

Reiko Heckel
Stefan Milius (Eds.)

LNCS 8089

Algebra and Coalgebra in Computer Science

5th International Conference, CALCO 2013
Warsaw, Poland, September 2013
Proceedings



Springer

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

Reiko Heckel Stefan Milius (Eds.)

Algebra and Coalgebra in Computer Science

5th International Conference, CALCO 2013

Warsaw, Poland, September 3-6, 2013

Proceedings



Springer

Volume Editors

Reiko Heckel
University of Leicester, Department of Computer Science
University Road
Leicester, LE1 7RH, UK
E-mail: reiko@mcs.le.ac.uk

Stefan Milius
Friedrich-Alexander Universität Erlangen-Nürnberg
Lehrstuhl für Theoretische Informatik
Martenstr. 3
91058 Erlangen-Nürnberg, Germany
E-mail: stefan.milius@fau.de

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-40205-0 e-ISBN 978-3-642-40206-7
DOI 10.1007/978-3-642-40206-7
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2013944581

CR Subject Classification (1998): F.4, F.3, D.2, F.1, F.2, I.1, G.2, G.4

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

© Springer-Verlag Berlin Heidelberg 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

CALCO, the International Conference on Algebra and Coalgebra in Computer Science, is a high-level, bi-annual event formed by joining CMCS (the International Workshop on Coalgebraic Methods in Computer Science) and WADT (the Workshop on Algebraic Development Techniques). CALCO aims to bring together researchers and practitioners with interests in foundational aspects, and both traditional and emerging uses of algebras and coalgebras in computer science. The study of algebra and coalgebra relates to the data, process, and structural aspects of software systems.

Previous CALCO editions took place in Swansea (UK, 2005), Bergen (Norway, 2007), Udine (Italy, 2009) and Winchester (UK, 2011). CALCO 2013, the fifth conference in the series, took place in Warsaw (Poland), September 3–6, 2013.

CALCO 2013 received 33 submissions, out of which 18 were selected for presentation at the conference. The standard of submissions was generally very high. The selection process was carried out by the Program Committee, taking into account the originality, quality, and relevance of the material presented in each submission, based on the opinions of three or four expert reviewers for each submission. The selected and revised papers are included in this volume, together with contributions by the invited speakers Andrej Bauer, Mikołaj Bojańczyk, Neil Ghani, and Damien Pous.

CALCO 2013 was co-located with two workshops. The CALCO Early Ideas Workshop, CALCO EI, was dedicated to presentation of work in progress and original research proposals. PhD students and young researchers were particularly encouraged to contribute. CALCO EI was organized by Monika Seisenberger. The CALCO Tools Workshop, organized by Lutz Schröder, is dedicated to tools based on algebraic and/or coalgebraic principles. These tool papers also appear in this volume.

We wish to thank all the authors for submitting their papers to CALCO 2013, the Program Committee for its diligent work in the selection process, and the external reviewers for their support in evaluating papers.

We are grateful to the University of Warsaw and the Polish Mathematical Society for hosting CALCO 2013 and to the Organizing Committee, chaired by Bartek Klin and Andrzej Tarlecki, for all the local arrangements. We also thank the Warsaw Centre of Mathematics and Computer Science for their financial support. At Springer, Alfred Hofmann and his team supported the publishing process. We gratefully acknowledge the use of EasyChair, the conference management system by Andrei Voronkov.

Organization

CALCO Steering Committee

Jiří Adámek	Technical University of Braunschweig, Germany
Michel Bidoit	CNRS and ENS de Cachan, France
Corina Cirstea	University of Southampton, UK
Andrea Corradini	University of Pisa, Italy
José Luiz Fiadeiro	Royal Holloway, University of London, UK
H. Peter Gumm (Co-chair)	Philipps University Marburg, Germany
Rolf Hennicker	Ludwig-Maximilians-University Munich, Germany
Bart Jacobs	Radboud University Nijmegen, Germany
Bartek Klin	University of Warsaw, Italy
Hans-Jörg Kreowski	University of Bremen, Germany
Alexander Kurz	University of Leicester, UK
Marina Lenisa	University of Udine, Italy
Ugo Montanari	University of Pisa, Italy
Larry Moss	Indiana University, Bloomington, USA
Till Mossakowski (Co-chair)	DFKI Lab Bremen and University of Bremen, Germany
Fernando Orejas	Politechnical University Catalonia, Barcelona, Spain
Francesco Parisi-Presicce	University of Rome La Sapienza, Italy
Dirk Pattinson	Australian National University, Australia
John Power	University of Bath, UK
Horst Reichel	Technical University of Dresden, Germany
Jan Rutten	CWI Amsterdam and Radboud University Nijmegen, The Netherlands
Lutz Schröder	Friedrich-Alexander University Erlangen-Nürnberg, Germany
Andrzej Tarlecki	University of Warsaw, Poland

Program Committee

Luca Aceto	Reykjavik University, Iceland
Jiří Adámek	Technical University of Braunschweig, Germany
Lars Birkedal	IT University of Copenhagen, Denmark
Filippo Bonchi	CNRS, ENS-Lyon, France

VIII Organization

Corina Cîrstea	University of Southampton, UK
Bob Coecke	University of Oxford, UK
Andrea Corradini	University of Pisa, Italy
Mai Gehrke	Université Paris Diderot – Paris 7, France
H. Peter Gumm	Philipps University Marburg, Germany
Gopal Gupta	University of Texas at Dallas, USA
Ichiro Hasuo	Tokyo University, Japan
Reiko Heckel (Co-chair)	University of Leicester, UK
Bart Jacobs	Radboud University Nijmegen, The Netherlands
Ekaterina Komendantskaya	University of Dundee, Scotland, UK
Barbara König	University of Duisburg-Essen, Germany
José Meseguer	University of Illinois, Urbana-Champaign, USA
Marino Miculan	University of Udine, Italy
Stefan Milius (Co-chair)	Friedrich-Alexander University Erlangen-Nürnberg, Germany
Larry Moss	Indiana University, Bloomington, USA
Till Mossakowski	DFKI Lab Bremen and University of Bremen, Germany
Prakash Panangaden	McGill University, Montreal, Canada
Dirk Pattinson	Australian National University, Australia
Dusko Pavlovic	Royal Holloway, University of London, UK
Daniela Petrişan	University of Leicester, UK
John Power	University of Bath, UK
Jan Rutten	CWI Amsterdam and Radboud University Nijmegen, The Netherlands
Lutz Schröder	Friedrich-Alexander University Erlangen-Nürnberg, Germany
Monika Seisenberger	Swansea University, UK
Alexandra Silva	Radboud University Nijmegen and CWI Amsterdam, The Netherlands
Paweł Sobociński	University of Southampton, UK
Sam Staton	University of Cambridge, UK
Yde Venema	University of Amsterdam, The Netherlands
Uwe Wolter	University of Bergen, Norway

Additional Reviewers

Giorgio Bacci	Mihai Codrescu	Giuseppe Greco
Ulrich Berger	Josée Desharnais	Helle Hvid Hansen
Arnar Birgisson	Brian Devries	Chris Heunen
Aleš Bizjak	Ross Duncan	Tom Hirschowitz
Roberto Bruni	Murdoch Gabbay	Jean-Baptiste Jeannin
Martin Churchill	Rajeev Goré	Henning Kerstan

Aleks Kissinger

Robbert Krebbers

Sebastian Küpper

Clemens Kupke

Alberto Lluch Lafuente

Guy McCusker

Robert Myers

Marco Peressotti

Vaughan Pratt

Joshua Sack

Mehrnoosh Sadrzadeh

Isar Stubbe

Viktor Winschel

Joost Winter

Fabio Zanasi

Table of Contents

Invited Talks

An Effect System for Algebraic Effects and Handlers	1
<i>Andrej Bauer and Matija Pretnar</i>	
Automata and Algebras for Infinite Words and Trees	17
<i>Mikołaj Bojańczyk</i>	
Positive Inductive-Recursive Definitions	19
<i>Neil Ghani, Lorenzo Malatesta, and Fredrik Nordvall Forsberg</i>	
Coalgebraic Up-to Techniques	34
<i>Damien Pous</i>	

Contributed Papers

Exploiting Algebraic Laws to Improve Mechanized Axiomatizations	36
<i>Luca Aceto, Eugen-Ioan Goriac, Anna Ingolfsdottir, Mohammad Reza Mousavi, and Michel A. Reniers</i>	
Positive Fragments of Coalgebraic Logics	51
<i>Adriana Balan, Alexander Kurz, and Jiří Velebil</i>	
Many-Valued Relation Lifting and Moss' Coalgebraic Logic	66
<i>Marta Bílková and Matěj Dostál</i>	
Saturated Semantics for Coalgebraic Logic Programming	80
<i>Filippo Bonchi and Fabio Zanasi</i>	
Presenting Distributive Laws	95
<i>Marcello M. Bonsangue, Helle Hvid Hansen, Alexander Kurz, and Jurriaan Rot</i>	
Interaction and Observation: Categorical Semantics of Reactive Systems Trough Dialgebras	110
<i>Vincenzo Ciancia</i>	
Homomorphisms of Coalgebras from Predicate Liftings	126
<i>Sebastian Enqvist</i>	
From Kleisli Categories to Commutative C^* -Algebras: Probabilistic Gelfand Duality	141
<i>Robert Furber and Bart Jacobs</i>	

Trace Semantics via Generic Observations	158
<i>Sergey Goncharov</i>	
Full Abstraction for Fair Testing in CCS	175
<i>Tom Hirschowitz</i>	
A Simple Case of Rationality of Escalation	191
<i>Pierre Lescanne</i>	
Coalgebras with Symmetries and Modelling Quantum Systems	205
<i>Daniel Marsden</i>	
From Operational Chu Duality to Coalgebraic Quantum Symmetry	220
<i>Yoshihiro Maruyama</i>	
Noninterfering Schedulers: When Possibilistic Noninterference Implies Probabilistic Noninterference	236
<i>Andrei Popescu, Johannes Hölzl, and Tobias Nipkow</i>	
Simulations and Bisimulations for Coalgebraic Modal Logics	253
<i>Daniel Gorín and Lutz Schröder</i>	
A Coalgebraic View of ε -Transitions	267
<i>Alexandra Silva and Bram Westerbaan</i>	
Nets, Relations and Linking Diagrams	282
<i>Paweł Sobociński</i>	
A Logic-Programming Semantics of Services	299
<i>Ionuț Țuțu and José Luiz Fiadeiro</i>	
CALCO-Tools Workshop	
Preface to CALCO-Tools	314
<i>Lutz Schröder</i>	
Checking Conservativity with HETS	315
<i>Mihai Codescu, Till Mossakowski, and Christian Maeder</i>	
The HI-Maude Tool	322
<i>Muhammad Fadlisyah and Peter Csaba Ölveczky</i>	
Constructor-Based Inductive Theorem Prover	328
<i>Daniel Găină, Min Zhang, Yuki Chiba, and Yasuhito Arimoto</i>	
A Timed CTL Model Checker for Real-Time Maude	334
<i>Daniela Lepri, Erika Ábrahám, and Peter Csaba Ölveczky</i>	
Hybridisation at Work	340
<i>Renato Neves, Alexandre Madeira, Manuel A. Martins, and Luís S. Barbosa</i>	

Penrose: Putting Compositionality to Work for Petri Net	
Reachability	346
<i>Paweł Sobociński and Owen Stephens</i>	
QStream: A Suite of Streams	353
<i>Joost Winter</i>	
Author Index	359

An Effect System for Algebraic Effects and Handlers

Andrej Bauer and Matija Pretnar

Faculty of Mathematics and Physics
University of Ljubljana

Abstract. We present an effect system for algebraic effects and handlers. Because handlers may transform an effectful computation into a pure one, the effect system is non-monotone in the sense that effects do not just accumulate, but may also be deleted from types or generally transformed. We also provide denotational semantics for the effect system, based on a domain-theoretic model with partial equivalence relations. The semantics validates equational reasoning about effectful computations.

1 Introduction

An *effect system* supplements a traditional type system for a programming language with information about which computational effects may, will, or will not happen when a piece of code is executed. A well designed and solidly implemented effect system helps the programmer understand the code and find mistakes, but it can also be used to safely rearrange, optimize, and parallelize code [1,2]. As many before us [1,3,4,5] we take on the task of striking just the right balance between simplicity and expressivity by devising an effect system for the programming language Eff [6]. The novelty here is that Eff has not only *first-class algebraic effects* [7], but also *effect handlers* [8]. Therefore, an effect system for Eff is by its nature *non-monotone* — an effectful computation may become pure when enclosed by a handler — so effects do not just accumulate in the types, but also get deleted and generally transformed. Another feature of our effect system is that its denotational semantics validates *equational* reasoning, which is traditionally thought to be in the dominion of pure languages, and can be tricky once effects are included [9].

The paper is organized as follows. In §2 we describe the *core Eff* and an effect system for it. In §3 we give a denotational semantics for core Eff, and use it to validate program transformation rules that can be used for equational reasoning about effectful computations.

2 Core Eff

Eff is a ML-style [10,11] programming language. Effects in ML are not visible in the types. For example, inserting a print statement in the middle of code does

not create any changes in the typing information. In contrast, in the monadic style such a change would taint all enclosing types with the I/O monad [12]. It is our intention to augment the types with information about computational effects which is unobtrusive for programmers, i.e., in the implementation we envision *effect inference* which serves primarily as a program analysis tool.

In ML-style languages effects are provided through built-in functions and special-purpose constructs such as exceptions and references. Eff is based on the *algebraic* approach in which effects are accessed uniformly and exclusively through *operations*, which are a primitive concept. Examples of operations are reading and writing on a communication channel, updating and looking up the contents of a reference, and raising an exception. Thus, in Eff each terminating computation results either in an (effect-free) value, or it triggers an operation. Each operation has an associated (delimited) *continuation*, which is a suspended computation expecting the result of the operation and doing whatever is to be done after the operation is performed.

Operations by themselves do not actually perform effects. Instead their behavior is controlled by a second primitive notion, the *effect handlers*. These are a direct generalization of exception handlers to other operations. The most significant difference between exception handlers and effect handlers is that the latter have access to the continuation of the handled operation. With handlers we may implement a great variety of computational effects, such as transactional memory, various non-deterministic execution strategies, stream redirection, cooperative multi-threading, delimited continuations, and others.

The current implementation of Eff includes a number of features, such as syntactic sugar, products, records, inductive types, type definitions, effect definitions, etc., which are inessential for a conceptual analysis. We therefore restrict attention to *core Eff* whose syntax is shown below.

Types

Effect type $E ::= \text{effect } (\text{operation } \text{op}_i : A_i \rightarrow B_i)_i \text{ end}$
 Expression type $A, B ::= \text{nat} \mid \text{bool} \mid \text{unit} \mid \text{empty} \mid E^{\{\iota_1, \dots, \iota_n\}} \mid$
 $A \rightarrow \underline{C} \mid \underline{C} \Rightarrow \underline{D}$
 Computation type $\underline{C}, \underline{D} ::= A! \{ \iota_1 \# \text{op}_1, \dots, \iota_n \# \text{op}_n \}$

Terms

Expression $e ::= x \mid 0 \mid \text{succ } e \mid \text{true} \mid \text{false} \mid$
 $() \mid \iota \mid \text{fun } x \mapsto c \mid e \# \text{op} \mid h$
 Handler $h ::= \text{handler val } x \mapsto c_v \mid (e_i \# \text{op}_i x_i k_i \mapsto c_i)_i$
 Computation $c ::= \text{val } e \mid \text{let } x = c_1 \text{ in } c_2 \mid \text{let rec } f x = c_1 \text{ in } c_2 \mid$
 $\text{iszero } e \mid \text{if } e \text{ then } c_1 \text{ else } c_2 \mid \text{absurd } e \mid e_1 e_2 \mid$
 $\text{with } e \text{ handle } c$

We describe informally what the various parts mean, and refer the readers to [6] for a more thorough introduction. Eff uses a fine grain call-by-value evaluation strategy [13], which means that it distinguishes effect-free *expressions* and possibly effectful *computations*. There are corresponding *expression* and *computation types*.

2.1 Expressions

An expression is either a variable x , zero 0 , a successor `succ e` , a boolean value `true` or `false`, the unit `()`, an effect instance ι , a function abstraction of a computation, an operation `e # op`, or a handler transforming computations to computations. We briefly comment on the ones that are peculiar to Eff.

An *effect type*

```
effect (operation opi : Ai → Bi)i end
```

declares operations `opi` with given types of parameters A_i and results B_i . (Here and elsewhere, $(\dots)_i$ indicates that \dots may be repeated finitely many times.) For example, the effect type of exceptions is

```
effect
  operation abort : unit → empty
end
```

and the effect type `ref` for a reference holding a natural number is

```
effect
  operation lookup : unit → nat
  operation update : nat → unit
end
```

(1)

Effect instances ι are a way of making several copies of the same computational effect. For example, there may be several communication channels, several mutable references, etc. In this respect Eff differs from [5], where bare operations are considered; in terms of Eff that is like having a single instance of each effect.

The expression type $E^{\{\iota_1, \dots, \iota_n\}}$ is inhabited by expressions which evaluate to one of the instances ι_1, \dots, ι_n , whose effect type is E . We call $\{\iota_1, \dots, \iota_n\}$ a *region* and abbreviate it as ρ . A smaller region is more informative, so ideally we would like all of them to be just singletons. But this is not possible because instances are first-class values and so we can write

```
let x = (if b then  $\iota_1$  else  $\iota_2$ ) in ...
```

The best we can say about x is that its type is $E^{\{\iota_1, \iota_2\}}$.

2.2 Computations

A *computation type* $A!\{\iota_1 \# \text{op}_1, \dots, \iota_n \# \text{op}_n\}$ means that the computation either produces a (pure) value of type A , or triggers one of the listed operations $\iota_1 \# \text{op}_1, \dots, \iota_n \# \text{op}_n$. We abbreviate such finite sets of operations with the letter δ and call them *dirt*. A computation is either a pure value $\text{val } e$, a **let** binding, a recursive function definition, a zero-test $\text{iszero } e$, a conditional statement, a destructor for the **empty** type, an application, or a **handle** construct.

The computation $\text{val } e$ is pure and indicates a “final” result e , while an operation applied to a parameter $\iota \# \text{op } e$ is the principal way of triggering an effect. By itself $\iota \# \text{op } e$ is just a suspended computation with an associated continuation. For an actual effect to take place it has to be handled by a handler, as described below. The continuation associated with $\iota \# \text{op } e$ is $\text{fun } x \mapsto \text{val } x$. Such operations are known as *generic effects* [7].

A binding $\text{let } x = c_1 \text{ in } c_2$ is evaluated as follows:

1. if c_1 evaluates to $\text{val } e$ then the binding evaluates to c_2 with x bound to e ,
2. if c_1 evaluates to an operation $\iota \# \text{op } e$ with continuation κ , then the binding evaluates to $\iota \# \text{op } e$ with continuation $\text{fun } y \mapsto (\text{let } x = \kappa y \text{ in } c_2)$.

It may help to think of **val** and **let** as being similar to Haskell **return** and **do**, respectively. In ML **val** is invisible, while **let** is essentially the same as ours.

The handling construct applies a handler to a computation. If h is the *handler*

$$\text{handler val } x \mapsto c_v \mid (e_i \# \text{op}_i x_i k_i \mapsto c_i)_i$$

and c is a computation then **with** h **handle** c first evaluates c and then evaluates the clause of the handler that matches the result given by c . If no clause matches, c evaluated to an operation which is propagated outwards. In all cases the handler wraps itself around the continuation so that subsequent operations are handled as well. More precisely:

1. if c evaluates to $\text{val } e$, then the handling construct evaluates to c_v with x bound to e ,
2. if c evaluates to $e_i \# \text{op}_i e'$ with continuation κ , then the handling construct evaluates to c_i with x_i and k_i bound to e' and $\text{fun } y \mapsto \text{with } h \text{ handle } \kappa y$, respectively.
3. if c evaluates to any other operation $\iota \# \text{op } e'$ with continuation κ , then the handling construct acts as if h contained the clause

$$\iota \# \text{op } x k \mapsto (\text{let } y = \iota \# \text{op } x \text{ in } k y).$$

Thus it evaluates to $\iota \# \text{op } e'$ with continuation $\text{fun } y \mapsto \text{with } h \text{ handle } (\kappa y)$.

We may wrap several handling constructs around c , in which case the inner handler takes precedence. Note that $\text{let } x = c_1 \text{ in } c_2$ is equivalent to

$$\text{with } (\text{handler val } x \mapsto c_2) \text{ handle } c_1$$

so we could theoretically omit **let**.

2.3 Typing Rules

The typing system of core Eff has two typing judgments,

$$\Gamma \vdash_e e : A \quad \text{and} \quad \Gamma \vdash_c c : \underline{C}$$

stating that an expression e has expression type A , and that a computation c has computation type \underline{C} , respectively. As usual, Γ is a typing context

$$x_1 : A_1, \dots, x_n : A_n$$

which assigns expression types A_i to distinct variables x_i . We also have subtyping of expression types $A \leq A'$ and computation types $\underline{C} \leq \underline{C}'$. We fix an assignment Ξ of effect types to instances, and assume implicitly that all dirts and regions appearing in the rules are well formed, i.e., if $\iota \# \text{op}$ appears then in fact $(\iota : E) \in \Xi$ and $\text{op} \in E$, and that a region ρ in E^ρ mentions only instances whose effect type is E .

The typing rules for variables, primitive constants and functions are standard:

$$\frac{(x : A) \in \Gamma}{\Gamma \vdash_e x : A} \quad \Gamma \vdash_e 0 : \text{nat} \quad \frac{\Gamma \vdash_e e : \text{nat}}{\Gamma \vdash_e \text{succ } e : \text{nat}} \quad \Gamma \vdash_e () : \text{unit}$$

$$\Gamma \vdash_e \text{true} : \text{bool} \quad \Gamma \vdash_e \text{false} : \text{bool} \quad \frac{\Gamma, x : A \vdash_c c : \underline{C}}{\Gamma \vdash_e (\text{fun } x \mapsto c) : A \rightarrow \underline{C}}$$

The typing rule for instances

$$\frac{\iota \in \rho}{\Gamma \vdash_e \iota : E^\rho}$$

verifies that the instance ι is in the region ρ and consults Ξ to check that ρ contains only instances whose effect type is E . An operation $e \# \text{op}$ has a function type, where the dirt in the codomain must contain operations $\iota \# \text{op}$ for ι ranging over the region associated with e :

$$\frac{\Gamma \vdash_e e : E^\rho \quad (\text{op} : A \rightarrow B) \in E \quad \forall \iota \in \rho. (\iota \# \text{op}) \in \delta}{\Gamma \vdash_e e \# \text{op} : A \rightarrow B! \delta}$$

The handler type $A! \delta \Rightarrow \underline{C}$ expresses the fact that a handler transforms computations of type $A! \delta$ to computations of type \underline{C} . The typing rule for handlers

$$\frac{\Gamma, x : A \vdash_c c_v : \underline{C} \quad \Gamma \vdash_e e_i : E_i^{\rho_i} \quad (\text{op}_i : A_i \rightarrow B_i) \in E_i \quad \Gamma, x_i : A_i, k_i : B_i \rightarrow \underline{C} \vdash_c c_i : \underline{C} \quad \forall (\iota \# \text{op}) \in \delta. \exists i. \rho_i = \{\iota\} \wedge \text{op}_i = \text{op}}{\Gamma \vdash_e (\text{handler val } x \mapsto c_v \mid (e_i \# \text{op}_i x_i k_i \mapsto c_i)_i) : A! \delta \Rightarrow \underline{C}}$$

checks that the instances and operations are paired up correctly according to their effect types, all clauses have computation type \underline{C} under suitable assumptions on bound variables, and every operation in δ is handled by a clause. Note

that an operation $\iota \# \text{op}$ is taken to be handled by a clause $e_i \# \text{op}_i x_i k_i \mapsto c_i$ only if $\text{op} = \text{op}_i$ and e_i is ascertained to have exactly the region $\{\iota\}$. If the region were $\{\iota, \iota'\}$ we could not tell whether the clause handles $\iota \# \text{op}$ or $\iota' \# \text{op}$.

The typing rules for computations hold no surprises at all:

$$\frac{\Gamma \vdash_e e : A}{\Gamma \vdash_c \text{val } e : A! \delta} \quad \frac{\Gamma \vdash_c c_1 : A! \delta \quad \Gamma, x : A \vdash_c c_2 : B! \delta}{\Gamma \vdash_c (\text{let } x = c_1 \text{ in } c_2) : B! \delta}$$

$$\frac{\Gamma, f : A \rightarrow \underline{C}, x : A \vdash_c c_1 : \underline{C} \quad \Gamma, f : A \rightarrow \underline{C} \vdash_c c_2 : \underline{D}}{\Gamma \vdash_c (\text{let rec } f x = c_1 \text{ in } c_2) : \underline{D}}$$

$$\frac{\Gamma \vdash_e e : \text{nat}}{\Gamma \vdash_c \text{iszero } e : \text{bool}! \delta} \quad \frac{\Gamma \vdash_e e : \text{bool} \quad \Gamma \vdash_c c_1 : \underline{C} \quad \Gamma \vdash_c c_2 : \underline{C}}{\Gamma \vdash_c (\text{if } e \text{ then } c_1 \text{ else } c_2) : \underline{C}}$$

$$\frac{\Gamma \vdash_e e : \text{empty}}{\Gamma \vdash_c (\text{absurd } e) : \underline{C}} \quad \frac{\Gamma \vdash_e e_1 : A \rightarrow \underline{C} \quad \Gamma \vdash_e e_2 : A}{\Gamma \vdash_c e_1 e_2 : \underline{C}}$$

$$\frac{\Gamma \vdash_e e : \underline{C} \Rightarrow \underline{D} \quad \Gamma \vdash_c c : \underline{C}}{\Gamma \vdash_c (\text{with } e \text{ handle } c) : \underline{D}}$$

The structural subtyping rules are expected as well [14]:

$$\frac{\Gamma \vdash_e e : A \quad A \leq A'}{\Gamma \vdash_e e : A'} \quad \frac{\Gamma \vdash_c c : \underline{C} \quad \underline{C} \leq \underline{C'}}{\Gamma \vdash_c c : \underline{C'}} \quad A \leq A$$

$$\frac{A \leq A' \quad A' \leq A''}{A \leq A''} \quad \frac{A' \leq A \quad \underline{C} \leq \underline{C'}}{(A \rightarrow \underline{C}) \leq (A' \rightarrow \underline{C'})} \quad \frac{\underline{C}' \leq \underline{C} \quad \underline{D} \leq \underline{D}'}{(\underline{C} \Rightarrow \underline{D}) \leq (\underline{C}' \Rightarrow \underline{D}')}$$

$$\frac{\rho \subseteq \rho'}{E^\rho \leq E^{\rho'}} \quad \frac{A \leq A' \quad \delta \subseteq \delta'}{A! \delta \leq A'! \delta'}$$

There are two subsumption rules, one for each typing judgment. The subtyping relation is reflexive and transitive. Function and handler types are contravariant in the domain and covariant in the codomain. Indeed, we may always pretend that a handler handles fewer operations, and triggers more operations than it actually does. An effect type is covariant in the region, while a computation type is covariant in both its parts.

Subtyping buys us additional expressivity. For instance, assuming ι_1 and ι_2 are instances of E ,

```

let a =  $\iota_1$  in
let b = (if p then a else  $\iota_2$ ) in
  val (handler val x  $\mapsto$  val () | a # op x k  $\mapsto$  val ())

```

we would like to give a the type $E^{\{\iota_1\}}$ so that the handler can have the more specific type $A!\{\iota_1 \# \text{op}\} \Rightarrow \text{unit}!\emptyset$, but then without subtyping $E^{\{\iota_1\}} \leq E^{\{\iota_1, \iota_2\}}$ we could not give b its type $E^{\{\iota_1, \iota_2\}}$.

2.4 Example: State Handlers

We demonstrate the features of our type system by looking at a state handler for the reference effect `ref` shown in (1). Define `state` to be the function

```
fun r ↦ (handler val x ↦ val (fun s ↦ val x
      | r # lookup x k ↦ val (fun s ↦ k s s)
      | r # update s' k ↦ val (fun s ↦ k () s'))
```

In words, `state` accepts an instance r and returns a handler which handles lookups and updates on r by using the standard functional encoding of the state monad. For any instance ι , expression type A and dirt δ , the function `state` has the type

$$\text{ref}^{\{\iota\}} \rightarrow (A!(\{\iota \# \text{lookup}, \iota \# \text{update}\} \cup \delta) \Rightarrow (\text{nat} \rightarrow A!\delta)!\delta).$$

Thus, if c is a computation of type

$$A!(\{\iota \# \text{lookup}, \iota \# \text{update}\} \cup \delta)$$

then

$$\text{let } h = \text{state } \iota \text{ in (with } h \text{ handle } c) \tag{2}$$

has the type $(\text{nat} \rightarrow A!\delta)!\delta$. That is, we obtained a computation which possibly triggers effects δ and returns a function. Upon application of the function to an initial state effects δ may be triggered again, after which a final result of type A is obtained. In particular, if $\delta = \emptyset$ then (2) is a pure computation.

If we weakened the domain of `state` to `ref` ^{$\{\iota_1, \iota_2\}$} then we would not be able to deduce that `state` handled anything, so we would only be able to assign to `state` the less useful type

$$\text{ref}^{\{\iota_1, \iota_2\}} \rightarrow (A!\delta \Rightarrow (\text{nat} \rightarrow A!\delta)!\delta).$$

We may handle several references by invoking several instantiations of `state`. For example, let c be the computation which swaps the contents of two references:

```
let x1 =  $\iota_1$  # lookup() in
let x2 =  $\iota_2$  # lookup() in
let u =  $\iota_1$  # update x2 in
let v =  $\iota_2$  # update x1 in ()
```

By itself, c has the type

$$\text{unit}!\{\iota_1 \# \text{lookup}, \iota_1 \# \text{update}, \iota_2 \# \text{lookup}, \iota_2 \# \text{update}\},$$

while

$$\text{let } h_1 = \text{state } \iota_1 \text{ in (with } h_1 \text{ handle } c)$$

has type $(\text{nat} \rightarrow \text{unit}!\{\iota_2 \# \text{lookup}, \iota_2 \# \text{update}\})!\{\iota_2 \# \text{lookup}, \iota_2 \# \text{update}\}$. If we handle both instances,

$$\begin{aligned} &\text{let } h_1 = \text{state } \iota_1 \text{ in} \\ &\text{let } h_2 = \text{state } \iota_2 \text{ in} \\ &\quad \text{with } h_2 \text{ handle (with } h_1 \text{ handle } c), \end{aligned}$$

we get the pure type $\text{nat} \rightarrow (\text{nat} \rightarrow \text{unit}!\emptyset)!\emptyset$.

3 Denotational Semantics

An outline of a set-theoretic algebraic semantics of computational effects, as developed by [6,15,16,7,8], is shown in the following table.

<i>Programming language</i>	<i>Algebra</i>
effect type	algebraic signature
expression type	set
expression	element
computation type	free algebra
pure computation	generator
effectful computation	algebraic operation
handler	homomorphism of algebras

To properly account for recursion and non-termination we adapt it to a *domain-theoretic* semantics of algebraic effects under which expression and computation types are *domains*. We first give Curry-style semantics in which terms are interpreted without being typed, and then, following John Reynolds [17], we provide a Church-style semantics in which types receive meanings as well. It does not matter much what kind of domains we use, they could be ω -cpo's, Scott domains, effective Scott domains, or any other kind of domains that model the basic type-theoretic operations (product, function space, coalesced sum).

3.1 Semantics of Expressions and Computations

Expressions could be modeled with *predomains*, because they are inert pieces of data, free from computational effects, including non-termination. However, the bottom element is useful for denotation of ill-typed expressions and runtime errors. In any case, the predomain nature of expressions will be captured later on by the partial equivalence relations. The domains for computation types are free in a suitable sense, i.e., they enjoy a recursion principle which acts as a substitute for the universality of free algebras. We assume a given set of all instances \mathbb{I} and

a set of all operation symbols \mathbb{O} , write $+$ for a disjoint sum, and \oplus for coalesced sum.

To interpret expressions and computations we need suitable domains of *values* V and *results* R , respectively. These have to be large enough to contain all possible denotations of expressions and computations, which is usually achieved by solving suitable recursive domain equations, such as

$$\begin{aligned} V &= \mathbb{N}_\perp \oplus \{0, 1\}_\perp \oplus \{\star\}_\perp \oplus \mathbb{I}_\perp \oplus R^V \oplus R^R, \\ R &= (V + \mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp. \end{aligned}$$

The summands in the equation for V correspond to various expression types. The recursive domain equation for R says that a non-bottom element of R is either a value, or a quadruple $(\iota, \text{op}, v, \kappa)$ corresponding to the operation $\iota \# \text{op}$ applied to parameter v and with continuation κ . However, as in [17], we shall assume only that V and R are large enough to contain the various components needed for the semantics as *retracts*:

$$\begin{array}{ccc} \mathbb{N}_\perp \begin{array}{c} \xrightarrow{s_{\text{nat}}} \\ \xleftarrow{r_{\text{nat}}} \end{array} V & \{0, 1\}_\perp \begin{array}{c} \xrightarrow{s_{\text{bool}}} \\ \xleftarrow{r_{\text{bool}}} \end{array} V & \{\star\}_\perp \begin{array}{c} \xrightarrow{s_{\text{unit}}} \\ \xleftarrow{r_{\text{unit}}} \end{array} V \\ \mathbb{I}_\perp \begin{array}{c} \xrightarrow{s_{\text{effect}}} \\ \xleftarrow{r_{\text{effect}}} \end{array} V & R^V \begin{array}{c} \xrightarrow{s_{\rightarrow}} \\ \xleftarrow{r_{\rightarrow}} \end{array} V & R^R \begin{array}{c} \xrightarrow{s_{\Rightarrow}} \\ \xleftarrow{r_{\Rightarrow}} \end{array} V \end{array}$$

and

$$(V + \mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp \begin{array}{c} \xrightarrow{s_{\text{res}}} \\ \xleftarrow{r_{\text{res}}} \end{array} R$$

Thus, R and V could be solutions to the above domain equations, but they could also both be a universal domain, or just a (non-trivial) reflexive domain.¹ Note that there are further canonical retractions

$$\begin{aligned} V &\begin{array}{c} \xrightarrow{s_{\text{val}}} \\ \xleftarrow{r_{\text{val}}} \end{array} (V + \mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp \\ (\mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp &\begin{array}{c} \xrightarrow{s_{\text{oper}}} \\ \xleftarrow{r_{\text{oper}}} \end{array} (V + \mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp \end{aligned}$$

which in combination with the section-retraction pair $(s_{\text{res}}, r_{\text{res}})$ convert elements of R to pure values and operations, respectively.

The domain $(V + \mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp$ has the following *recursion principle*. Given a domain D and maps

$$h_{\text{val}} : V \longrightarrow D \quad \text{and} \quad h_{\text{oper}} : \mathbb{I} \times \mathbb{O} \times V \times R^V \longrightarrow D,$$

there is a unique strict map $h : (V + \mathbb{I} \times \mathbb{O} \times V \times R^V)_\perp \longrightarrow D$ such that, for all $v \in V$, $\iota \in \mathbb{I}$, $\text{op} \in \mathbb{O}$, $\kappa : V \rightarrow R$,

$$\begin{aligned} h(s_{\text{val}}(v)) &= h_{\text{val}}(v), \\ h(s_{\text{oper}}(\iota, \text{op}, v, \kappa)) &= h_{\text{oper}}(\iota, \text{op}, v, h \circ r_{\text{res}} \circ \kappa). \end{aligned}$$

¹ A domain D is *reflexive* if it contains its functions space D^D as a retract, and is therefore a model of the untyped λ -calculus.