

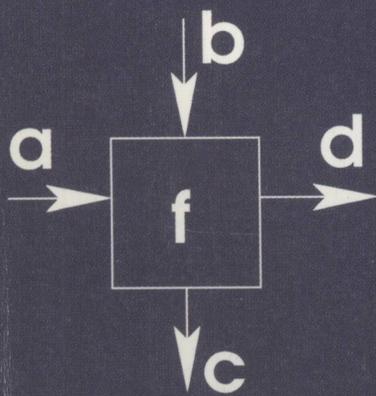
Tutorial

LNCS 3622

Varmo Vene
Tarmo Uustalu (Eds.)

Advanced Functional Programming

5th International School, AFP 2004
Tartu, Estonia, August 2004
Revised Lectures



Springer

TP311.1-53

A257
2004

Varmo Vene Tarmo Uustalu (Eds.)

Advanced Functional Programming

5th International School, AFP 2004
Tartu, Estonia, August 14 – 21, 2004
Revised Lectures



E200600007



Springer

Volume Editors

Varmo Vene
University of Tartu
Department of Computer Science
J. Liivi 2, EE-50409 Tartu, Estonia
E-mail: varmo@cs.ut.ee

Tarmo Uustalu
Institute of Cybernetics
Akadeemia tee 21, EE-12618 Tallinn, Estonia
E-mail: tarmo@cs.ioc.ee

Library of Congress Control Number: 2005931987

CR Subject Classification (1998): D.1.1, D.1, D.3, F.3, D.2

ISSN 0302-9743
ISBN-10 3-540-28540-7 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-28540-3 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 2005
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11546382 06/3142 5 4 3 2 1 0

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

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

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

New York University, NY, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Preface

This volume contains the revised lecture notes corresponding to nine of the lecture courses presented at the 5th International School on Advanced Functional Programming, AFP 2004, held in Tartu, Estonia, August 14–21, 2004.

The goal of the AFP schools is to inform the wide international communities of computer science students and software production professionals about the new and important developments in the area of functional programming. The schools put a special emphasis on practical applications of advanced techniques. The Tartu school was preceded by four earlier schools in Båstad, Sweden (1995, LNCS 925), Olympia, WA, USA (1996, LNCS 1129), Braga, Portugal (1998, LNCS 1608) and Oxford, UK (2002, LNCS 2638).

The scientific programme of AFP 2004 consisted of five preparatory (“intermediate”) courses, given by John Hughes (Chalmers University of Technology, Göteborg, Sweden), Doaitse Swierstra (Universiteit Utrecht, The Netherlands) and Rinus Plasmeijer (Radboud Universiteit Nijmegen, The Netherlands), and nine regular (“advanced”) courses, presented by Atze Dijkstra (Universiteit Utrecht, The Netherlands), Doaitse Swierstra, John Hughes, Conor McBride (University of Nottingham, UK), Alberto Pardo (Universidade de la República, Montevideo, Uruguay), Rinus Plasmeijer, Bernard Pope (University of Melbourne, Australia), Peter Thiemann (Universität Freiburg, Germany), and Simon Thompson (University of Kent, UK). There was also a student session.

The school attracted a record number of 68 participants from 16 countries (inclusive of the lecturers and organizers).

This volume contains the notes for the advanced courses. Following the school, the lecturers revised the notes they had prepared for the school. The revised notes were each carefully checked by two or three second readers selected from among the most qualified available and then revised once more by the lecturers. We are proud to commend the final texts to everyone wishing to acquire first-hand knowledge about some of the exciting and trendsetting developments in functional programming.

We are grateful to our sponsors, to the Faculty of Mathematics and Computer Science of the University of Tartu, to the lecturers and the second readers for their hard work on the oral presentations, and the notes, and to all our participants. You made the school what it was.

Tartu and Tallinn, June 2005

Varmo Vene
Tarmo Uustalu

Organization

Host Institution

AFP 2004 was organized by the Department of Computer Science of the University of Tartu in cooperation with the Center for Dependable Computing (CDC), an Estonian center of excellence in research.

Programme Committee

Varmo Vene (University of Tartu, Estonia) (chairman)
Johan Jeuring (Universiteit Utrecht, The Netherlands)
Tarmo Uustalu (Institute of Cybernetics, Tallinn, Estonia)

Organizing Committee

Varmo Vene (University of Tartu, Estonia) (chairman)
Härmel Nestra (University of Tartu, Estonia)
Vesal Vojdani (University of Tartu, Estonia)
Tarmo Uustalu (Institute of Cybernetics, Tallinn, Estonia)

Second Readers

Venanzio Capretta (University of Ottawa, Canada)
James Cheney (University of Edinburgh, UK)
Catarina Coquand (Chalmers University of Technology, Sweden)
Jeremy Gibbons (University of Oxford, UK)
Thomas Hallgren (Oregon Graduate Institute, Portland, OR, USA)
Michael Hanus (Christian-Albrechts-Universität zu Kiel, Germany)
Johan Jeuring (Universiteit Utrecht, The Netherlands)
Jerzy Karczmarczuk (Université Caen, France)
Ralf Lämmel (CWI, Amsterdam, The Netherlands)
Andres Löb (Universiteit Utrecht, The Netherlands)
Nicolas Magaud (University of New South Wales, Sydney, Australia)
Simon Marlow (Microsoft Research, Cambridge, UK)
Ross Paterson (City University, London, UK)
Simon Peyton Jones (Microsoft Research, Cambridge, UK)
Colin Runciman (University of York, UK)
Tim Sheard (Portland State University, Portland, OR, USA)
Joost Visser (Universidade do Minho, Braga, Portugal)
Eric Van Wyk (University of Minnesota, Minneapolis, MN, USA)

Sponsoring Institutions

Tiigriülikool programme of the Estonian Information Technology Foundation

National Centers of Excellence programme of the Estonian Ministry of Education and Research

EU FP5 IST programme via the thematic network project APPSEM II

Lecture Notes in Computer Science

For information about Vols. 1–3591

please contact your bookseller or Springer

Vol. 3710: M. Barni, I. Cox, T. Kalker, H.J. Kim (Eds.), *Digital Watermarking*. XII, 485 pages. 2005.

Vol. 3703: F. Fages, S. Soliman (Eds.), *Principles and Practice of Semantic Web Reasoning*. VIII, 163 pages. 2005.

Vol. 3702: B. Beckert (Ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*. XIII, 343 pages. 2005. (Subseries LNAI).

Vol. 3698: U. Furbach (Ed.), *KI 2005: Advances in Artificial Intelligence*. XIII, 409 pages. 2005. (Subseries LNAI).

Vol. 3697: W. Duch, J. Kacprzyk, E. Oja, S. Zadrozny (Eds.), *Artificial Neural Networks: Formal Models and Their Applications - ICANN 2005, Part II*. XXXII, 1045 pages. 2005.

Vol. 3696: W. Duch, J. Kacprzyk, E. Oja, S. Zadrozny (Eds.), *Artificial Neural Networks: Biological Inspirations - ICANN 2005, Part I*. XXXI, 703 pages. 2005.

Vol. 3687: S. Singh, M. Singh, C. Apte, P. Perner (Eds.), *Pattern Recognition and Image Analysis, Part II*. XXV, 809 pages. 2005.

Vol. 3686: S. Singh, M. Singh, C. Apte, P. Perner (Eds.), *Pattern Recognition and Data Mining, Part I*. XXVI, 689 pages. 2005.

Vol. 3684: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part IV*. LXXIX, 933 pages. 2005. (Subseries LNAI).

Vol. 3683: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. LXXIX, 1397 pages. 2005. (Subseries LNAI).

Vol. 3682: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. LXXIX, 1371 pages. 2005. (Subseries LNAI).

Vol. 3681: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part I*. LXXX, 1319 pages. 2005. (Subseries LNAI).

Vol. 3678: A. McLysaght, D.H. Huson (Eds.), *Comparative Genomics*. VIII, 167 pages. 2005. (Subseries LNBI).

Vol. 3677: J. Dittmann, S. Katzenbeisser, A. Uhl (Eds.), *Communications and Multimedia Security*. XIII, 360 pages. 2005.

Vol. 3675: Y. Luo (Ed.), *Cooperative Design, Visualization, and Engineering*. XI, 264 pages. 2005.

Vol. 3674: W. Jonker, M. Petković (Eds.), *Secure Data Management*. X, 241 pages. 2005.

Vol. 3672: C. Hankin, I. Siveroni (Eds.), *Static Analysis*. X, 369 pages. 2005.

Vol. 3671: S. Bressan, S. Ceri, E. Hunt, Z.G. Ives, Z. Belahsene, M. Rys, R. Unland (Eds.), *Database and XML Technologies*. X, 239 pages. 2005.

Vol. 3670: M. Bravetti, L. Kloul, G. Zavattaro (Eds.), *Formal Techniques for Computer Systems and Business Processes*. XIII, 349 pages. 2005.

Vol. 3665: K. S. Candan, A. Celentano (Eds.), *Advances in Multimedia Information Systems*. X, 221 pages. 2005.

Vol. 3664: C. Türker, M. Agosti, H.-J. Schek (Eds.), *Peer-to-Peer, Grid, and Service-Oriented in Digital Library Architectures*. X, 261 pages. 2005.

Vol. 3663: W.G. Kropatsch, R. Sablatnig, A. Hanbury (Eds.), *Pattern Recognition*. XIV, 512 pages. 2005.

Vol. 3662: C. Baral, G. Greco, N. Leone, G. Terracina (Eds.), *Logic Programming and Nonmonotonic Reasoning*. XIII, 454 pages. 2005. (Subseries LNAI).

Vol. 3661: T. Panayiotopoulos, J. Gratch, R. Aylett, D. Ballin, P. Olivier, T. Rist (Eds.), *Intelligent Virtual Agents*. XIII, 506 pages. 2005. (Subseries LNAI).

Vol. 3660: M. Beigl, S. Intille, J. Kekimoto, H. Tokuda (Eds.), *UbiComp 2005: Ubiquitous Computing*. XVII, 394 pages. 2005.

Vol. 3659: J.R. Rao, B. Sunar (Eds.), *Cryptographic Hardware and Embedded Systems - CHES 2005*. XIV, 458 pages. 2005.

Vol. 3658: V. Matoušek, P. Mautner, T. Pavelka (Eds.), *Text, Speech and Dialogue*. XV, 460 pages. 2005. (Subseries LNAI).

Vol. 3655: A. Aldini, R. Gorrieri, F. Martinelli (Eds.), *Foundations of Security Analysis and Design III*. VII, 273 pages. 2005.

Vol. 3654: S. Jajodia, D. Wijesekera (Eds.), *Data and Applications Security XIX*. X, 353 pages. 2005.

Vol. 3653: M. Abadi, L.d. Alfaro (Eds.), *CONCUR 2005 - Concurrency Theory*. XIV, 578 pages. 2005.

Vol. 3652: A. Rauber, S. Christodoulakis, A. M. Tjoa (Eds.), *Research and Advanced Technology for Digital Libraries*. XVIII, 545 pages. 2005.

Vol. 3649: W.M.P. van der Aalst, B. Benatallah, F. Casati, F. Curbera (Eds.), *Business Process Management*. XII, 472 pages. 2005.

Vol. 3648: J.C. Cunha, P.D. Medeiros (Eds.), *Euro-Par 2005 Parallel Processing*. XXXVI, 1299 pages. 2005.

Vol. 3646: A. F. Famili, J.N. Kok, J.M. Peña, A. Siebes, A. Feelders (Eds.), *Advances in Intelligent Data Analysis VI*. XIV, 522 pages. 2005.

Vol. 3645: D.-S. Huang, X.-P. Zhang, G.-B. Huang (Eds.), *Advances in Intelligent Computing, Part II*. XIII, 1010 pages. 2005.

- Vol. 3644: D.-S. Huang, X.-P. Zhang, G.-B. Huang (Eds.), *Advances in Intelligent Computing, Part I. XXVII*, 1101 pages. 2005.
- Vol. 3642: D. Ślęzak, J. Yao, J.F. Peters, W. Ziarko, X. Hu (Eds.), *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part II. XXIII*, 738 pages. 2005. (Subseries LNAI).
- Vol. 3641: D. Ślęzak, G. Wang, M. Szczuka, I. Düntsch, Y. Yao (Eds.), *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part I. XXIV*, 742 pages. 2005. (Subseries LNAI).
- Vol. 3639: P. Godefroid (Ed.), *Model Checking Software. XI*, 289 pages. 2005.
- Vol. 3638: A. Butz, B. Fisher, A. Krüger, P. Olivier (Eds.), *Smart Graphics. XI*, 269 pages. 2005.
- Vol. 3637: J. M. Moreno, J. Madrenas, J. Cosp (Eds.), *Evolvable Systems: From Biology to Hardware. XI*, 227 pages. 2005.
- Vol. 3636: M.J. Blesa, C. Blum, A. Roli, M. Sampels (Eds.), *Hybrid Metaheuristics. XII*, 155 pages. 2005.
- Vol. 3634: L. Ong (Ed.), *Computer Science Logic. XI*, 567 pages. 2005.
- Vol. 3633: C. Bauzer Medeiros, M. Egenhofer, E. Bertino (Eds.), *Advances in Spatial and Temporal Databases. XIII*, 433 pages. 2005.
- Vol. 3632: R. Nieuwenhuis (Ed.), *Automated Deduction – CADE-20. XIII*, 459 pages. 2005. (Subseries LNAI).
- Vol. 3631: J. Eder, H.-M. Haav, A. Kalja, J. Penjam (Eds.), *Advances in Databases and Information Systems. XIII*, 393 pages. 2005.
- Vol. 3630: M.S. Capcarrere, A.A. Freitas, P.J. Bentley, C.G. Johnson, J. Timmis (Eds.), *Advances in Artificial Life. XIX*, 949 pages. 2005. (Subseries LNAI).
- Vol. 3629: J.L. Fiadeiro, N. Harman, M. Roggenbach, J. Rutten (Eds.), *Algebra and Coalgebra in Computer Science. XI*, 457 pages. 2005.
- Vol. 3628: T. Gschwind, U. Aßmann, O. Nierstrasz (Eds.), *Software Composition. X*, 199 pages. 2005.
- Vol. 3627: C. Jacob, M.L. Pilat, P.J. Bentley, J. Timmis (Eds.), *Artificial Immune Systems. XII*, 500 pages. 2005.
- Vol. 3626: B. Ganter, G. Stumme, R. Wille (Eds.), *Formal Concept Analysis. X*, 349 pages. 2005. (Subseries LNAI).
- Vol. 3625: S. Kramer, B. Pfahring (Eds.), *Inductive Logic Programming. XIII*, 427 pages. 2005. (Subseries LNAI).
- Vol. 3624: C. Chekuri, K. Jansen, J.D.P. Rolim, L. Trevisan (Eds.), *Approximation, Randomization and Combinatorial Optimization. XI*, 495 pages. 2005.
- Vol. 3623: M. Liškiewicz, R. Reischuk (Eds.), *Fundamentals of Computation Theory. XV*, 576 pages. 2005.
- Vol. 3622: V. Vene, T. Uustalu (Eds.), *Advanced Functional Programming. IX*, 359 pages. 2005.
- Vol. 3621: V. Shoup (Ed.), *Advances in Cryptology – CRYPTO 2005. XI*, 568 pages. 2005.
- Vol. 3620: H. Muñoz-Avila, F. Ricci (Eds.), *Case-Based Reasoning Research and Development. XV*, 654 pages. 2005. (Subseries LNAI).
- Vol. 3619: X. Lu, W. Zhao (Eds.), *Networking and Mobile Computing. XXIV*, 1299 pages. 2005.
- Vol. 3618: J. Jedrzejowicz, A. Szepletowski (Eds.), *Mathematical Foundations of Computer Science 2005. XVI*, 814 pages. 2005.
- Vol. 3617: F. Roli, S. Vitulano (Eds.), *Image Analysis and Processing – ICIAAP 2005. XXIV*, 1219 pages. 2005.
- Vol. 3615: B. Ludäscher, L. Raschid (Eds.), *Data Integration in the Life Sciences. XII*, 344 pages. 2005. (Subseries LNBI).
- Vol. 3614: L. Wang, Y. Jin (Eds.), *Fuzzy Systems and Knowledge Discovery, Part II. XLI*, 1314 pages. 2005. (Subseries LNAI).
- Vol. 3613: L. Wang, Y. Jin (Eds.), *Fuzzy Systems and Knowledge Discovery, Part I. XLI*, 1334 pages. 2005. (Subseries LNAI).
- Vol. 3612: L. Wang, K. Chen, Y. S. Ong (Eds.), *Advances in Natural Computation, Part III. LXI*, 1326 pages. 2005.
- Vol. 3611: L. Wang, K. Chen, Y. S. Ong (Eds.), *Advances in Natural Computation, Part II. LXI*, 1292 pages. 2005.
- Vol. 3610: L. Wang, K. Chen, Y. S. Ong (Eds.), *Advances in Natural Computation, Part I. LXI*, 1302 pages. 2005.
- Vol. 3608: F. Dehne, A. López-Ortiz, J.-R. Sack (Eds.), *Algorithms and Data Structures. XIV*, 446 pages. 2005.
- Vol. 3607: J.-D. Zucker, L. Saitta (Eds.), *Abstraction, Reformulation and Approximation. XII*, 376 pages. 2005. (Subseries LNAI).
- Vol. 3606: V. Malyshekin (Ed.), *Parallel Computing Technologies. XII*, 470 pages. 2005.
- Vol. 3604: R. Martin, H. Bez, M. Sabin (Eds.), *Mathematics of Surfaces XI. IX*, 473 pages. 2005.
- Vol. 3603: J. Hurd, T. Melham (Eds.), *Theorem Proving in Higher Order Logics. IX*, 409 pages. 2005.
- Vol. 3602: R. Eigenmann, Z. Li, S.P. Midkiff (Eds.), *Languages and Compilers for High Performance Computing. IX*, 486 pages. 2005.
- Vol. 3599: U. Aßmann, M. Aksit, A. Resink (Eds.), *Model Driven Architecture. X*, 235 pages. 2005.
- Vol. 3598: H. Murakami, H. Nakashima, H. Tokuda, M. Yasumura, *Ubiquitous Computing Systems. XIII*, 275 pages. 2005.
- Vol. 3597: S. Shimojo, S. Ichii, T.W. Ling, K.-H. Song (Eds.), *Web and Communication Technologies and Internet-Related Social Issues - HSI 2005. XIX*, 368 pages. 2005.
- Vol. 3596: F. Dau, M.-L. Mugnier, G. Stumme (Eds.), *Conceptual Structures: Common Semantics for Sharing Knowledge. XI*, 467 pages. 2005. (Subseries LNAI).
- Vol. 3595: L. Wang (Ed.), *Computing and Combinatorics. XVI*, 995 pages. 2005.
- Vol. 3594: J.C. Setubal, S. Verjovski-Almeida (Eds.), *Advances in Bioinformatics and Computational Biology. XIV*, 258 pages. 2005. (Subseries LNBI).
- Vol. 3593: V. Mařík, R. W. Brennan, M. Pěchouček (Eds.), *Holonic and Multi-Agent Systems for Manufacturing. XI*, 269 pages. 2005. (Subseries LNAI).
- Vol. 3592: S. Katsikas, J. Lopez, G. Pernul (Eds.), *Trust, Privacy and Security in Digital Business. XII*, 332 pages. 2005.

¥462.56元

Table of Contents

Typing Haskell with an Attribute Grammar <i>Atze Dijkstra, S. Doaitse Swierstra</i>	1
Programming with Arrows <i>John Hughes</i>	73
Epigram: Practical Programming with Dependent Types <i>Conor McBride</i>	130
Combining Datatypes and Effects <i>Alberto Pardo</i>	171
GEC: A Toolkit for Generic Rapid Prototyping of Type Safe Interactive Applications <i>Peter Achten, Marko van Eekelen, Rinus Plasmeijer, Arjen van Weelden</i>	210
A Functional Shell That Operates on Typed and Compiled Applications <i>Rinus Plasmeijer, Arjen van Weelden</i>	245
Declarative Debugging with Buddha <i>Bernard Pope</i>	273
Server-Side Web Programming in WASH <i>Peter Thiemann</i>	309
Refactoring Functional Programs <i>Simon Thompson</i>	331
Author Index	359

Typing Haskell with an Attribute Grammar

Atze Dijkstra and S. Doaitse Swierstra

Institute of Information and Computing Sciences,
Utrecht University, P.O.Box 80.089,
3508 TB Utrecht, Netherlands
{atze, doaitse}@cs.uu.nl

Abstract. A great deal has been written about type systems. Much less has been written about implementing them. Even less has been written about implementations of complete compilers in which all aspects come together. This paper fills this gap by describing the implementation of a series of compilers for a simplified variant of Haskell. By using an attribute grammar system, aspects of a compiler implementation can be described separately and added in a sequence of steps, thereby giving a series of increasingly complex (working) compilers. Also, the source text of both this paper and the executable compilers come from the same source files by an underlying minimal weaving system. Therefore, source and explanation is kept consistent.

1 Introduction and Overview

Haskell98 [31] is a complex language, not to mention its more experimental incarnations. Though also intended as a research platform, realistic compilers for Haskell [1] have grown over the years and understanding and experimenting with those compilers is not an easy task. Experimentation on a smaller scale usually is based upon relatively simple and restricted implementations [20], often focusing only on a particular aspect of the language and/or its implementation. This paper aims at walking somewhere between this complexity and simplicity by

- Describing the implementation of essential aspects of Haskell (or any other (functional) programming language), hence the name Essential Haskell (EH) used for simplified variants of Haskell¹ in these notes.
- Describing these aspects separately in order to provide a better understanding.
- Adding these aspects on top of each other in an incremental way, thus leading to a sequence of compilers, each for a larger subset of complete Haskell (and extensions).
- Using tools like the Utrecht University Attribute Grammar (UUAG) system [3], hereafter referred to as the AG system, to allow for separate descriptions for the various aspects.

¹ The 'E' in EH might also be expanded to other aspects of the compiler, like being an Example.

The remaining sections of this introduction will expand on this by looking at the intentions, purpose and limitations of these notes in more detail. This is followed by a short description of the individual languages for which we develop compilers throughout these notes. The last part of the introduction contains a small tutorial on the AG system used in these notes. After the introduction we continue with discussing the implementation of the first three compilers (sections 2, 3 and 4) out of a (currently) sequence of ten compilers. On the web site [11] for this project the full distribution of the code for these compilers can be found. We conclude these notes by reflecting upon our experiences with the AG system and the creation of these notes (section 5).

1.1 Purpose

For whom is this material intended?

- For students who wish to learn more about the implementation of functional languages. This paper also informally explains the required theory, in particular about type systems.
- For researchers who want to build (e.g.) a prototype and to experiment with extensions to the type system and need a non-trivial and realistic starting point. This paper provides documentation, design rationales and an implementation for such a starting point.
- For those who wish to study a larger example of the tools used to build the compilers in these notes. We demonstrate the use of the AG system, which allows us to separately describe the various aspects of a language implementation. Other tools for maintaining consistency between different versions of the resulting compilers and the source code text included in these notes are also used, but will not be discussed.

For this intended audience these notes provide:

- A description of the implementation of a type checker/inferencer for a subset of Haskell. We describe the first three languages of a (currently) sequence of ten, that end in a full implementation of an extended Haskell.
- A description of the semantics of Haskell, lying between the more formal [16,14] and more implementation oriented [21,33] and similar to other combinations of theory and practice [34].
- A gradual instead of a big bang explanation.
- Empirical support for the belief that the complexity of a compiler can be managed by splitting the implementation of the compiler into separate aspects.
- A working combination of otherwise usually separately proven or implemented features.

We will come back to this in our conclusion (see section 5).

We restrict ourselves in the following ways, partly because of space limitations, partly by design:

- We do not discuss extensions to Haskell implemented in versions beyond the last version presented in these notes. See section 1.3 for a preciser description of what can and cannot be found in these notes with respect to Haskell features.
- We concern ourselves with typing only. Other aspects, like pretty printing and parsing, are not discussed. However, the introduction to the AG system (see section 1.4) gives some examples of the pretty printing and the interaction between parsing, AG code and Haskell code.
- We do not deal with type theory or parsing theory as a subject on its own. This paper is intended to describe “how to implement” and will use theory from that point of view. Theoretical aspects are touched upon from a more intuitive point of view.

Although informally and concisely introduced where necessary, familiarity with the following will make reading and understanding these notes easier:

- Functional programming, in particular using Haskell
- Compiler construction in general
- Type systems, λ -calculus
- Parser combinator library and AG system [3,38]

For those not familiar with the AG system a short tutorial has been included at the end of this introduction (see section 1.4). It also demonstrates the use of the parser combinators used throughout the implementation of all EH versions.

We expect that by finding a balance between theory and implementation, we serve both those who want to learn and those who want to do research. It is also our belief that by splitting the big problem into smaller aspects the combination can be explained in an easier way.

In the following sections we give examples of the Haskell features present in the series of compilers described in the following chapters. Only short examples are given, so the reader gets an impression of what is explained in more detail and implemented in the relevant versions of the compiler.

1.2 A Short Tour

Though all compilers described in these notes deal with a different issue, they all have in common that they are based on the λ -calculus, most of the time using the syntax and semantics of Haskell. The first version of our series of compilers therefore accepts a language that most closely resembles the λ -calculus, in particular typed λ -calculus extended with **let** expressions and some basic types and type constructors such as *Int*, *Char* and tuples.

EH version 1: λ -calculus. An EH program is a single expression, contrary to a Haskell program which consists of a set of declarations forming a module.

```
let  $i :: Int$ 
     $i = 5$ 
in  $i$ 
```

All variables need to be typed explicitly; absence of an explicit type is considered to be an error. The corresponding compiler (EH version 1, section 2) checks the explicit types against actual types.

Besides the basic types *Int* and *Char*, composite types can be formed by building tuples and defining functions:

```
let  $id :: Int \rightarrow Int$ 
     $id = \lambda x \rightarrow x$ 
     $fst :: (Int, Char) \rightarrow Int$ 
     $fst = \lambda(a, b) \rightarrow a$ 
in  $id (fst (id 3, 'x'))$ 
```

Functions accept one parameter only, which can be a pattern. All types are monomorphic.

EH version 2: Explicit/implicit typing. The next version (EH version 2, section 3) no longer requires the explicit type specifications, which thus may have to be inferred by the compiler.

The reconstructed type information is monomorphic, for example the identity function in:

```
let  $id = \lambda x \rightarrow x$ 
in let  $v = id 3$ 
    in  $id$ 
```

is inferred to have the type $id :: Int \rightarrow Int$.

EH version 3: Polymorphism. The third version (EH version 3, section 4) performs standard Hindley-Milner type inferencing [8,9] which also supports parametric polymorphism. For example,

```
let  $id = \lambda x \rightarrow x$ 
in  $id 3$ 
```

is inferred to have type $id :: \forall a.a \rightarrow a$.

1.3 Haskell Language Elements Not Described

As mentioned before, only a subset of the full sequence of compilers is described in these notes. Currently, as part of an ongoing work [11], in the compilers following the compilers described in these notes, the following Haskell features are dealt with:

EH 4. Quantifiers everywhere: higher ranked types [36,32,7,28] and existentials [30,25,27]. See also the longer version of these notes handed out during the AFP04 summerschool [13].

EH 5. Data types.

EH 6. Kinds, kind inference, kind checking, kind polymorphism.

EH 7. Non extensible records, subsuming tuples.

EH 8. Code generation for a GRIN (Graph Reduction Intermediate Notation) like backend [6,5].

EH 9. Class system, explicit implicit parameters [12].

EH 10. Extensible records [15,22].

Also missing are features which fall in the category syntactic sugar, programming in the large and the like. Haskell incorporates many features which make programming easier and/or manageable. Just to mention a few:

- Binding group analysis
- Syntax directives like infix declarations
- Modules [10,37].
- Type synonyms
- Syntactic sugar for **if**, **do**, list notation and comprehension.

We have deliberately not dealt with these issues. Though necessary and convenient we feel that these features should be added after all else has been dealt with, so as not to make understanding and implementing essential features more difficult.

1.4 An AG Mini Tutorial

The remaining part of the introduction contains a small tutorial on the AG system. The tutorial explains the basic features of the AG system. The explanation of remaining features is postponed to its first use throughout the main text. These places are marked with *AG*. The tutorial can safely be skipped if the reader is already familiar with the AG system.

Haskell and Attribute Grammars (AG). Attribute grammars can be mapped onto functional programs [23,19,4]. Vice versa, the class of functional programs (catamorphisms

[39]) mapped onto can be described by attribute grammars. The AG system exploits this correspondence by providing a notation (attribute grammar) for computations over trees which additionally allows program fragments to be described separately. The AG compiler gathers these fragments, combines these fragments, and generates a corresponding Haskell program.

In this AG tutorial we start with a small example Haskell program (of the right form) to show how the computation described by this program can be expressed in the AG notation and how the resulting Haskell program generated by the AG compiler can be used. The ‘repmin’ problem [4] is used for this purpose. A second example describing a ‘pocket calculator’ (that is, expressions) focusses on more advanced features and typical AG usage patterns.

Repmin a la Haskell. Repmin stands for “replacing the integer valued leaves of a tree by the minimal integer value found in the leaves”. The solution to this problem requires two passes over a tree structure, computing the minimum and computing a new tree with the minimum as its leaves respectively. It is often used as the typical example of a circular program which lends itself well to be described by the AG notation. When described in Haskell it is expressed as a computation over a tree structure:

```
data Tree = Tree_Leaf Int
          | Tree_Bin Tree Tree
deriving Show
```

The computation itself simultaneously computes the minimum of all integers found in the leaves of the tree and the new tree with this minimum value. The result is returned as a tuple computed by function *r*:

```
repmin :: Tree → Tree
repmin t
  = t'
  where (t', tmin) = r t tmin
        r (Tree_Leaf i) m = (Tree_Leaf m, i)
        r (Tree_Bin lt rt) m = (Tree_Bin lt' rt', lmin 'min' rmin)
        where (lt', lmin) = r lt m
              (rt', rmin) = r rt m
```

We can use this function in some setting, for example:

```
tr = Tree_Bin (Tree_Leaf 3) (Tree_Bin (Tree_Leaf 4) (Tree_Leaf 5))
tr' = repmin tr
main :: IO ()
main = print tr'
```

The resulting program produces the following output:

```
Tree_Bin (Tree_Leaf 3) (Tree_Bin (Tree_Leaf 3) (Tree_Leaf 3))
```

The computation of the new tree requires the minimum. This minimum is passed as a parameter m to r at the root of the tree by extracting it from the result of r . The result tuple of the invocation $r \ t \ tmin$ depends on itself via the minimum $tmin$ so it would seem we have a cyclic definition. However, the real dependency is not on the tupled result of r but on its elements because it is the element $tmin$ of the result tuple which is passed back and not the tuple itself. The elements are not cyclically dependent so Haskell's laziness prevents a too eager computation of the elements of the tuple which might otherwise have caused an infinite loop during execution. Note that we have two more or less independent computations that both follow the tree structure, and a weak interaction, when passing the $tmin$ value back in the tree.

Repmim a la AG. The structure of *repmim* is similar to the structure required by a compiler. A compiler performs several computations over an *abstract syntax tree* (AST), for example for computing its type and code. This corresponds to the *Tree* structure used by *repmim* and the tupled results. In the context of attribute grammars the elements of this tuple are called *attribute's*. Occasionally the word *aspect* is used as well, but an aspect may also refer to a group of attributes associated with one particular feature of the AST, language or problem at hand.

Result elements are called *synthesized* attributes. On the other hand, a compiler may also require information that becomes available at higher nodes in an AST to be available at lower nodes in an AST. The m parameter passed to r in *repmim* is an example of this situation. In the context of attribute grammars this is called an *inherited attribute*.

Using AG notation we first define the AST corresponding to our problem (for which the complete compilable solution is given in Fig. 1):

DATA Tree

```
| Leaf int : {Int}
| Bin  lt  : Tree
      rt  : Tree
```

The **DATA** keyword is used to introduce the equivalent of Haskell's **data** type. A **DATA**(*node*) defines a *node* (*node*) (or *nonterminal*) of an AST. Its alternatives, enumerated one by one after the vertical bar |, are called *variants*, *productions*. The term *constructor* is occasionally used to stress the similarity with its Haskell counterpart. Each variant has members, called *children* if they refer to other nodes of the AST and *fields* otherwise. Each child and field has a name (before the colon) and a type (after the colon). The type may be either another **DATA** node (if a child) or a monomorphic Haskell type (if a field), delimited by curly braces. The curly braces may be omitted if the Haskell type is a single identifier. For example, the **DATA** definition for the *repmim* problem introduces a node (nonterminal) *Tree*, with variants (productions) *Leaf* and *Bin*. A *Bin* has children *lt* and *rt* of type *Tree*. A *Leaf* has no children but contains only a field *int* holding a Haskell *Int* value.