \\
\llbracket \Gamma \vdash M \blacktriangledown N : \tau_1 + \tau_2 \xrightarrow{k} \tau \rrbracket_{\delta} &= \\
&\quad \uparrow(\langle \downarrow(\llbracket \Gamma \vdash M : \tau_1 \xrightarrow{k} \tau \rrbracket_{\delta}), \downarrow(\llbracket \Gamma \vdash N : \tau_2 \xrightarrow{k} \tau \rrbracket_{\delta}) \rangle) \\
\llbracket \Gamma \vdash \mathbf{tt} : \mathbb{1} \rrbracket_{\delta} &= ! \quad (\text{the unique morphism to the terminal object}) \\
\llbracket \Gamma \vdash \mathbf{absurd} : \Rightarrow \tau \rrbracket_{\delta} &= \mathbf{curry}(\mathbf{h} \circ \pi_2)
\end{aligned}$$

Figure 23: Semantics of Well-Typed Terms

α in $\llbracket \Gamma \vdash M : \sigma \rrbracket_{\delta, (\alpha, \alpha)} : \llbracket \Gamma \rrbracket (\delta, (\alpha, \alpha)) \rightarrow \llbracket \sigma \rrbracket (\delta, (\alpha, \alpha))$, and thus we can view $\llbracket \Gamma \rrbracket$ as a constant when applying the isomorphism. The semantics of a type application $M @ \tau$ is then given by the projection map at component $\llbracket \tau \rrbracket$ of the end interpreting the type of M . For the introduction and elimination forms of (co)product types, and the unit and empty type, we define the semantics in terms of the corresponding (co)limits in SET_1 , applying the currying isomorphism defined in Equation (2) to mediate with arrow types. Similarly, a semantics for the mapping and folding primitives also follows from the currying isomorphism defined in Equation (2).

Both the denotation function $\llbracket - \rrbracket$ as well as the function it computes are total. Consequently, a well-typed value can be computed from every well-typed term. In this sense, the categorical model provides us with a sound computational model of the calculus, which we could implement by writing a definitional interpreter [Reynolds, 1998]. In the next section, we will discuss how a more traditional small-step operational semantics can be derived from the same categorical model.

5.5 OPERATIONAL SEMANTICS

The previous section gave an overview of a categorical semantics of our calculus. In this section, we define a small-step operational semantics for our calculus, and discuss how it relates to the categorical model.

5.5.1 Reduction Rules

We define our operational semantics as a reduction semantics in the style of Felleisen and Hieb [1992]. Figure 24 shows the definition of values and evaluation contexts. In our definition of values, we must account for the fact that language primitives can exist at any kind. For example, the primitive \mathbf{v}_1 by itself is a value of type $\tau_1 \xrightarrow{k} \tau_1 + \tau_2$. Simultaneously, applying \mathbf{v}_1 with a value and/or a sequence of type arguments (the number of

$$\begin{aligned}
v & := \lambda x.M \mid \Lambda \alpha.M \mid \mathbf{in} \ \bar{\tau} \ v \mid \mathbf{unin} \ \bar{\tau} \ v \mid (v_1 \ \blacktriangle \ v_2) \ \bar{\tau} \ v & \text{(Values)} \\
& \mid \iota_1 \ \bar{\tau} \ v \mid \iota_2 \ \bar{\tau} \ v \mid \mathbf{map}\langle v \rangle^{\tau'} \ \bar{\tau} \mid (v) \ \tau' \ \bar{\tau} \mid \pi_1 \ \bar{\tau} \mid \pi_2 \ \bar{\tau} \\
& \mid (v_1 \ \blacktriangledown \ v_2) \ \bar{\tau} \mid \mathbf{tt} \ \bar{\tau} \mid \mathbf{absurd} \ \bar{\tau} \ v \\
\\
E & := [] \mid E \ M \mid v \ E \mid E \ \tau \mid \mathbf{let} \ (x : \sigma) = E \ \mathbf{in} \ M \mid \mathbf{let} \ (x : \sigma) = v \ \mathbf{in} \ E & \text{(Contexts)} \\
& \mid \mathbf{map}\langle E \rangle^{\tau} \mid (E) \ \tau \mid E \ \blacktriangle \ M \mid v \ \blacktriangle \ E \mid E \ \blacktriangledown \ M \mid v \ \blacktriangledown \ E
\end{aligned}$$

which depends on the kind of its arrow type), also yields a value. In fact, all the *partial applications* of ι_1 with only some of its type arguments, or all type arguments but no value argument, are also values. We use gray highlights to indicate such an optional application with type and/or value arguments in the definition of values.

Figure 25 defines the reduction rules. We split the rules in two categories: the first set describes β -reduction⁶⁰ for the various type formers, while the second set determines how the $\mathbf{map}\langle - \rangle^-$ primitive computes. Similar to the definition of values and contexts in Figure 24, we use the notation $\bar{\tau}$ to depict a sequence of zero or more type applications. Unlike for values, these type arguments are not optional; terms typed by an arrow types must be fully applied with all their type arguments before they reduce. The notation $N \bullet M$ is used as a syntactic shorthand for the composition of two arrow types, which is defined through η -expansion of all its type arguments and the term argument. The reduction rules for the $\mathbf{map}\langle \tau \rangle^M$ primitive are type directed, in the sense that the selected reduction depends on τ . This is necessary, because in an application of $\mathbf{map}\langle - \rangle^-$ to a value, there is no way to decide whether to apply the function or to push the $\mathbf{map}\langle - \rangle^-$ further inwards by only looking at the value.

Figure 24: Values and Evaluation Contexts. Highlights indicate optional occurrences of (type) arguments.

⁶⁰ Here, we mean “ β -reduction” in the more general sense of simplifying an application of an elimination form to an introduction form.

$$\begin{aligned}
& ((\lambda x.M) v) \longrightarrow M[v/x] & (1) \\
\mathbf{let} (x : \sigma) = v \mathbf{in} M & \longrightarrow M[v/x] & (2) \\
& (\Lambda \alpha.M) \tau \longrightarrow M[\tau/\alpha] & (3) \\
& \mathbf{unin} \bar{\tau} (\mathbf{in} \bar{\tau} v) \longrightarrow v & (4) \\
& \langle v_1 \rangle^{\tau'} \bar{\tau} (\mathbf{in} \bar{\tau} v_2) \longrightarrow v_1 \bar{\tau} (\mathbf{map}\langle v_1 \rangle^{\tau'} \bar{\tau} v_2) & (5) \\
\pi_1 \bar{\tau} ((v_1 \blacktriangle v_2) \bar{\tau} v) & \longrightarrow v_1 \bar{\tau} v & (6) \\
\pi_2 \bar{\tau} ((v_1 \blacktriangle v_2) \bar{\tau} v) & \longrightarrow v_2 \bar{\tau} v & (7) \\
(v_1 \blacktriangledown v_2) \bar{\tau} (\iota_1 \bar{\tau} v) & \longrightarrow v_1 \bar{\tau} v & (8) \\
(v_1 \blacktriangledown v_2) \bar{\tau} (\iota_2 \bar{\tau} v) & \longrightarrow v_2 \bar{\tau} v & (9) \\
& \mathbf{map}\langle v_1 \rangle^{(\lambda x.x)} \bar{\tau} v_2 \longrightarrow v_1 \bar{\tau} v_2 & (10) \\
& \mathbf{map}\langle v_1 \rangle^{\mu(\tau')} \bar{\tau} (\mathbf{in} \bar{\tau} v_2) \longrightarrow \mathbf{in} \bar{\tau} (\mathbf{map}\langle v_1 \rangle^{(\tau' \mu(\tau'))} \bar{\tau} v_2) & (11) \\
\mathbf{map}\langle v \rangle^{\tau_1 \times \tau_2} \bar{\tau} ((v_1 \blacktriangle v_2) \bar{\tau} v_3) & \\
\longrightarrow ((\mathbf{map}\langle v \rangle^{\tau_1} \bullet v_1) \blacktriangle (\mathbf{map}\langle v \rangle^{\tau_2} \bullet v_2)) \bar{\tau} v_3 & (12) \\
\mathbf{map}\langle v_1 \rangle^{\tau_1 + \tau_2} \bar{\tau} (\iota_1 \bar{\tau} v_2) & \longrightarrow \iota_1 \bar{\tau} (\mathbf{map}\langle v_1 \rangle^{\tau_1} \bar{\tau} v_2) & (13) \\
\mathbf{map}\langle v_1 \rangle^{\tau_1 + \tau_2} \bar{\tau} (\iota_2 \bar{\tau} v_2) & \longrightarrow \iota_2 \bar{\tau} (\mathbf{map}\langle v_1 \rangle^{\tau_2} \bar{\tau} v_2) & (14) \\
& \mathbf{map}\langle v \rangle^{\mathbf{1}} \bar{\tau} (\mathbf{tt} \bar{\tau}) \longrightarrow \mathbf{tt} \bar{\tau} & (15) \\
& N \bullet M \triangleq \overline{\Lambda \alpha. \lambda x. N} \bar{\alpha} (M \bar{\alpha} x)
\end{aligned}$$

Figure 25: Reduction rules

5.5.2 Relation to the Denotational Model

⁶¹ This property implies what Devesas Campos and Levy [2018] call soundness of the denotational model with respect to the operational model. Their soundness property is about a big-step relation; ours is small-step.

The reduction rules shown in Figure 25 define a computational model for our calculus. We now discuss how this model arises from the denotational model discussed in Section 5.4. Informally speaking, reducing a term should not change its meaning. This intuition is reflected by the following implication, which states if M reduces to N , their semantics should be equal.⁶¹

$$M \longrightarrow N \implies \llbracket M \rrbracket = \llbracket N \rrbracket \quad (4)$$

While we do not give a formal proof of the implication above, by relying on the categorical model to inform how terms compute we can be reasonably confident that our semantics does not contain any reductions that violate this property. That is, all the reductions shown in Figure 25 are supported by an equality of morphisms in the categorical model.

What does this mean, specifically? The semantics of well-typed terms is given by a natural transformation, so if $M \rightarrow N$, M and N should be interpreted as the same natural transformation. Equivalence of natural transformations is defined pointwise in terms of the equality relation for morphisms in the underlying category. In our case, this is the category \mathbf{SET} , as terms are interpreted as natural transformations between functors into \mathbf{SET} . By studying the properties—expressed as equalities between morphisms—of the constructions that give a semantics to the different type formers, and reifying these equalities as syntactic reduction rules, we obtain an operational model that we conjecture respects the denotational model by construction.

Let us illustrate this principle with a concrete example. The semantics of a sum type $\tau_1 + \tau_2 : k$ is given by a coproduct in the category $\llbracket k \rrbracket$. The universal property of coproducts tells us that $[f, g] \circ \iota_1 = f$ and $[f, g] \circ \iota_2 = g$, or in other words, constructing and then immediately deconstructing a coproduct is the same as doing nothing. Rules (8) and (9) in Figure 25 reflect these equations. That is, since the ι_1 , ι_2 , and $[-, -]$ primitives are interpreted as the injections ι_1 , ι_2 , and unique morphism $[-, -]$ respectively, the universal property of coproducts tells us that the left-hand side and right-hand side of rule (8) and (9) in Figure 25 are interpreted to equal morphism in the categorical domain.

The remaining reduction rules are justified by the categorical model in a similar fashion. More specifically:

- Rules (1,2) follow from the β -law for exponential objects, which states that $\text{eval} \circ \langle \text{curry}(f), \text{id} \rangle = f$.

- Rule (3) holds definitionally, assuming type substitution is appropriately defined such that it corresponds to functor application.
- Rule (4) follows from Lambek’s lemma, which states that the component of an initial algebra is always an isomorphism. That is, there exists a morphism unin such that $\text{unin} \circ \text{in} = \text{id}$.
- Rule (5) reflects the universal property of folds, i.e., $\text{cata}(f) \circ \text{in} = f \circ F(\text{cata}(f))$.
- Rules (6,7) follow from the universal property of products, which states that $\pi_1 \circ \langle f, g \rangle = f$ and $\pi_2 \circ \langle f, g \rangle = g$.
- Rule (10) mirrors the identity law for functors, i.e. $F(\text{id}) = \text{id}$.
- Rule (11) is derived from naturality of the component of the initial algebra of higher-order functors, which states that $\mu(F)(f) \circ \text{in} = \text{in} \circ F(\mu(F))(f)$.
- Rule (12,13,14,15) are derived from the way (co)-limits are computed pointwise in functor categories. For example, the morphism action of the product of two functors F and G is defined as $(F \times G)(f) = \langle F(f) \circ \pi_1, G(f) \circ \pi_2 \rangle$, which gives rise to rule (12).

5.6 RELATED WORK

The problem of equipping functional languages with better support for modularity as been studied extensively in the literature. One of the earlier instances is the *Algebraic Design Language* (ADL) by [Kieburtz and Lewis \[1995\]](#), which features language primitives for specifying computable functions in terms of algebras. ADL overlaps to a large extent with the first-order fragment of our calculus, but lacks support for defining nested data types. [Zhang et al. \[2021\]](#) recently proposed a calculus and language for *compositional programming*, called CP. Their

language design is inspired by *object algebras*, which in turn is based on the *tagless final* approach [Carette et al., 2009a, Kiselyov, 2010] and *final algebra semantics* [Wand, 1979], which, according to [Wand, 1979, §7], is an extension of *initial algebra semantics*. These lines of work thus provide similar modularity as initial algebra semantics, but in a way that does not require *tagged values*. While the categorical foundations of Zhang et al.’s CP language seems to be an open question, the language provides flexible support for modular programming, in part due to its powerful notion of subtyping. We are not aware of attempts to model (higher-order) effects and handlers using CP. In contrast, our calculus is designed to have a clear categorical semantics. This semantics makes it straightforward to define state of the art type safe modular (higher-order) effects and handlers. Morris and McKinna [2019a] define a language that has built-in support for *row types*, which supports both extensible records and variants. While their language captures many known flavors of extensibility, due to parameterizing the type system over a so-called *row theory* describing how row types behave under composition, rows are restricted to first order types. Consequently, they cannot describe any modularity that hinges on the composition of (higher-order) signature functors.

The question of including nested data types in a language’s support for modularity has received some attention as well. For example, Cai et al. [2016] develop an extension of F_ω with equirecursive types tailored to describe patterns from datatype generic programming. Their calculus is expressive enough to capture the modularity abstractions discussed in this paper, including those requiring nested data types, but lacks a denotational model; a correspondence between a subset of types in their calculus and (traversable) functors is discussed informally. Similarly, Abel et al. [2005] consider an operational perspective of traversals over nested datatypes by studying several extensions of F_ω with primitives for (*generalized*) *Mendler iteration and coiteration*. Although these are expressive enough to describe modular higher-order effects and handlers, their semantic foun-

dation is very different from the semantics of the primitive fold operation in our calculus. It is future work to investigate how our calculus can be extended with support for codata.

A major source of inspiration for the work in this paper are recent works by [Johann and Polonsky \[2019\]](#), [Johann et al. \[2021\]](#), and [Johann and Ghiorzi \[2021\]](#), which respectively study the semantics and parametricity of nested data types and GADTs. For the latter, the authors develop a dedicated calculus with a design and semantics that is very similar to ours. Still, there are some subtle but key differences between the designs; for example, their calculus does not include general notions of \forall -types and function types, but rather integrates these into a single type representing natural transformations between type constructors. While their setup does not require the same stratification of the type syntax we adopt here, it is also slightly less expressive, as the built-in type of transformations is restricted to closing over 0-arity arguments.

Data type generic programming commonly uses a *universe of descriptions* [[Benke et al., 2003](#)], which is a data type whose inhabitants correspond to a signature functor. Generic functions are commonly defined by induction over these descriptions, ranging over a semantic reflection of the input description in the type system of a dependently-typed host language [[Dagand, 2013a](#)]. In fact, [Chapman et al. \[2010a\]](#) considered the integration of descriptions in a language’s design by developing a type theory with native support for generic programming. We are, however, not aware of any notion of descriptions that corresponds to our syntax of well-formed types.

5.7 CONCLUSION AND FUTURE WORK

In this paper, we presented the design and semantics of a calculus with support for modularity. We demonstrated it can serve as a basis for capturing several well-known programming patterns for retrofitting type-safe modularity to functional languages, such as modular interpreters in the style of Data Types

à la Carte, and modular (higher-order) algebraic effects. The formal semantics associates these patterns with their motivating concepts, creating the possibility for a compiler to benefit from their properties such as by performing fusion-based optimizations.

POSTSCRIPT

This chapter contributed a core calculus that nicely captures the essence of how extensible data types can be incorporated in the design of a functional meta language in a principled way. How do we connect this result to the overarching goal of enabling language designers to define reusable programming language components as intrinsically-typed definitional interpreters?

There are several steps to take. First and foremost, the results should be extended to be able to describe modularity of *indexed data types*, as discussed in Section 2.2, to allow us to encode intrinsically-typed interpreters in the language. This introduces additional challenges, as functions that eliminate indexed data types locally rely on additional equalities to rule out redundant cases for ill-typed inputs. Beyond that, a key selling point of intrinsically-typed definitional interpreters is their readability, which is mostly due to *dependent pattern matching* [Cockx, 2017]. A meta language design should include this, and other syntactical conveniences, to be a feasible option for defining reusable programming language components.

CONCLUSIONS



CONCLUSIONS

The core problem we addressed in this thesis is the high development cost associated with the development of programming languages with formally specified type systems and verified type soundness. To reduce this cost, we pinpointed two areas in which it is necessary to make progress:

1. the development of semantic techniques that allow us to specify the semantics of language constructs in a way that is independent of the context (i.e., language) in which they are used, and
2. the design of (functional) meta-languages that support modularity out-of-the-box, that can be used to define modular and reusable language components without having to deal with the additional overhead that is usually incurred by modularity.

In the remainder of this chapter, we will revisit the contributions of this thesis, and reflect on these contributions in light of the hypotheses from Section [1.2](#).

6.1 SUMMARY OF THE CONTRIBUTIONS

Each chapter in this thesis contains its own contributions. Here, we summarize the core contributions of this thesis.

In Chapter [2](#), we showed how to embed modular intrinsically-typed definitional interpreters in Agda, and developed a no-

tion of *intrinsically-typed language fragments*, that form a self-contained, type safe, and modular specification of a language construct. Language fragments can be freely composed and reused to build bigger programming languages in a way that all these properties are preserved.

In Chapter 3, we defined abstractions for defining *modular semantics for higher-order effects*, and developed reasoning infrastructure to prove these semantics correct with respect to equational theories that give an axiomatic specification of an effect's intended behavior. This enables us to give a modular semantics to language constructs such as λ -abstraction or exception catching, and reason about the correctness of these semantics.

In Chapter 4, we developed the design of a domain-specific meta language for defining reusable programming language components. Key features of the design are an *effect system that supports higher-order effects*, and *native support for modular data types*. As a result, the language allows us to write modular interpreters (effectively solving the *expression problem* [Wadler, 1998]) without having to resort to shallow embeddings of initial algebra semantics.

Finally, in Chapter 5, we presented a core calculus with *built-in primitives for modularity*, together with a type system and semantics. The calculus is expressive enough to capture many familiar programming abstractions for modularity, such as modular data types in the style of Data Types á la Carte [Swierstra, 2008], or inductively-defined free monads. The calculus provides a formal basis for understanding languages with built-in support for modular data types, such as the language design from Chapter 4.

6.2 HYPOTHESIS 1: INTRINSICALLY-TYPED INTERPRETERS

In part I of this thesis (Chapters 2 and 3), we explored the following hypothesis:

reusable programming language components should be concise, readable, and safe-by-construction. Intrinsically-typed definitional interpreters are an excellent match for these requirements.

In Chapter 2, we developed techniques for writing modular intrinsically-typed definitional interpreters. Crucially, we pinpointed *canonicity-preserving* separation and inclusion predicates (Section 2.3) that witness respectively type-safe union and extension of value domains. These turned out to be the key ingredients for defining intrinsically-typed interpreters in a way that they can be composed in a safe-by-construction manner.

This is, however, a mere first step towards using intrinsically-typed definitional interpreters for the specification of reusable programming language components. As we pointed out in Section 2.5, the abstractions developed in Chapter 2 necessarily presuppose both the notion of typing and semantics used to define a language fragment. Although the techniques we developed for modularizing intrinsically-typed definitional interpreters can be generalized to languages with lexical binding that denote into a more expressive semantic domain (Section 2.5), this leaves the question of whether the techniques apply to language components that do not fit this abstraction. Beyond that it is also still an open question whether language fragments could be generalized even further to allow the language designer to choose a flavor of typing and semantics that fits their language project.

ANOTHER QUESTION LEFT open at the end of Chapter 2 is how to adequately deal with side effects of the interpreter. In the setup used in Section 2.5, we make the assumption that a suitable monad exists that encapsulates all the required side effects for denoting a particular combination of language fragment. Depending on the combination of effects, constructing this monad may actually be tricky. Specifically, if we require higher-order effects and/or *intrinsically-typed effects*⁶² to denote a language, existing treatments of modular effects are insufficient.

⁶² An example of a monad modelling intrinsically-typed effects is the monad used by Bach et al. [2018] to denote a language with polymorphic references.

In Chapter 3, we give a partial answer to this question, by developing a unified approach for defining modular semantics for higher-order effects by elaborating them to algebraic effects. While the approach solves the modularity problem for higher-order effects, it is not equipped to deal with intrinsically-typed effects yet. As a result, we cannot yet use it to define the monad(s) assumed in Section 2.5 in a modular fashion. Most likely, this would require us to generalize the current approach, where we construct monad morphisms between different free monads in the category of Agda sets, to categories of monotone predicates to add the necessary *world-indexing* for modelling intrinsically-typed effects such as well-typedness of polymorphic references.

All in all, we conclude that the contributions presented in Chapters 2 and 3 provide us with a solid basis for understanding modularity of intrinsically-typed interpreters and higher-order effects. Within the current state of affairs we can already construct reusable language components for many different language constructs by writing intrinsically-typed definitional interpreters. However, before we can scale the approach to more advanced type systems, such as the interpreters for imperative languages by Bach et al. [2018] or the interpreter for linear session-typed languages by Rouvoet et al. [2020], there are some open research questions to be addressed first. Even more expressive systems, such as System F and dependent type theory, remain out of reach until the intrinsically-typed approach has been extended to these languages.

6.3 HYPOTHESIS 2: META LANGUAGE DESIGN

In part II (Chapters 4 and 5) of this thesis, we explored the following hypothesis:

modularity adds significant syntactic and interpretative overhead when using state-of-the art (dependently-typed) programming languages to define reusable programming language components. Incorporating

modular inductive data types in the design of functional languages is therefore an essential first step in the development of meta languages for the purpose of defining reusable programming language components.

In Chapter 4, we presented a meta language design for defining reusable programming language components, that features built-in support for modular data types. The language supports the definition of modular interpreters, solving the expression problem [Wadler, 1998], without incurring the usual syntactic overhead associated with the encodings of initial algebra semantics that are typically used to retrofit this kind of type-safe modularity to functional languages. However, the language lacks a formal specification of its type system and semantics.

In Chapter 5, we worked towards addressing this shortcoming by developing a calculus with built in support for modularity, together with a type system and semantics. While this does allow us to encode many common programming abstractions for modularity in a type-safe manner (including the kind of language components we defined in Chapter 4), there is still a gap between the languages. To close this gap we would need to, for example, define a desugaring from the meta language defined in Chapter 4 to the calculus defined in Chapter 5. Furthermore, the modular data types in these languages are not expressive enough to capture the modular interpreters discussed in Chapter 2. This would require the design of a meta language with built-in support for modular *indexed data types*.

6.4 FUTURE WORK

This thesis presented contributions towards making formally verified language components more accessible, with the goal of making formally specified and verified programming languages attainable for a broader audience of language engineers. To get there, however, there are still some challenges that need to be addressed.

6.4.1 *Modular semantics*

When revisiting the first hypothesis in Section 6.2, we already hinted at the open challenges when it comes to using intrinsically-typed definitional interpreters to specify verified and reusable programming language components. In particular, we would need to improve the expressivity of the approach by further generalizing both the notion of typing as well as the semantic domain of language fragments.

To generalize the notion of typing permitted, generalizations of effect systems such as the one by Mycroft et al. [2016] seem like a good starting point. This would, however, require us to further the state-of-the-art in intrinsically-typed definitional interpreters to this class of typings, before we can think about extending the definition of language fragments accordingly. Furthermore, this induces the question of how to compose fragments that use different typing. While it is conceivable that we could leverage similar separation and inclusion predicates as we used to capture assumptions about canonicity to phrase a fragment's typing in terms of assumptions (rather than fixing it upfront), we have yet to explore this direction.

WILE THE ABOVE concerns *static effects*, similar modularity concerns arise when considering *dynamic effects*. To give a semantics to languages with polymorphic references, for example, we denote into a category of monotone predicates over the memory state. Other language fragments, however, may maintain different invariants, and it is still an open question how we would mediate the different requirements on the semantic domain imposed by such fragments.

The setup in Section 2.5 already hints at a possible solution direction for this problem by factoring all effects into a monadic interface. Behind the scenes, for the languages that have polymorphic references, this monad is implemented by constructing an inductively-defined free monad in a category of monotone predicates. Similarly, the intrinsically-typed interpreter for session-typed λ -calculus by Rouvoet et al. [2020]

relies on a inductively defined free monad in the category of predicates over (runtime) session types. Finding a general and modular type of command tree [Hancock and Setzer, 2000] that subsumes these and other intrinsically-typed dynamic effects would be a crucial step in achieving a modular treatment of effects.

6.4.2 *Meta Language Design*

When it comes to meta language design, the languages presented in Chapter 4 and Chapter 5 occupy two very different points in the design space. Where the former works backwards from the programmer’s perspective and focusses more on usability, the second puts more emphasis on mathematical rigor and a formal specification. This leaves use with a clear goal of combining these two different perspectives in a single language design. We could get there by adding a formal specification of a type system to the language presented in Chapter 4, or alternatively define a desugaring into the calculus presented in Chapter 5.

Beyond that, when it comes to usability there are more language features we could consider adding which do not currently appear in either language. For example, *functor subtyping* or *row typing* are obvious candidate features that would improve usability.

6.4.3 *Connecting the Dots*

Aside from the work that is still to be done on the individual parts, we may look ahead and think about how to connect the contributions of the two parts in this thesis. After all, Chapters 2 and 3 are concerned with developing the semantic tools necessary to describe reusable language components as intrinsically-typed definitional interpreters, whereas Chapters 4 and 5 investigate the design of meta languages that support the modular definition of untyped definitional interpreters.

The vision that connects these lines of work is to develop a dedicated meta language for developing reusable programming language components by writing intrinsically-typed definitional interpreters. Such a language would allow us write intrinsically-typed definitional interpreters much like we write them in Agda today, but with the key difference that they are modular and composable. Before we are ready to embark on designing such a language, however, we must have a solid understanding of the semantics of composing intrinsically-typed definitional interpreters, as well as how to incorporate modularity for data types and their operations in a language's design and type system. In this thesis, we have laid the necessary foundation, meaning that the time is ripe to start working towards making this vision a reality.

APPENDIX

BIBLIOGRAPHY

- Michael Gordon Abbott, Thorsten Altenkirch, and Neil Ghani. Categories of containers. In Andrew D. Gordon, editor, *Foundations of Software Science and Computational Structures, 6th International Conference, FOSSACS 2003 Held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings*, volume 2620 of *Lecture Notes in Computer Science*, pages 23–38. Springer, 2003. URL https://doi.org/10.1007/3-540-36576-1_2.
- Michael Gordon Abbott, Thorsten Altenkirch, and Neil Ghani. Containers: Constructing strictly positive types. *Theor. Comput. Sci.*, 342(1):3–27, 2005a. URL <https://doi.org/10.1016/j.tcs.2005.06.002>.
- Michael Gordon Abbott, Thorsten Altenkirch, and Neil Ghani. Containers: Constructing strictly positive types. *Theor. Comput. Sci.*, 342(1):3–27, 2005b. URL <https://doi.org/10.1016/j.tcs.2005.06.002>.
- Andreas Abel and Ralph Matthes. (co-)iteration for higher-order nested datatypes. In Herman Geuvers and Freek Wiedijk, editors, *Types for Proofs and Programs, Second International Workshop, TYPES 2002, Berg en Dal, The Netherlands, April 24-28, 2002, Selected Papers*, volume 2646 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2002. URL https://doi.org/10.1007/3-540-39185-1_1.
- Andreas Abel, Ralph Matthes, and Tarmo Uustalu. Generalized iteration and coiteration for higher-order nested datatypes. In Andrew D. Gordon, editor, *Foundations of Software Science and Computational Structures, 6th International Conference, FOSSACS 2003 Held as Part of the Joint European Conference on Theory and*

- Practice of Software, ETAPS 2003, Warsaw, Poland, April 7-11, 2003, Proceedings*, volume 2620 of *Lecture Notes in Computer Science*, pages 54–69. Springer, 2003. URL https://doi.org/10.1007/3-540-36576-1_4.
- Andreas Abel, Ralph Matthes, and Tarmo Uustalu. Iteration and coiteration schemes for higher-order and nested datatypes. *Theor. Comput. Sci.*, 333(1-2):3–66, 2005. URL <https://doi.org/10.1016/j.tcs.2004.10.017>.
- Jiří Adámek. Free algebras and automata realizations in the language of categories. *Commentationes Mathematicae Universitatis Carolinae*, 15(4):589–602, 1974.
- Amal J. Ahmed, Andrew W. Appel, and Roberto Virga. A stratified semantics of general references A stratified semantics of general references. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, page 75. IEEE Computer Society, 2002. URL <https://doi.org/10.1109/LICS.2002.1029818>.
- Amal Jamil Ahmed. *Semantics of Types for Mutable State*. PhD thesis, USA, 2004. AAI3136691.
- Guillaume Allais, Robert Atkey, James Chapman, Conor McBride, and James McKinna. A type and scope safe universe of syntaxes with binding: their semantics and proofs. *Proc. ACM Program. Lang.*, 2(ICFP):90:1–90:30, 2018. URL <https://doi.org/10.1145/3236785>.
- Guillaume Allais, Robert Atkey, James Chapman, Conor McBride, and James McKinna. A type- and scope-safe universe of syntaxes with binding: their semantics and proofs. *J. Funct. Program.*, 31:e22, 2021. URL <https://doi.org/10.1017/S0956796820000076>.
- Thorsten Altenkirch and Ambrus Kaposi. Type theory in type theory using quotient inductive types. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM*

- SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 18–29. ACM, 2016. doi: 10.1145/2837614.2837638. URL <https://doi.org/10.1145/2837614.2837638>.
- Thorsten Altenkirch, Neil Ghani, Peter G. Hancock, Conor McBride, and Peter Morris. Indexed containers. *J. Funct. Program.*, 25, 2015. URL <https://doi.org/10.1017/S095679681500009X>.
- Nada Amin and Tiark Rompf. Type soundness proofs with definitional interpreters. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 666–679. ACM, 2017. URL <https://doi.org/10.1145/3009837.3009866>.
- Michael A. Arbib and Ernest G. Manes. *Arrows, Structures, and Functors: The Categorical Imperative*. Academic Press, 1975.
- Robert Atkey. Syntax and semantics of quantitative type theory. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 56–65. ACM, 2018. URL <https://doi.org/10.1145/3209108.3209189>.
- Lennart Augustsson and Magnus Carlsson. An exercise in dependent types: A well-typed interpreter. In *In Workshop on Dependent Types in Programming, Gothenburg, 1999*.
- Steve Awodey. *Category Theory*. Oxford Logic Guides. Oxford University Press, Oxford, New York, second edition, second edition edition, June 2010. ISBN 978-0-19-923718-0.
- 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 Joe Hurd and Thomas F. Melham, editors, *Theorem Proving in Higher Order Logics, 18th International Conference, TPHOLs 2005, Oxford, UK, August 22-25, 2005, Proceedings*, volume 3603 of *Lecture Notes in Computer Science*, pages 50–65. Springer, 2005. URL https://doi.org/10.1007/11541868_4.
- Casper Bach and Cas van der Rest. Hefty algebras: Modular elaboration of higher-order algebraic effects. *Proc. ACM Program. Lang.*, 7(POPL):1801–1831, 2023. URL <https://doi.org/10.1145/3571255>.
- Casper Bach, Arjen Rouvoet, Andrew Tolmach, Robbert Krebbers, and Eelco Visser. Intrinsically-typed definitional interpreters for imperative languages. *Proc. ACM Program. Lang.*, 2(POPL):16:1–16:34, 2018. URL <https://doi.org/10.1145/3158104>.
- Casper Bach, Cas van der Rest, and Tom Schrijvers. Staged effects and handlers for modular languages with abstraction. In *Workshop on Partial Evaluation and Program Manipulation (PEPM)*. <http://casperbp.net/store/staged-effects-and-handlers.pdf>, 2021.
- Patrick Bahr. Composing and decomposing data types: a closed type families implementation of data types à la carte. In José Pedro Magalhães and Tiark Rompf, editors, *Proceedings of the 10th ACM SIGPLAN workshop on Generic programming, WGP 2014, Gothenburg, Sweden, August 31, 2014*, pages 71–82. ACM, 2014. URL <https://doi.org/10.1145/2633628.2633635>.
- Patrick Bahr and Tom Hvitved. Compositional data types. In Jaakko Järvi and Shin-Cheng Mu, editors, *Proceedings of the seventh ACM SIGPLAN workshop on Generic programming, WGP@ICFP 2011, Tokyo, Japan, September 19-21, 2011*, pages 83–94. ACM, 2011. URL <https://doi.org/10.1145/2036918.2036930>.

- Patrick Bahr and Tom Hvitved. Parametric compositional data types. In James Chapman and Paul Blain Levy, editors, *Proceedings Fourth Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2012, Tallinn, Estonia, 25 March 2012*, volume 76 of *EPTCS*, pages 3–24, 2012a. URL <https://doi.org/10.4204/EPTCS.76.3>.
- Patrick Bahr and Tom Hvitved. Parametric compositional data types. In James Chapman and Paul Blain Levy, editors, *Proceedings Fourth Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2012, Tallinn, Estonia, 25 March 2012*, volume 76 of *EPTCS*, pages 3–24, 2012b. URL <https://doi.org/10.4204/EPTCS.76.3>.
- Andrej Bauer and Matija Pretnar. An effect system for algebraic effects and handlers. *Log. Methods Comput. Sci.*, 10(4), 2014. URL [https://doi.org/10.2168/LMCS-10\(4:9\)2014](https://doi.org/10.2168/LMCS-10(4:9)2014).
- Andrej Bauer and Matija Pretnar. Programming with algebraic effects and handlers. *J. Log. Algebraic Methods Program.*, 84(1):108–123, 2015. URL <https://doi.org/10.1016/j.jlamp.2014.02.001>.
- Marcin Benke, Peter Dybjer, and Patrik Jansson. Universes for generic programs and proofs in dependent type theory. *Nord. J. Comput.*, 10(4):265–289, 2003.
- Birthe van den Berg and Tom Schrijvers. A framework for higher-order effects & handlers. *CoRR*, abs/2302.01415, 2023. URL <https://doi.org/10.48550/arXiv.2302.01415>.
- Birthe van den Berg, Tom Schrijvers, Casper Bach, and Nicolas Wu. Latent effects for reusable language components. In Hakjoo Oh, editor, *Programming Languages and Systems - 19th Asian Symposium, APLAS 2021, Chicago, IL, USA, October 17-18, 2021, Proceedings*, volume 13008 of *Lecture Notes in Computer Science*, pages 182–201. Springer, 2021a. URL https://doi.org/10.1007/978-3-030-89051-3_11.

- Birthe van den Berg, Tom Schrijvers, Casper Bach, and Nicolas Wu. Latent effects for reusable language components: Extended version. *CoRR*, abs/2108.11155, 2021b. URL <https://arxiv.org/abs/2108.11155>.
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. Handle with care: relational interpretation of algebraic effects and handlers. *Proc. ACM Program. Lang.*, 2(POPL):8:1–8:30, 2018. URL <https://doi.org/10.1145/3158096>.
- Richard S. Bird and Lambert G. L. T. Meertens. Nested datatypes. In Johan Jeuring, editor, *Mathematics of Program Construction, MPC'98, Marstrand, Sweden, June 15-17, 1998, Proceedings*, volume 1422 of *Lecture Notes in Computer Science*, pages 52–67. Springer, 1998. URL <https://doi.org/10.1007/BFb0054285>.
- Richard S. Bird and Ross Paterson. Generalised folds for nested datatypes. *Formal Aspects Comput.*, 11(2):200–222, 1999. URL <https://doi.org/10.1007/s001650050047>.
- Matthias Blume, Umut A. Acar, and Wonseok Chae. Extensible programming with first-class cases. In John H. Reppy and Julia L. Lawall, editors, *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming, ICFP 2006, Portland, Oregon, USA, September 16-21, 2006*, pages 239–250. ACM, 2006. URL <https://doi.org/10.1145/1159803.1159836>.
- Jonathan Immanuel Brachthäuser, Philipp Schuster, and Klaus Ostermann. Effects as capabilities: effect handlers and lightweight effect polymorphism. *Proc. ACM Program. Lang.*, 4(OOPSLA):126:1–126:30, 2020. URL <https://doi.org/10.1145/3428194>.
- Edwin C. Brady. Programming and reasoning with algebraic effects and dependent types. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA*,

- USA - September 25 - 27, 2013, pages 133–144, 2013a. URL <https://doi.org/10.1145/2500365.2500581>.
- Edwin C. Brady. Idris, a general-purpose dependently typed programming language: Design and implementation. *J. Funct. Program.*, 23(5):552–593, 2013b. URL <https://doi.org/10.1017/S095679681300018X>.
- Yufei Cai, Paolo G. Giarrusso, and Klaus Ostermann. System f-omega with equirecursive types for datatype-generic programming. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 30–43. ACM, 2016. doi: 10.1145/2837614.2837660. URL <https://doi.org/10.1145/2837614.2837660>.
- Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *J. Funct. Program.*, 19(5):509–543, 2009a. URL <https://doi.org/10.1017/S0956796809007205>.
- Jacques Carette, Oleg Kiselyov, and Chung-chieh Shan. Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *J. Funct. Program.*, 19(5):509–543, 2009b. URL <https://doi.org/10.1017/S0956796809007205>.
- Pietro Cenciarelli and Eugenio Moggi. A syntactic approach to modularity in denotational semantics. 1993.
- James Chapman, Pierre-Évariste Dagand, Conor McBride, and Peter Morris. The gentle art of levitation. In Paul Hudak and Stephanie Weirich, editors, *Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010*, pages 3–14. ACM, 2010a. ISBN 978-1-60558-794-3. URL <https://doi.org/10.1145/1863543.1863547>.
- James Chapman, Pierre-Évariste Dagand, Conor McBride, and Peter Morris. The gentle art of levitation. In Paul Hudak

- and Stephanie Weirich, editors, *Proceeding of the 15th ACM SIGPLAN international conference on Functional programming, ICFP 2010, Baltimore, Maryland, USA, September 27-29, 2010*, pages 3–14. ACM, 2010b. ISBN 978-1-60558-794-3. URL <https://doi.org/10.1145/1863543.1863547>.
- James Chapman, Roman Kireev, Chad Nester, and Philip Wadler. System F in agda, for fun and profit. In Graham Hutton, editor, *Mathematics of Program Construction - 13th International Conference, MPC 2019, Porto, Portugal, October 7-9, 2019, Proceedings*, volume 11825 of *Lecture Notes in Computer Science*, pages 255–297. Springer, 2019. doi: 10.1007/978-3-030-33636-3_10. URL https://doi.org/10.1007/978-3-030-33636-3_10.
- Adam Chlipala. Parametric higher-order abstract syntax for mechanized semantics. In James Hook and Peter Thiemann, editors, *Proceeding of the 13th ACM SIGPLAN international conference on Functional programming, ICFP 2008, Victoria, BC, Canada, September 20-28, 2008*, pages 143–156. ACM, 2008. URL <https://doi.org/10.1145/1411204.1411226>.
- Martin Churchill and Peter D. Mosses. Modular bisimulation theory for computations and values. In Frank Pfenning, editor, *FOSSACS 2013*, volume 7794 of *LNCS*, pages 97–112. Springer, 2013.
- Martin Churchill, Peter D. Mosses, Neil Sculthorpe, and Paolo Torrini. Reusable components of semantic specifications. *LNCS Trans. Aspect Oriented Softw. Dev.*, 12:132–179, 2015.
- Matteo Cimini, Dale Miller, and Jeremy G. Siek. Extrinsically typed operational semantics for functional languages. In Ralf Lämmel, Laurence Tratt, and Juan de Lara, editors, *Proceedings of the 13th ACM SIGPLAN International Conference on Software Language Engineering, SLE 2020, Virtual Event, USA, November 16-17, 2020*, pages 108–125. ACM, 2020. URL <https://doi.org/10.1145/3426425.3426936>.

- Koen Claessen. A poor man's concurrency monad. *J. Funct. Program.*, 9(3):313–323, 1999. URL <https://doi.org/10.1017/S0956796899003342>.
- Jesper Cockx. Dependent pattern matching and proof-relevant unification. 2017. URL <https://lirias.kuleuven.be/handle/123456789/583556>.
- Lukas Convent, Sam Lindley, Conor McBride, and Craig McLaughlin. Doo bee doo bee doo. *J. Funct. Program.*, 30:e9, 2020. URL <https://doi.org/10.1017/S0956796820000039>.
- Thierry Coquand. Pattern matching with dependent types. In *Proceedings of the Workshop on Types for Proofs and Programs*, pages 71–83. Citeseer, 1992.
- Thierry Coquand and Christine Paulin. Inductively defined types. In Per Martin-Löf and Grigori Mints, editors, *COLOG-88, International Conference on Computer Logic, Tallinn, USSR, December 1988, Proceedings*, volume 417 of *Lecture Notes in Computer Science*, pages 50–66. Springer, 1988. ISBN 3-540-52335-9. URL https://doi.org/10.1007/3-540-52335-9_47.
- Bruno C. d. S. Oliveira and William R. Cook. Extensibility for the masses - practical extensibility with object algebras. In James Noble, editor, *ECOOP 2012 - Object-Oriented Programming - 26th European Conference, Beijing, China, June 11-16, 2012. Proceedings*, volume 7313 of *Lecture Notes in Computer Science*, pages 2–27. Springer, 2012. ISBN 978-3-642-31056-0. URL https://doi.org/10.1007/978-3-642-31057-7_2.
- Bruno C. d. S. Oliveira, Shin-Cheng Mu, and Shu-Hung You. Modular reifiable matching: a list-of-functors approach to two-level types. In Ben Lippmeier, editor, *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, September 3-4, 2015*, pages 82–93. ACM, 2015. URL <https://doi.org/10.1145/2804302.2804315>.
- Pierre-Évariste Dagand. *A cosmology of datatypes : reusability and dependent types*. PhD thesis, University of Strathclyde,

- Glasgow, UK, 2013a. URL http://oleg.lib.strath.ac.uk/R/?func=dbin-jump-full&object_id=22713.
- Pierre-Évariste Dagand. *A cosmology of datatypes : reusability and dependent types*. PhD thesis, University of Strathclyde, Glasgow, UK, 2013b. URL http://oleg.lib.strath.ac.uk/R/?func=dbin-jump-full&object_id=22713.
- Pierre-Évariste Dagand. The essence of ornaments. *J. Funct. Program.*, 27:e9, 2017. URL <https://doi.org/10.1017/S0956796816000356>.
- Pierre-Évariste Dagand and Conor McBride. Transporting functions across ornaments. *J. Funct. Program.*, 24(2-3):316–383, 2014. URL <https://doi.org/10.1017/S0956796814000069>.
- Benjamin Delaware, Bruno C. d. S. Oliveira, and Tom Schrijvers. Meta-theory à la carte. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, pages 207–218. ACM, 2013a. URL <https://doi.org/10.1145/2429069.2429094>.
- Benjamin Delaware, Bruno C. d. S. Oliveira, and Tom Schrijvers. Meta-theory à la carte. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, pages 207–218. ACM, 2013b. URL <https://doi.org/10.1145/2429069.2429094>.
- Benjamin Delaware, Steven Keuchel, Tom Schrijvers, and Bruno C. d. S. Oliveira. Modular monadic meta-theory. pages 319–330, 2013c. URL <https://doi.org/10.1145/2500365.2500587>.
- Marco Devesas Campos and Paul Blain Levy. A syntactic view of computational adequacy. In Christel Baier and Ugo Dal Lago, editors, *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held*

as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, *Proceedings*, volume 10803 of *Lecture Notes in Computer Science*, pages 71–87. Springer, 2018. ISBN 978-3-319-89365-5. URL https://doi.org/10.1007/978-3-319-89366-2_4.

Dominique Devriese and Frank Piessens. On the bright side of type classes: instance arguments in agda. In *ACM SIGPLAN international conference on Functional Programming (ICFP)*, pages 143–155, 2011.

Matthias Felleisen and Robert Hieb. The revised report on the syntactic theories of sequential control and state. *Theor. Comput. Sci.*, 103(2):235–271, 1992. URL [https://doi.org/10.1016/0304-3975\(92\)90014-7](https://doi.org/10.1016/0304-3975(92)90014-7).

Andrzej Filinski. Representing layered monads. In *POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999*, pages 175–188, 1999. URL <https://doi.org/10.1145/292540.292557>.

Marcelo P. Fiore and Sam Staton. Substitution, jumps, and algebraic effects. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 41:1–41:10, 2014. URL <https://doi.org/10.1145/2603088.2603163>.

Benedict R Gaster and Mark P Jones. A polymorphic type system for extensible records and variants. Technical report, Citeseer, 1996.

Johan Glimming and Neil Ghani. Difunctorial semantics of object calculus. In Viviana Bono, Michele Bugliesi, and Sophia Drossopoulou, editors, *Proceedings of the Second Workshop on Object Oriented Developments, WOOD 2004, London, UK, August 30, 2004*, volume 138 of *Electronic Notes in Theoretical Computer*

- Science*, pages 79–94. Elsevier, 2004. URL <https://doi.org/10.1016/j.entcs.2005.09.012>.
- Joseph A Goguen. An initial algebra approach to the specification, correctness and implementation of abstract data types. *IBM Research Report*, 6487, 1976.
- Peter G. Hancock and Anton Setzer. Interactive programs in dependent type theory. In Peter Clote and Helmut Schwichtenberg, editors, *Computer Science Logic, 14th Annual Conference of the EACSL, Fischbachau, Germany, August 21-26, 2000, Proceedings*, volume 1862 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 2000. URL https://doi.org/10.1007/3-540-44622-2_21.
- Robert Harper. A simplified account of polymorphic references. *Inf. Process. Lett.*, 51(4):201–206, 1994.
- C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, 1969. URL <https://doi.org/10.1145/363235.363259>.
- Kohei Honda, Vasco Thudichum Vasconcelos, and Makoto Kubo. Language primitives and type discipline for structured communication-based programming. In Chris Hankin, editor, *Programming Languages and Systems - ESOP'98, 7th European Symposium on Programming, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS'98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings*, volume 1381 of *Lecture Notes in Computer Science*, pages 122–138. Springer, 1998. URL <https://doi.org/10.1007/BFb0053567>.
- Alex Hubers and J. Garrett Morris. Generic programming with extensible data types: Or, making ad hoc extensible data types less ad hoc. *Proc. ACM Program. Lang.*, 7(ICFP), aug 2023. URL <https://doi.org/10.1145/3607843>.
- Paul Hudak, Simon L. Peyton Jones, Philip Wadler, Brian Boutel, Jon Fairbairn, Joseph H. Fasel, María M. Guzmán, Kevin Hammond, John Hughes, Thomas Johnsson, Richard B. Kieburtz,

- Rishiyur S. Nikhil, Will Partain, and John Peterson. Report on the programming language haskell, A non-strict, purely functional language. *ACM SIGPLAN Notices*, 27(5):1, 1992. URL <https://doi.org/10.1145/130697.130699>.
- Martin Hyland, Gordon D. Plotkin, and John Power. Combining effects: Sum and tensor. *Theor. Comput. Sci.*, 357(1-3):70–99, 2006. URL <https://doi.org/10.1016/j.tcs.2006.03.013>.
- Mauro Jaskelioff. Monatron: An extensible monad transformer library. In *Implementation and Application of Functional Languages - 20th International Symposium, IFL 2008, Hatfield, UK, September 10-12, 2008. Revised Selected Papers*, pages 233–248, 2008. URL https://doi.org/10.1007/978-3-642-24452-0_13.
- Johan Jeuring and Erik Meijer, editors. *Advanced Functional Programming, First International Spring School on Advanced Functional Programming Techniques, Båstad, Sweden, May 24-30, 1995, Tutorial Text*, volume 925 of *Lecture Notes in Computer Science*, 1995. Springer. ISBN 3-540-59451-5. URL <https://doi.org/10.1007/3-540-59451-5>.
- Patricia Johann and Neil Ghani. Initial algebra semantics is enough! In Simona Ronchi Della Rocca, editor, *Typed Lambda Calculi and Applications, 8th International Conference, TLCA 2007, Paris, France, June 26-28, 2007, Proceedings*, volume 4583 of *Lecture Notes in Computer Science*, pages 207–222. Springer, 2007. URL https://doi.org/10.1007/978-3-540-73228-0_16.
- Patricia Johann and Enrico Ghiorzi. Parametricity for nested types and gads. *Log. Methods Comput. Sci.*, 17(4), 2021. URL [https://doi.org/10.46298/lmcs-17\(4:23\)2021](https://doi.org/10.46298/lmcs-17(4:23)2021).
- Patricia Johann and Andrew Polonsky. Higher-kinded data types: Syntax and semantics. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver*,

- BC, Canada, June 24-27, 2019, pages 1–13. IEEE, 2019. URL <https://doi.org/10.1109/LICS.2019.8785657>.
- Patricia Johann, Enrico Ghiorzi, and Daniel Jeffries. Parametricity for primitive nested types. In Stefan Kiefer and Christine Tasson, editors, *Foundations of Software Science and Computation Structures - 24th International Conference, FOSSACS 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021, Proceedings*, volume 12650 of *Lecture Notes in Computer Science*, pages 324–343. Springer, 2021. URL https://doi.org/10.1007/978-3-030-71995-1_17.
- Thomas Johnsson. Attribute grammars as a functional programming paradigm. In Gilles Kahn, editor, *Functional Programming Languages and Computer Architecture, Portland, Oregon, USA, September 14-16, 1987, Proceedings*, volume 274 of *Lecture Notes in Computer Science*, pages 154–173. Springer, 1987. URL https://doi.org/10.1007/3-540-18317-5_10.
- Mark P. Jones. Functional programming with overloading and higher-order polymorphism. In [Jeuring and Meijer \[1995\]](#), pages 97–136. ISBN 3-540-59451-5. URL https://doi.org/10.1007/3-540-59451-5_4.
- Mark P. Jones and Luc Duponcheel. Composing monads. Research Report YALEU/DCS/RR-1004, Yale University, New Haven, Connecticut, USA, December 1993. URL <http://web.cecs.pdx.edu/~mpj/pubs/RR-1004.pdf>.
- Guy L. Steele Jr. Building interpreters by composing monads. In Hans-Juergen Boehm, Bernard Lang, and Daniel M. Yellin, editors, *Conference Record of POPL'94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon, USA, January 17-21, 1994*, pages 472–492. ACM Press, 1994. URL <https://doi.org/10.1145/174675.178068>.

- Ohad Kammar, Sam Lindley, and Nicolas Oury. Handlers in action. In Greg Morrisett and Tarmo Uustalu, editors, *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27, 2013*, pages 145–158. ACM, 2013. URL <https://doi.org/10.1145/2500365.2500590>.
- Steven Keuchel and Tom Schrijvers. Generic datatypes à la carte. In Jacques Carette and Jeremiah Willcock, editors, *Proceedings of the 9th ACM SIGPLAN workshop on Generic programming, WGP 2013, Boston, Massachusetts, USA, September 28, 2013*, pages 13–24. ACM, 2013. URL <https://doi.org/10.1145/2502488.2502491>.
- Richard B. Kieburtz and Jeffrey Lewis. Programming with algebras. In [Jeuring and Meijer \[1995\]](#), pages 267–307. ISBN 3-540-59451-5. URL https://doi.org/10.1007/3-540-59451-5_8.
- Oleg Kiselyov. Typed tagless final interpreters. In Jeremy Gibbons, editor, *Generic and Indexed Programming - International Spring School, SSGIP 2010, Oxford, UK, March 22-26, 2010, Revised Lectures*, volume 7470 of *Lecture Notes in Computer Science*, pages 130–174. Springer, 2010. ISBN 978-3-642-32201-3. URL https://doi.org/10.1007/978-3-642-32202-0_3.
- Oleg Kiselyov and Hiromi Ishii. Freer monads, more extensible effects. In Ben Lippmeier, editor, *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, September 3-4, 2015*, pages 94–105. ACM, 2015. URL <https://doi.org/10.1145/2804302.2804319>.
- Oleg Kiselyov, Amr Sabry, and Cameron Swords. Extensible effects: an alternative to monad transformers. In Chungchieh Shan, editor, *Proceedings of the 2013 ACM SIGPLAN Symposium on Haskell, Boston, MA, USA, September 23-24, 2013*, pages 59–70. ACM, 2013. URL <https://doi.org/10.1145/2503778.2503791>.

- Donald E. Knuth. *Literate programming*, volume 27 of *CSLI lecture notes series*. Center for the Study of Language and Information, 1992. ISBN 978-0-937073-81-0.
- Hsiang-Shang Ko and Jeremy Gibbons. Programming with ornaments. *J. Funct. Program.*, 27:e2, 2017. URL <https://doi.org/10.1017/S0956796816000307>.
- Saul A Kripke. Semantical analysis of modal logic i normal modal propositional calculi. *Mathematical Logic Quarterly*, 9 (5-6):67–96, 1963.
- Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. Cakeml: a verified implementation of ML. In Suresh Jagannathan and Peter Sewell, editors, *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, pages 179–192. ACM, 2014. URL <https://doi.org/10.1145/2535838.2535841>.
- Daan Leijen. Type directed compilation of row-typed algebraic effects. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 486–499. ACM, 2017. URL <https://doi.org/10.1145/3009837.3009872>.
- Xavier Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, 2009. URL <https://doi.org/10.1145/1538788.1538814>.
- Paul Blain Levy. *Call-By-Push-Value: A Functional/Imperative Synthesis*, volume 2 of *Semantics Structures in Computation*. Springer, 2004. ISBN 1-4020-1730-8.
- Sheng Liang, Paul Hudak, and Mark P. Jones. Monad transformers and modular interpreters. In Ron K. Cytron and Peter Lee, editors, *Conference Record of POPL'95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming*

- Languages, San Francisco, California, USA, January 23-25, 1995*, pages 333–343. ACM Press, 1995a. URL <https://doi.org/10.1145/199448.199528>.
- Sheng Liang, Paul Hudak, and Mark P. Jones. Monad transformers and modular interpreters. In Ron K. Cytron and Peter Lee, editors, *Conference Record of POPL'95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*, pages 333–343. ACM Press, 1995b. ISBN 0-89791-692-1. URL <https://doi.org/10.1145/199448.199528>.
- Sam Lindley and James Cheney. Row-based effect types for database integration. In Benjamin C. Pierce, editor, *Proceedings of TLDI 2012: The Seventh ACM SIGPLAN Workshop on Types in Languages Design and Implementation, Philadelphia, PA, USA, Saturday, January 28, 2012*, pages 91–102. ACM, 2012. URL <https://doi.org/10.1145/2103786.2103798>.
- Saunders MacLane. *Categories for the Working Mathematician*. Springer-Verlag, New York, 1971. Graduate Texts in Mathematics, Vol. 5.
- Ken Madlener, Sjaak Smetsers, and Marko C. J. D. van Eekelen. Formal component-based semantics. In Michel A. Reniers and Pawel Sobocinski, editors, *Proceedings Eight Workshop on Structural Operational Semantics 2011, SOS 2011, Aachen, Germany, 5th September 2011*, volume 62 of *EPTCS*, pages 17–29, 2011. URL <https://doi.org/10.4204/EPTCS.62.2>.
- Per Martin-Löf. *Intuitionistic type theory*, volume 1 of *Studies in proof theory*. Bibliopolis, 1984. ISBN 978-88-7088-228-5.
- Conor McBride. Ornamental algebras, algebraic ornaments. Unpublished manuscript, 2011.
- Lambert G. L. T. Meertens. Paramorphisms. *Formal Aspects Comput.*, 4(5):413–424, 1992. URL <https://doi.org/10.1007/BF01211391>.

- Erik Meijer, Maarten M. Fokkinga, and Ross Paterson. Functional programming with bananas, lenses, envelopes and barbed wire. In John Hughes, editor, *Functional Programming Languages and Computer Architecture, 5th ACM Conference, Cambridge, MA, USA, August 26-30, 1991, Proceedings*, volume 523 of *Lecture Notes in Computer Science*, pages 124–144. Springer, 1991. URL https://doi.org/10.1007/3540543961_7.
- Robin Milner. A theory of type polymorphism in programming. *J. Comput. Syst. Sci.*, 17(3):348–375, 1978. URL [https://doi.org/10.1016/0022-0000\(78\)90014-4](https://doi.org/10.1016/0022-0000(78)90014-4).
- Eugenio Moggi. Computational lambda-calculus and monads. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science (LICS '89), Pacific Grove, California, USA, June 5-8, 1989*, pages 14–23. IEEE Computer Society, 1989. URL <https://doi.org/10.1109/LICS.1989.39155>.
- Eugenio Moggi. An abstract view of programming languages. Technical Report ECS-LFCS-90-113, LFCS, University of Edinburgh, 1990. URL <http://www.disi.unige.it/person/MoggiE/ftp/abs-view.ps.gz>.
- Eugenio Moggi. Notions of computation and monads. *Inf. Comput.*, 93(1):55–92, 1991. URL [https://doi.org/10.1016/0890-5401\(91\)90052-4](https://doi.org/10.1016/0890-5401(91)90052-4).
- J. Garrett Morris. Variations on variants. In Ben Lippmeier, editor, *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, September 3-4, 2015*, pages 71–81. ACM, 2015. URL <https://doi.org/10.1145/2804302.2804320>.
- J. Garrett Morris and James McKinna. Abstracting extensible data types: or, rows by any other name. *Proc. ACM Program. Lang.*, 3(POPL):12:1–12:28, 2019a. URL <https://doi.org/10.1145/3290325>.
- J. Garrett Morris and James McKinna. Abstracting extensible data types: or, rows by any other name. *Proc. ACM Program.*

Lang., 3(POPL):12:1–12:28, 2019b. URL <https://doi.org/10.1145/3290325>.

Peter D. Mosses. Modular structural operational semantics. *J. Log. Algebraic Methods Program.*, 60-61:195–228, 2004. URL <https://doi.org/10.1016/j.jlap.2004.03.008>.

Alan Mycroft, Dominic A. Orchard, and Tomas Petricek. Effect systems revisited - control-flow algebra and semantics. In Christian W. Probst, Chris Hankin, and René Rydhof Hansen, editors, *Semantics, Logics, and Calculi - Essays Dedicated to Hanne Riis Nielson and Flemming Nielson on the Occasion of Their 60th Birthdays*, volume 9560 of *Lecture Notes in Computer Science*, pages 1–32. Springer, 2016. ISBN 978-3-319-27809-4. URL https://doi.org/10.1007/978-3-319-27810-0_1.

Flemming Nielson and Hanne Riis Nielson. Type and effect systems. In Ernst-Rüdiger Olderog and Bernhard Steffen, editors, *Correct System Design, Recent Insight and Advances, (to Hans Langmaack on the occasion of his retirement from his professorship at the University of Kiel)*, volume 1710 of *Lecture Notes in Computer Science*, pages 114–136. Springer, 1999. URL https://doi.org/10.1007/3-540-48092-7_6.

Ulf Norell. Dependently typed programming in agda. In Pieter W. M. Koopman, Rinus Plasmeijer, and S. Doaitse Swierstra, editors, *Advanced Functional Programming, 6th International School, AFP 2008, Heijen, The Netherlands, May 2008, Revised Lectures*, volume 5832 of *Lecture Notes in Computer Science*, pages 230–266. Springer, 2008. ISBN 978-3-642-04651-3. URL https://doi.org/10.1007/978-3-642-04652-0_5.

Ulf Norell. Dependently typed programming in agda. In Andrew Kennedy and Amal Ahmed, editors, *Proceedings of TLDI'09: 2009 ACM SIGPLAN International Workshop on Types in Languages Design and Implementation, Savannah, GA, USA, January 24, 2009*, pages 1–2. ACM, 2009. URL <https://doi.org/10.1145/1481861.1481862>.

- Scott Owens, Magnus O. Myreen, Ramana Kumar, and Yong Kiam Tan. Functional big-step semantics. In Peter Thiemann, editor, *Programming Languages and Systems - 25th European Symposium on Programming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*, volume 9632 of *Lecture Notes in Computer Science*, pages 589–615. Springer, 2016. URL https://doi.org/10.1007/978-3-662-49498-1_23.
- Lionel Parreaux, Aleksander Boruch-Gruszecki, and Paolo G. Giarrusso. Towards improved GADT reasoning in scala. In Jonathan Immanuel Brachthäuser, Sukyoung Ryu, and Nathaniel Nystrom, editors, *Proceedings of the Tenth ACM SIGPLAN Symposium on Scala, Scala@ECOOP 2019, London, UK, July 17, 2019*, pages 12–16. ACM, 2019. URL <https://doi.org/10.1145/3337932.3338813>.
- Frank Pfenning and Conal Elliott. Higher-order abstract syntax. In Richard L. Wexelblat, editor, *Proceedings of the ACM SIGPLAN'88 Conference on Programming Language Design and Implementation (PLDI), Atlanta, Georgia, USA, June 22-24, 1988*, pages 199–208. ACM, 1988. URL <https://doi.org/10.1145/53990.54010>.
- Benjamin C. Pierce. *Basic category theory for computer scientists*. Foundations of computing. MIT Press, 1991. ISBN 978-0-262-66071-6.
- Benjamin C. Pierce. *Types and programming languages*. MIT Press, 2002. ISBN 978-0-262-16209-8.
- Maciej Piróg and Jeremy Gibbons. The coinductive resumption monad. In Bart Jacobs, Alexandra Silva, and Sam Staton, editors, *Proceedings of the 30th Conference on the Mathematical Foundations of Programming Semantics, MFPS 2014, Ithaca, NY, USA, June 12-15, 2014*, volume 308 of *Electronic Notes in Theoretical Computer Science*, pages 273–288. Elsevier, 2014. URL <https://doi.org/10.1016/j.entcs.2014.10.015>.

- Maciej Piróg, Tom Schrijvers, Nicolas Wu, and Mauro Jaskielioff. Syntax and semantics for operations with scopes. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 809–818. ACM, 2018. URL <https://doi.org/10.1145/3209108.3209166>.
- Gordon D. Plotkin. A structural approach to operational semantics. *J. Log. Algebraic Methods Program.*, 60-61:17–139, 2004.
- Gordon D. Plotkin and John Power. Notions of computation determine monads. In Mogens Nielsen and Uffe Engberg, editors, *Foundations of Software Science and Computation Structures, 5th International Conference, FOSSACS 2002. Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 8-12, 2002, Proceedings*, volume 2303 of *Lecture Notes in Computer Science*, pages 342–356. Springer, 2002. ISBN 3-540-43366-X. URL https://doi.org/10.1007/3-540-45931-6_24.
- Gordon D. Plotkin and John Power. Algebraic operations and generic effects. *Appl. Categorical Struct.*, 11(1):69–94, 2003. URL <https://doi.org/10.1023/A:1023064908962>.
- Gordon D. Plotkin and Matija Pretnar. Handlers of algebraic effects. In Giuseppe Castagna, editor, *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, volume 5502 of *Lecture Notes in Computer Science*, pages 80–94. Springer, 2009a. URL https://doi.org/10.1007/978-3-642-00590-9_7.
- Gordon D. Plotkin and Matija Pretnar. Handlers of algebraic effects. In Giuseppe Castagna, editor, *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, volume 5502 of *Lecture Notes in Computer*

- Science*, pages 80–94. Springer, 2009b. ISBN 978-3-642-00589-3. URL https://doi.org/10.1007/978-3-642-00590-9_7.
- Matija Pretnar. Inferring algebraic effects. *Log. Methods Comput. Sci.*, 10(3), 2014. URL [https://doi.org/10.2168/LMCS-10\(3:21\)2014](https://doi.org/10.2168/LMCS-10(3:21)2014).
- Matija Pretnar. An introduction to algebraic effects and handlers. invited tutorial paper. In Dan R. Ghica, editor, *The 31st Conference on the Mathematical Foundations of Programming Semantics, MFPS 2015, Nijmegen, The Netherlands, June 22-25, 2015*, volume 319 of *Electronic Notes in Theoretical Computer Science*, pages 19–35. Elsevier, 2015. URL <https://doi.org/10.1016/j.entcs.2015.12.003>.
- Didier Rémy. Typechecking records and variants in a natural extension of ML. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 11-13, 1989*, pages 77–88. ACM Press, 1989. URL <https://doi.org/10.1145/75277.75284>.
- Cas van der Rest and Casper Bach. Towards a language for defining reusable programming language components - (project paper). In Wouter Swierstra and Nicolas Wu, editors, *Trends in Functional Programming - 23rd International Symposium, TFP 2022, Virtual Event, March 17-18, 2022, Revised Selected Papers*, volume 13401 of *Lecture Notes in Computer Science*, pages 18–38. Springer, 2022. URL https://doi.org/10.1007/978-3-031-21314-4_2.
- Cas van der Rest and Casper Bach. Types and semantics for extensible data types. In Chung-Kil Hur, editor, *Programming Languages and Systems - 21st Asian Symposium, APLAS 2023, Taipei, Taiwan, November 26-29, 2023, Proceedings*, volume 14405 of *Lecture Notes in Computer Science*, pages 46–66. Springer, 2023a. URL https://doi.org/10.1007/978-981-99-8311-7_3.

- Cas van der Rest and Casper Bach. Types and semantics for extensible data types (extended version). *CoRR*, abs/2309.14985, 2023b. URL <https://doi.org/10.48550/arXiv.2309.14985>.
- Cas van der Rest and Casper Bach. GitHub - heft-lang/hefty-equations: Modular reasoning about (elaborations of) higher-order effects — github.com. <https://github.com/heft-lang/hefty-equations>, 2024.
- Cas van der Rest and Wouter Swierstra. A completely unique account of enumeration. *Proc. ACM Program. Lang.*, 6(ICFP):411–437, 2022. URL <https://doi.org/10.1145/3547636>.
- Cas van der Rest, Casper Bach, Arjen Rouvoet, Eelco Visser, and Peter Mosses. Intrinsically-typed definitional interpreters à la carte (artifact), 2022a. URL <https://doi.org/10.5281/zenodo.7074690>.
- Cas van der Rest, Casper Bach, Arjen Rouvoet, Eelco Visser, and Peter D. Mosses. Intrinsically-typed definitional interpreters à la carte. *Proc. ACM Program. Lang.*, 6(OOPSLA2):1903–1932, 2022b. URL <https://doi.org/10.1145/3563355>.
- John C. Reynolds. Definitional interpreters for higher-order programming languages. *High. Order Symb. Comput.*, 11(4):363–397, 1998. URL <https://doi.org/10.1023/A:1010027404223>.
- Grigore Rosu and Traian-Florin Serbanuta. An overview of the K semantic framework. *J. Log. Algebraic Methods Program.*, 79(6):397–434, 2010. URL <https://doi.org/10.1016/j.jlap.2010.03.012>.
- Arjen Rouvoet, Casper Bach, Robbert Krebbers, and Eelco Visser. Intrinsically-typed definitional interpreters for linear, session-typed languages. In Jasmin Blanchette and Catalin Hritcu, editors, *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2020, New Orleans, LA, USA, January 20–21, 2020*, pages 284–298. ACM, 2020. URL <https://doi.org/10.1145/3372885.3373818>.

- Arjen Rouvoet, Robbert Krebbers, and Eelco Visser. Intrinsically typed compilation with nameless labels. *Proc. ACM Program. Lang.*, 5(POPL):1–28, 2021. doi: 10.1145/3434303. URL <https://doi.org/10.1145/3434303>.
- David Schmidt. *Denotational Semantics*. Allyn and Bacon, 1986.
- Tom Schrijvers, Nicolas Wu, Benoit Desouter, and Bart Demoen. Heuristics entwined with handlers combined: From functional specification to logic programming implementation. In *Proceedings of the 16th International Symposium on Principles and Practice of Declarative Programming, Kent, Canterbury, United Kingdom, September 8-10, 2014*, pages 259–270, 2014. URL <https://doi.org/10.1145/2643135.2643145>.
- Tom Schrijvers, Maciej Piróg, Nicolas Wu, and Mauro Jaskielioff. Monad transformers and modular algebraic effects: what binds them together. In Richard A. Eisenberg, editor, *Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2019, Berlin, Germany, August 18-23, 2019*, pages 98–113. ACM, 2019. URL <https://doi.org/10.1145/3331545.3342595>.
- Christopher Schwaab and Jeremy G. Siek. Modular type-safety proofs in agda. In Matthew Might, David Van Horn, Andreas Abel, and Tim Sheard, editors, *Proceedings of the 7th Workshop on Programming languages meets program verification, PLPV 2013, Rome, Italy, January 22, 2013*, pages 3–12. ACM, 2013. URL <https://doi.org/10.1145/2428116.2428120>.
- Dana Scott and Christopher Strachey. Towards a mathematical semantics for computer languages. *Proceedings of the Symposium on Computers and Automata*, 21, 01 1971.
- Neil Sculthorpe, Paolo Torrini, and Peter D. Mosses. A modular structural operational semantics for delimited continuations. In Olivier Danvy and Ugo de'Liguoro, editors, *Proceedings of the Workshop on Continuations, WoC 2016, London, UK, April*

- 12th 2015, volume 212 of *EPTCS*, pages 63–80, 2015. URL <https://doi.org/10.4204/EPTCS.212.5>.
- Christopher Strachey. Towards a formal semantics. 1966.
- Christopher S. Strachey. Fundamental concepts in programming languages. *High. Order Symb. Comput.*, 13(1/2):11–49, 2000. URL <https://doi.org/10.1023/A:1010000313106>.
- Wouter Swierstra. Data types à la carte. *J. Funct. Program.*, 18(4):423–436, 2008. URL <https://doi.org/10.1017/S0956796808006758>.
- Walid Taha and Tim Sheard. Metaml and multi-stage programming with explicit annotations. *Theor. Comput. Sci.*, 248(1-2):211–242, 2000. URL [https://doi.org/10.1016/S0304-3975\(00\)00053-0](https://doi.org/10.1016/S0304-3975(00)00053-0).
- Hayo Thielecke. *Categorical Structure of Continuation Passing Style*. PhD thesis, University of Edinburgh, 1997.
- Paolo Torrini and Tom Schrijvers. Reasoning about modular datatypes with mendler induction. In Ralph Matthes and Matteo Mio, editors, *Proceedings Tenth International Workshop on Fixed Points in Computer Science, FICS 2015, Berlin, Germany, September 11-12, 2015*, volume 191 of *EPTCS*, pages 143–157, 2015. URL <https://doi.org/10.4204/EPTCS.191.13>.
- Tarmo Uustalu and Varmo Vene. Mendler-style inductive types, categorically. *Nord. J. Comput.*, 6(3):343, 1999.
- Cas van der Rest and Casper Bach. Hefty algebras: Modular elaboration of higher-order effects. *J. Funct. Program.*, 35, 2025. doi: 10.1017/S0956796825100142. URL <https://doi.org/10.1017/s0956796825100142>.
- Phil Wadler. The expression problem. <http://homepages.inf.ed.ac.uk/wadler/papers/expression/expression.txt>, 1998. Accessed: 2020-07-01.

- Philip Wadler. The essence of functional programming. In Ravi Sethi, editor, *Conference Record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Albuquerque, New Mexico, USA, January 19-22, 1992*, pages 1–14. ACM Press, 1992. URL <https://doi.org/10.1145/143165.143169>.
- Philip Wadler, Wen Kokke, and Jeremy G. Siek. *Programming Language Foundations in Agda*. July 2020. URL <http://plfa.inf.ed.ac.uk/20.07/>.
- Mitchell Wand. Final algebra semantics and data type extensions. *J. Comput. Syst. Sci.*, 19(1):27–44, 1979. URL [https://doi.org/10.1016/0022-0000\(79\)90011-4](https://doi.org/10.1016/0022-0000(79)90011-4).
- Mitchell Wand. Type inference for record concatenation and multiple inheritance. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science (LICS '89), Pacific Grove, California, USA, June 5-8, 1989*, pages 92–97. IEEE Computer Society, 1989. URL <https://doi.org/10.1109/LICS.1989.39162>.
- Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Inf. Comput.*, 115(1):38–94, 1994. URL <https://doi.org/10.1006/inco.1994.1093>.
- Nicolas Wu, Tom Schrijvers, and Ralf Hinze. Effect handlers in scope. In Wouter Swierstra, editor, *Proceedings of the 2014 ACM SIGPLAN symposium on Haskell, Gothenburg, Sweden, September 4-5, 2014*, pages 1–12. ACM, 2014. URL <https://doi.org/10.1145/2633357.2633358>.
- Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. Interaction trees: representing recursive and impure programs in coq. *Proc. ACM Program. Lang.*, 4(POPL):51:1–51:32, 2020. URL <https://doi.org/10.1145/3371119>.
- Zhixuan Yang and Nicolas Wu. Reasoning about effect interaction by fusion. *Proc. ACM Program. Lang.*, 5(ICFP):1–29, 2021. URL <https://doi.org/10.1145/3473578>.

- Zhixuan Yang, Marco Paviotti, Nicolas Wu, Birthe van den Berg, and Tom Schrijvers. Structured handling of scoped effects. In Ilya Sergey, editor, *Programming Languages and Systems - 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings*, volume 13240 of *Lecture Notes in Computer Science*, pages 462–491. Springer, 2022. URL https://doi.org/10.1007/978-3-030-99336-8_17.
- Yannick Zakowski, Calvin Beck, Irene Yoon, Ilia Zaichuk, Vadim Zaliva, and Steve Zdancewic. Modular, compositional, and executable formal semantics for LLVM IR. *Proc. ACM Program. Lang.*, 5(ICFP):1–30, 2021. URL <https://doi.org/10.1145/3473572>.
- Weixin Zhang, Yaozhu Sun, and Bruno C. d. S. Oliveira. Compositional programming. *ACM Trans. Program. Lang. Syst.*, 43(3):9:1–9:61, 2021. URL <https://doi.org/10.1145/3460228>.
- Yizhou Zhang and Andrew C. Myers. Abstraction-safe effect handlers via tunneling. *Proc. ACM Program. Lang.*, 3(POPL):5:1–5:29, 2019. URL <https://doi.org/10.1145/3290318>.

PUBLICATIONS

- Casper Bach, Cas van der Rest, and Tom Schrijvers. Staged effects and handlers for modular languages with abstraction. In *Workshop on Partial Evaluation and Program Manipulation (PEPM)*. <http://casperbp.net/store/staged-effects-and-handlers.pdf>, 2021
- Cas van der Rest and Casper Bach. Towards a language for defining reusable programming language components - (project paper). In Wouter Swierstra and Nicolas Wu, editors, *Trends in Functional Programming - 23rd International Symposium, TFP 2022, Virtual Event, March 17-18, 2022, Revised Selected Papers*, volume 13401 of *Lecture Notes in Computer Science*, pages 18–38. Springer, 2022. URL https://doi.org/10.1007/978-3-031-21314-4_2
- Cas van der Rest and Wouter Swierstra. A completely unique account of enumeration. *Proc. ACM Program. Lang.*, 6(ICFP):411–437, 2022. URL <https://doi.org/10.1145/3547636>
- Cas van der Rest, Casper Bach, Arjen Rouvoet, Eelco Visser, and Peter D. Mosses. Intrinsically-typed definitional interpreters à la carte. *Proc. ACM Program. Lang.*, 6(OOPSLA2):1903–1932, 2022b. URL <https://doi.org/10.1145/3563355>
- Casper Bach and Cas van der Rest. Hefty algebras: Modular elaboration of higher-order algebraic effects. *Proc. ACM Program. Lang.*, 7(POPL):1801–1831, 2023. URL <https://doi.org/10.1145/3571255>

- Cas van der Rest and Casper Bach. Types and semantics for extensible data types. In Chung-Kil Hur, editor, *Programming Languages and Systems - 21st Asian Symposium, APLAS 2023, Taipei, Taiwan, November 26-29, 2023, Proceedings*, volume 14405 of *Lecture Notes in Computer Science*, pages 46–66. Springer, 2023a. URL https://doi.org/10.1007/978-981-99-8311-7_3
- Cas van der Rest and Casper Bach. Hefty algebras: Modular elaboration of higher-order effects. *J. Funct. Program.*, 35, 2025. doi: 10.1017/S0956796825100142. URL <https://doi.org/10.1017/s0956796825100142>

SUMMARY

Type systems are a tool for preventing software errors, by classifying (sub)terms according to how they are evaluated. This way, mistakes can be caught at *compile-time*, ruling out the existence of entire classes of mistakes altogether. Using a programming language with a strong type system to develop critical software can dramatically reduce the prevalence and impact of bugs.

In light of the potentially enormous impact of bugs, it is important that we can trust a type system to be successful in preventing errors. A key property of type systems that reflects this criterium is *type soundness*, which establishes that “well-typed programs cannot go wrong”. That is, programs that are deemed safe by the type system should not exhibit certain wrong behaviour when executed. To gain trust in a type system’s ability to prevent errors, we can give a formal system of both a language’s type system and *semantics*, and mathematically verify that the type system is sound with respect to the defined semantics. While this provides airtight evidence that a type system is successful in ruling out certain mistakes, the formal specification and verification of programming languages requires a formidable amount of time and expertise on behalf of the language designer, and is therefore infeasible in most cases.

This thesis aims to cut down on the cost associated with the formal specification and verification of programming languages, through reuse of (parts of) existing specifications and proofs. Specifically, we investigate how to build *reusable programming language components*. Underlying this research is the belief that the effort expended to specify and verify a new programming language should focus on the novel features of the language, and that we should benefit as much as possible from pre-existing verification efforts for other parts of the language. To realize this vision, it is necessary to make progress on two fronts. First, we

must understand how the semantics of programming languages can be formally specified and verified in a modular way. An important challenge in defining a language's semantics modularly is to give a modular description of its *side effects*. While *algebraic effects and handlers* provide a solution to this problem for many familiar side effects, *higher-order* effects cannot be described using the approach. Second, the *meta language* used to develop specifications should feature adequate support for reuse in order to eliminate syntactic overhead incurred by modularity as much as possible.

Part I of this thesis contributes several semantic techniques for defining *modular intrinsically-typed definitional interpreters* in a dependently-typed host language, such as Agda. More specifically, in Chapter 2 we explain how intrinsically-typed definitional interpreters can be defined modularly, such that the composition of interpreters preserves the encoded soundness invariant. Then, in Chapter 3, we explore how to define and reason about higher-order effects in a modular way. A modular treatment of higher-order effects is crucial for defining the semantics of many widely-used language constructs in a modular way, such as λ -abstraction or exception catching.

In part II of this thesis we investigate the design of *meta languages* for the purpose of defining reusable programming language components. In Chapter 4, we discuss the design of a language that features built-in support for modular algebraic data types and higher-order effects. This combination of features allows the definition of reusable language components, while incurring minimal syntactic overhead resulting from modularity of the definition. Chapter 5 puts this design on a stronger formal foundation, by developing a typing discipline for modular data types, together with a categorical semantics, illustrating how this feature can be incorporated in the design of meta-languages in a principled manner.

The contributions of this thesis addresses several important open challenges for developing reusable programming language components, such as how to preserve type safety invariants

when composing intrinsically-typed interpreters, modular definitions and reasoning for higher-order effects, and how to integrate modularity in the design of functional meta languages. However, before reuse becomes a feasible approach for large-scale language projects, there is still work to be done. Most importantly, we still lack a uniform treatment of intrinsically-typed side effects, such as polymorphic references, and the support for modularity in meta languages would have to be extended to include intrinsically-typed definitional interpreters as well.

SAMENVATTING

Typesystemen zijn een middel om softwarefouten te voorkomen, door (sub)termen te classificeren aan de hand van hoe ze geëvalueerd worden. Op deze manier kunnen fouten gesignaleerd worden *tijdens het compileren*, waarmee bepaalde soorten fouten in het geheel uitgesloten kunnen worden. Door een programmeertaal te gebruiken die uitgerust is met een sterk typesysteem voor het ontwikkelen van kritieke software, kan de frequentie en invloed van bugs dramatisch gereduceerd worden.

In het licht van de potentieel enorme invloed van bugs is het belangrijk dat we erop kunnen vertrouwen dat een typesysteem er inderdaad in slaagt om fouten te voorkomen. Een sleuteleigenschap van typesystemen die dit criterium vangt is *typejuistheid*, wat vaststelt dat “correct getypeerde programma’s niet fout kunnen gaan”. Dat wil zeggen, programma’s die als veilig beoordeeld worden door het typesysteem vertonen inderdaad bepaald verkeerd gedrag niet wanneer ze uitgevoerd worden. Het vertrouwen in het vermogen van een typesysteem om fouten te voorkomen kan verhoogd worden door een formele specificatie te geven van zowel het typesysteem alsmede de *semantiek* van de programmeertaal in kwestie, en vervolgens juistheid van het typesysteem ten opzichte van de gedefiniëerde semantiek wiskundig te verifiëren. Hoewel dit ons waterdicht bewijs verschaft dat een typesysteem bepaalde fouten uitsluit, is het in de meeste gevallen onuitvoerbaar gezien de uitzonderlijke hoeveelheid tijd en expertise die nodig is voor de formele specificatie en verificatie van programmeertalen.

Dit proefschrift heeft als doel de kosten van het formeel specificeren en verifiëren van programmeertalen te beperken middels hergebruik van (delen van) bestaande specificaties en bewijzen. Meer specifiek onderzoeken we hoe *herbruikbare programmeertaalonderdelen* gebouwd kunnen worden. Ten grondslag aan dit

onderzoek ligt de overtuiging dat de inspanning die verricht wordt om een nieuwe programmeertaal te specificeren en verifiëren zich zou moeten richten op de vernieuwende onderdelen van de taal, en dat voor andere delen van de taal zoveel mogelijk geprofiteerd zou moeten worden van bestaande bewijsinspanningen. Om deze visie te realiseren is het noodzakelijk dat we progressie boeken op twee fronten. Ten eerste dienen we te begrijpen hoe de semantiek van programmeertalen op een modulaire manier gespecificeerd en geverifieerd kan worden. Een belangrijke uitdaging als het aankomt op het modulair specificeren van de semantiek van een taal, is om haar *neveneffecten* op een modulaire manier te beschrijven. Hoewel *algebraïsche effecten en handlers* het modulariteitsprobleem oplossen voor veel mogelijke neveneffecten, kunnen *effecten van een hogere orde* niet met deze aanpak beschreven worden. Ten tweede zou de *meta-taal* die gebruikt wordt om specificaties in te ontwikkelen uitgerust moeten zijn met adequate ondersteuning voor hergebruik zodat de syntactische kosten ten gevolge van modulariteit tot een minimum beperkt worden.

In deel I van dit proefschrift dragen we verscheidene semantische technieken bij om modulaire *intrinsiek getypeerde definitie-interpretators* in een afhankelijk-getypeerde gasttaal, zoals Agda, te kunnen definiëren. Meer specifiek bespreken we in hoofdstuk 2 hoe intrinsiek getypeerde definitie-interpretators op modulaire manier gedefinieerd kunnen worden, zodanig dat het samenvoegen van interpretators hun typejuistheid behoudt. Vervolgens onderzoeken we in hoofdstuk 3 het op modulaire wijze beschrijven van en redeneren over effecten van een hogere orde. Het vermogen om effecten van een hogere orde op een modulaire manier te kunnen beschrijven is cruciaal voor het definiëren van veelvoorkomende taalonderdelen, zoals λ -abstractie of functionaliteit voor het afvangen van excepties.

In deel II van dit proefschrift onderzoeken we het ontwerp van *meta-talen* die als doel hebben om herbruikbare programmeertaalonderdelen te beschrijven. In hoofdstuk 4 bespreken we het ontwerp van een taal die ingebouwde ondersteuning

voor modulaire algebraïsche datatypes en effecten van een hogere orde bevat. Deze combinatie van functionaliteiten stelt ons in staat om herbruikbare programmeertaalonderdelen te definiëren terwijl de syntactische kosten als gevolg van modulariteit van de definities minimaal blijven. Hoofdstuk 5 verschaft dit taalontwerp van een sterkere formele fundering door een typesysteem te ontwikkelen voor modulaire datatypes in combinatie met een categorische semantiek, waarmee we illustreren hoe dergelijke functionaliteit op een principiële manier opgenomen kan worden in het ontwerp van een meta-taal.

De bijdragen in dit proefschrift geven antwoord op een aantal belangrijke open uitdagingen voor het ontwikkelen van herbruikbare programmeertaalonderdelen, zoals het behouden van de typejuistheidsinvariant wanneer intrinsiek getypeerde definitie-interpretators worden samengevoegd, het modulair definiëren en redeneren over effecten van een hogere orde, en hoe modulariteit in het ontwerp van functionele meta-talen geïntegreerd kan worden. Desondanks is er voldoende werk te verzetten voordat hergebruik een uitvoerbare strategie kan worden voor grootschalige taalprojecten. In het bijzonder missen we nog een uniforme beschrijving van intrinsiek getypeerde neveneffecten, zoals polymorfe referenties, en de ondersteuning van modulariteit in meta-talen zal moeten worden uitgebreid naar intrinsiek getypeerde definitie-interpretators.

CURRICULUM VITAE

1995

Born on June 1st in Rotterdam, The Netherlands.

2007 - 2013

High School (VWO/gymnasium), Marnix Gymnasium, Rotterdam, The Netherlands.

2013 - 2017

Bachelor Computer Science, Utrecht University, Utrecht, The Netherlands

2017 - 2019

Master Computer Science, Utrecht University, Utrecht, The Netherlands. Thesis title: "Generating Constrained Test Data Using Datatype Generic Programming", supervised by dr. Wouter Swierstra. Graduated with *cum laude* distinction.

2019 - 2024

Doctoral candidate in Computer Science, Delft University of Technology, Delft, The Netherlands. Thesis title: "Reusable Programming Language Components", supervised by dr. Casper Bach.

2024 - present

Formal Methods Engineer at Input Output Global/Shielded Technologies.

Titles in the IPA Dissertation Series since 2023

S. Kochanthara. *A Changing Landscape: On Safety & Open Source in Automated and Connected Driving.* Faculty of Mathematics and Computer Science, TU/e. 2023-01

L.M. Ochoa Venegas. *Break the Code? Breaking Changes and Their Impact on Software Evolution.* Faculty of Mathematics and Computer Science, TU/e. 2023-02

N. Yang. *Logs and models in engineering complex embedded production software systems.* Faculty of Mathematics and Computer Science, TU/e. 2023-03

J. Cao. *An Independent Timing Analysis for Credit-Based Shaping in Ethernet TSN.* Faculty of Mathematics and Computer Science, TU/e. 2023-04

K. Dokter. *Scheduled Protocol Programming.* Faculty of Mathematics and Natural Sciences, UL. 2023-05

J. Smits. *Strategic Language Workbench Improvements.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2023-06

A. Arslanagić. *Minimal Structures for Program Analysis and Verification.* Faculty of Science and Engineering, RUG. 2023-07

M.S. Bouwman. *Supporting Railway Standardisation with Formal Verification.* Faculty of Mathematics and Computer Science, TU/e. 2023-08

S.A.M. Lathouwers. *Exploring Annotations for Deductive Verification.* Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2023-09

J.H. Stoel. *Solving the Bank, Lightweight Specification and Verification Techniques for Enterprise Software.* Faculty of Mathematics and Computer Science, TU/e. 2023-10

D.M. Groenewegen. *WebDSL: Linguistic Abstractions for Web Programming.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2023-11

D.R. do Vale. *On Semantical Methods for Higher-Order Complexity Analysis.* Faculty of Sci-

ence, Mathematics and Computer Science, RU. 2024-01

M.J.G. Olsthoorn. *More Effective Test Case Generation with Multiple Tribes of AI.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-02

B. van den Heuvel. *Correctly Communicating Software: Distributed, Asynchronous, and Beyond.* Faculty of Science and Engineering, RUG. 2024-03

H.A. Hiep. *New Foundations for Separation Logic.* Faculty of Mathematics and Natural Sciences, UL. 2024-04

C.E. Brandt. *Test Amplification For and With Developers.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-05

J.I. Hejderup. *Fine-Grained Analysis of Software Supply Chains.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-06

J. Jacobs. *Guarantees by construction.* Faculty of Science, Mathematics and Computer Science, RU. 2024-07

O. Bunte. *Cracking OIL: A Formal Perspective on an Indus-*

trial DSL for Modelling Control Software. Faculty of Mathematics and Computer Science, TU/e. 2024-08

R.J.A. Erkens. *Automaton-based Techniques for Optimized Term Rewriting.* Faculty of Mathematics and Computer Science, TU/e. 2024-09

J.J.M. Martens. *The Complexity of Bisimilarity by Partition Refinement.* Faculty of Mathematics and Computer Science, TU/e. 2024-10

L.J. Edixhoven. *Expressive Specification and Verification of Choreographies.* Faculty of Science, OU. 2024-11

J.W.N. Paulus. *On the Expressivity of Typed Concurrent Calculi.* Faculty of Science and Engineering, RUG. 2024-12

J. Denkers. *Domain-Specific Languages for Digital Printing Systems.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-13

L.H. Applis. *Tool-Driven Quality Assurance for Functional Programming and Machine Learning.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2024-14

- P. Karkhanis.** *Driving the Future: Facilitating C-ITS Service Deployment for Connected and Smart Roadways.* Faculty of Mathematics and Computer Science, TU/e. 2024-15
- N.W. Cassee.** *Sentiment in Software Engineering.* Faculty of Mathematics and Computer Science, TU/e. 2024-16
- H. van Antwerpen.** *Declarative Name Binding for Type System Specifications.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2025-01
- I.N. Mulder.** *Proof Automation for Fine-Grained Concurrent Separation Logic.* Faculty of Science, Mathematics and Computer Science, RU. 2025-02
- T.S. Badings.** *Robust Verification of Stochastic Systems: Guarantees in the Presence of Uncertainty.* Faculty of Science, Mathematics and Computer Science, RU. 2025-03
- A.M. Mir.** *Machine Learning-assisted Software Analysis.* Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2025-04
- L.T. Vinkhuijzen.** *Data Structures for Quantum Circuit Verification and How To Compare Them.* Faculty of Mathematics and Natural Sciences, UL. 2025-05
- D. van der Wal.** *What is the Point? Single-Input-Change Testing a EULYNX Controller.* Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2025-06
- A. Rosset.** *Uniform Monad Presentations and Graph Quasitoposes.* Faculty of Sciences, Department of Computer Science, VU. 2025-07
- L. Guo.** *Higher-Order Termination with Logical Constraints.* Faculty of Science, Mathematics and Computer Science, RU. 2025-08
- D.A. Manrique Negrin.** *A Model Orchestra in Digital Twins: A Model-Driven Approach to Integration and Orchestration.* Faculty of Mathematics and Computer Science, TU/e. 2025-09
- C.A. Esterhuysen.** *Specification-Centric Multi-Agent Systems.* Faculty of Science, UvA. 2025-10
- H.M. Muctadir.** *Consistency Matters: Building Consistent Digital Twin Virtual Enti-*

ties. Faculty of Mathematics and Computer Science, TU/e. 2025-11

A. Stramaglia. *Model Checking Machine-Control Applications*. Faculty of Mathematics and Computer Science, TU/e. 2025-12

M. Saeedi Nikoo. *Supporting business process management: clone detection and recommendation techniques*. Faculty of Mathematics and Computer Science, TU/e. 2025-13

R.B. Rubbens. *Bridging the Implementation Gap: Advances in Model-Based Concurrent Program Verification*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2025-14

F.I. van der Berg. *DMC Model Checker: Delta-Driven Variable-Length Next-State Generation via Recursive Compression*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2025-15

D.G. Sprokholt. *Correct Translation between Weak Memory Model Architectures*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2025-16

B. Kohlen. *Verified & Fast Probabilistic Verification: Step-wise Refinement of Probabilistic Model Checking into an IEEE 754 Floating-Point Implementation*. Faculty of Electrical Engineering, Mathematics & Computer Science, UT. 2026-01

D. Frolich. *Exploratory Language Development*. Faculty of Science, UvA. 2026-02

L.B. van den Haak. *Verifying Optimised Parallel Code*. Faculty of Mathematics and Computer Science, TU/e. 2026-03

A. Papotti. *Human Acceptance of Machine Generated Code Security Analysis*. Faculty of Sciences, Department of Computer Science, VU. 2026-04

C.R. van der Rest. *Reusable Programming Language Components*. Faculty of Electrical Engineering, Mathematics, and Computer Science, TUD. 2026-04

N. Doorn. *O bug, where art thou? Improving software testing education*. Faculty of Science, OU. 2026-05

