

HEAD-ORDER TECHNIQUES AND OTHER PRAGMATICS OF LAMBDA CALCULUS GRAPH REDUCTION

Nikos B. Troullinos

DISSERTATION.COM



Boca Raton

Head-Order Techniques and Other Pragmatics of Lambda Calculus Graph Reduction

Copyright © 1993 Nikos B. Troullinos

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher.

Dissertation.com
Boca Raton, Florida
USA • 2011

ISBN-10: 1-61233-757-0
ISBN-13: 978-1-61233-757-9

Cover image © Michael Travers/Cutcaster.com

Abstract

In this dissertation Lambda Calculus reduction is studied as a means of improving the support for declarative computing. We consider systems having reduction semantics; i.e., systems in which computations consist of equivalence-preserving transformations between expressions. The approach becomes possible by reducing expressions beyond weak normal form, allowing expression-level output values, and avoiding compilation-centered transformations. In particular, we develop reduction algorithms which, although not optimal, are highly efficient.

A minimal linear notation for lambda expressions and for certain runtime structures is introduced for explaining operational features. This notation is related to recent theories which formalize the notion of substitution.

Our main reduction technique is Berkling's Head Order Reduction (HOR), a delayed substitution algorithm which emphasizes the extended left spine. HOR uses the de Bruijn representation for variables and a mechanism for artificially binding relatively free variables. HOR produces a lazy variant of the head normal form, the natural midway point of reduction. It is shown that beta reduction in the scope of relative free variables is not hard. Full normalization suggests new applications by not relegating partial evaluation to a meta level. Variations of HOR are presented, including a conservative "breadth-first" one which takes advantage of the inherent parallelism of the head normal form.

A reduction system must be capable of sharing intermediate results. Sharing under HOR has not received attention to date. In this dissertation variations of HOR which achieve sharing are described. Sharing is made possible via the special treatment of expressions referred to by head variables. The reduction strategy is based on normal order, achieves low reduction counts, but is shown to be incomplete.

Head Order Reduction with and without sharing, as well as other competing algorithms are evaluated on several test sets. Our results indicate that reduction rates in excess of one million reductions/second can be achieved on current processors in interpretive mode and with minimal pre- and post-processing. By extending the efficient algorithms for the pure calculus presented in this dissertation with primitives and data structures it is now possible to build useful reduction systems. We present some suggestions on how such systems can be designed.

Dedication

Στους γονείς μου, Μαίρη και Βασίλη

Preface

Foremost among all my teachers, Klaus Berkling taught me to question every assumption and not to rest until all aspects of a problem have been considered. For the last five years he has guided me through the maze of the practical aspects of Lambda Calculus reduction. With his judgment he reinforced my conviction that minimality and elegance are not luxuries but bare necessities.

Improving the practicality of Michael Hilton's *THOR/HORSE* was a strong influence and a well-defined first goal. It turned out that the problems are of a greater scope than I anticipated. Our interaction during the early stages of my study is greatly appreciated.

Inspiration, and the incentive for edging a bit closer to the "science" of computer science, was provided by the classes, seminars and talks that I was exposed to during my long career as a graduate student at Syracuse University. Ken Bowen, Lockwood Morris, John Reynolds, Alan Robinson, Ernest Sibert and Ed Storm come to mind, among many others. Academic committee members, Shiu-Kai Chin and Peter O'Hearn provided valuable feedback at every step.

Chuck Stormon and Coherent Research, Inc. afforded me the uncommon opportunity of continuing my graduate work while helping forge the future of this remarkable company. All my colleagues at CRi are to be thanked for their continuous interest, encouragement and good cheer.

Bruce Berra, Gideon Frieder and Brad Strait kindly provided financial assistance during my days at ECE, CIS and at the CASE Center. Their support is gratefully acknowledged.

From one-third across the globe away, loving parents Mairi and Vasilis were a vital source of encouragement, support and patience; enterprising sister Maria ensured that I stayed focused and balanced most of the time; Liebling Susanne E.(mail) Bohl was instrumental in making the journey more enjoyable. Lightheartedly, some credit also goes to a certain telecommunications company whose tariffs made it a bit more affordable to live and work in one continent while having so many of one's "friends & family" in another.

Table of Contents

Preface	vii
Table of Contents	ix
List of Figures	xiv
List of Tables	xv
1 LAMBDA CALCULUS AND REDUCTION COMPUTING	1
1.1 Introduction - Motivation	1
1.2 Operational Lambda Calculus as Computational/Machine Model	3
Lambda Calculus as Basis of a Reduction System	4
An Example of Reduction Semantics	5
Sharing is Sine Qua Non.....	6
1.3 Characteristics of Reduction Computing	7
Example of Reduction Programming	8
1.4 Direct Lambda Calculus Reduction	10
1.5 Applied Reduction Systems	12
1.6 Applications of Reduction Computing	15
Process Formalization	15
1.7 Results - Tour	16
Highlights and Results	16
Historic Lineage	18
Reducer Classification	21
Chapter Tour	21
Remarks	22
2 SYNTAX, NOTATION AND REPRESENTATIONS.....	25
2.1 Syntax, Semantics and Pragmatics	25
Syntax	25
Reduction	26

Normal Forms	28
Delayed Substitutions	29
2.2 Abstract Reducer Representations	29
Lambda Expression Representations	30
Generic Structural Tagging	30
Suspensions and Linking of Environment Elements	32
An Example of two Representations	33
2.3 BTF - A Generic Operational Notation	36
Syntax for Lambda Expressions	36
Simplification Rules	37
Reduction Example	37
Environments and Suspensions	40
Unbound Variable Counts	41
Abstraction-Depth Context Tagging	42
BTF Syntax Summary	42
2.4 Translations to	43
h2b: HOR to BTF	44
r2b: REP to BTF	44
3 GRAPH REDUCTION ALGORITHMS	47
3.1 Introduction	47
3.2 McCarthy's Eval -Apply	48
Essence of Graph Reduction	51
Style of Concrete Definitions	51
3.3 Landin's SECD machine	52
3.4 Wadsworth Reduction	55
RTNF/RTLF strategy	55
Copying of Operator Graphs	56
3.5 Aiello-Prini Environment Graph Reduction	59
AP Reduction Narrative	59
A Compact AP Machine	60
AP Reduction Example	61
3.6 Berking's Head Order Reduction	62
HOR in-the-large	64
Example of in-the-large rules	66
Head Normal Forms	67
Meaning of Suspensions and Environments	67

Weak Head-Order Reduction	70
Strong Head-Order Reduction	71
HOR Abstract Machine	72
HOR Narrative	73
4 FACETS OF HEAD ORDER REDUCTION	77
4.1 Head-Normal Reduction	77
Depictions of Lazy Head Normal Forms	78
4.2 Assurances of Correctness for HOR	81
4.3 Single Variable Arguments	82
4.4 Breadth-First HOR Strategy	84
Breadth-First Traversal	85
Successive hnfs under BF-HOR	86
4.5 Incremental Reduction and Expression Navigation	90
Controlling Reduction with a Quantum	90
Expression Navigation	93
4.6 RTNF/RTLF strategy with HOR machinery	94
HOR strategy under a REP-style	96
5 RELATED THEORIES	97
5.1 de Bruijn's Lambda Calculus	97
5.2 Explicit Substitutions - ACCL	100
5.3 Example and Comparison with HOR	106
6 SHARING STRATEGIES	109
6.1 General	109
6.2 Sharing in RTNF/RTLF Reducers	111
6.3 Head Order Reduction with Sharing	113
Sharing under Weak HOR	113
Sharing under Strong HOR	115

Abstract Machine for HOR with Sharing	115
Role of the "Unbound Variable Counts"	117
Meaning of a "Dummy" Environment	117
6.4 Variations of HOR	120
HOR-SH with Early VA	120
Shared Expressions as Operands	122
Tabulating the Variations of HOR	123
6.5 Incompleteness of HOR-SH	125
7 COMPARATIVE ANALYSIS AND EVALUATION	127
7.1 General	127
7.2 Comparison of Weak Reducers	129
7.3 Reduction Counts and Costs	131
Adopting a Cost Metric	131
Reduction Counts	132
Reduction Costs	135
Reducing/Checking of Normal Forms	139
Problem/Result Cost Ratios	141
Church Numeral Predecessor Tests	142
7.4 Evaluation of Variations of HOR	147
7.5 Tests of Recursion & Sharing	150
7.6 Comparisons of HORSE and π -RED	153
7.7 Remarks	158
Reduction Rates	158
Value of a Reduction	158
Reduction Counts as Measure of Efficiency	159
Nature and Cost of Copying	160
Optimal = Efficient ?	161
8 TOWARDS PRACTICAL REDUCTION SYSTEMS	163
8.1 Active vs. Passive Graphs	163
Lambda Expressions as Instruction Streams	164
Linear Coding of Graphs	164
Naming Expressions	166
Linear Organization of the Environment	167

8.2 Low-level Procedural Casting of Pure HOR	169
8.3 Enriched Calculus	172
Firing of Primitives	173
In-Line Compilation	174
Recursion	174
8.4 Implications of BF-HOR for Parallel Graph Reduction	177
Causal interdependencies and parallelism	180
8.5 Final Remarks	180
Future Work	180
Conclusions-Summary	181

A PPENDICES..... **185**

A. Abstract Reducers in Nearly Pure Lisp	185
B. Reduction Traces/Graphs	195
C. The BTRD Experimentation System	205
Highlights	205
Menus	209
Internal Structures	213
BETAR3 Example Trace	215
D. Test Expressions and Runtime Statistics	219
Bibliography	223
Index	227
Biographical Data	230

List of Figures

Figure 1.1	Combinator reduction contrasted to direct lambda reduction.	11
Figure 1.2	Contributions to the theory and the practice of Lambda Calculus Reduction.	19
Figure 2.1	Representation of hnf of test expression $tw7$	34
Figure 2.2	Four representations of $Y = (\lambda x.(\lambda z.zz)(\lambda y.x(yy)))$, with corresponding linear forms.	35
Figure 2.3	Normal order sequence of $((-2\ 0)(-2\ 1\ 0)(-2\ 1\ (1\ 0))(-2\ 1\ (1\ (1\ 0))))$ in tree form.	39
Figure 2.4	The translation function $h2b$ from HOR graph representations to BTF.	43
Figure 2.5	The translation function $r2b$ from REP-style graph representations to BTF.	44
Figure 3.1	Result graph of EV-AP reduction with pointer variables $v1, v2$ for nodes #1 and #2.	51
Figure 3.2	McCarthy's Eval-apply Interpreter as Lambda Calculus Reducer	52
Figure 3.3	Wadsworth's beta contraction via pointer manipulation.	56
Figure 3.4	Non-admissible graph for $(\lambda x.bx (\lambda z.za))(\lambda x.bx (\lambda z.za))(\lambda z.za)$	57
Figure 3.5	Full copying of the operator subgraph.	57
Figure 3.6	Partial copying of the operator subgraph.	58
Figure 3.7	Compact RTNF/RTL normal order reducer with a "force" flag.	61
Figure 3.8	Abstraction-application and application-abstraction corners	63
Figure 3.9	Left Spine tagged with a count of binding levels.	63
Figure 3.10	Symbolic depiction of the rules of HOR in-the-large.	65
Figure 3.11	Picture of a Head Normal Form.	68
Figure 3.12	Meaning of a suspension as a Lambda Calculus wfe.	68
Figure 3.13	Structure of an operational hnf with the environment represented via recursively "pushed" abstraction-application triangles.	69
Figure 4.1	Operational hnf of $((-1\ 0\ 0)\ (-2\ (-1\ 0\ (1\ 0))\ ((-1\ 0\ 1\ 0)\ 1)))$ via recursively nested environment triangles.	79
Figure 4.2	Abstract machine DAG of the previous operational hnf.	80
Figure 4.3	Calculating the correct dB index of an uninstantiated variable.	82
Figure 4.4	Order of traversal for BF-HOR.	86

Figure 4.5 Tree form of $tw10 = (-2 (-1 (-1 0 (0 1))(-1 2 0 1))(-1 0 2))$	87
Figure 4.6 Sequence of ohnfs under BF-HOR for $(-2 (-1 (-1 0 (0 1)) (-1 2 0 1)) (-1 0 2))$	88
Figure 4.7 Memory and time for the all* suite under four pure normal-order reducers.	95
Figure 5.1 Example expression Ω in de Bruijn and “raw” BTF tree forms.	98
Figure 6.1 Shared open subterm in different abstraction contexts.	119
Figure 7.1 Memory vs. time plots of three abstract reducers (all* tests).	131
Figure 7.2 Memory vs. time plots of the BETAR3 simulator (all* tests).	132
Figure 7.3 Reduction counts for normal-order and the two sharing reducers for the all* test suite.	134
Figure 7.4 Reduction costs for three non-sharing and the two sharing reducers for the all* test suite.	136
Figure 7.5 Cost and reduction ratios of AP-SH vs. HOR-SH for the all* test suite.	138
Figure 7.6 Cost ratios via AP-SH-VA and HOR-SH -VA reduction for problem-result pairs.	140
Figure 7.7 Ratios cost of “reduction” vs. nf checking for the all* test suite under BF-HOR, HOR and AP reduction.	141
Figure 7.8 Reduction costs for the pred* test suite on a logarithmic scale.	143
Figure 7.9 Ratios of AP-SH/HOR-SH for the pred* test suite.	144
Figure 7.10 Ratios of HOR/HOR-SH for the pred* test suite.	144
Figure 7.11 Memory, time and reductions for the pred* test suite.	146
Figure 7.12 Costs of reducing the all* suite under twelve variations of HOR.	149
Figure 7.13 BETAR3 reduction counts for the tfac* test suite.	151
Figure 7.14 BETAR3 times including a cyclic Y rule for the tfac* test suite.	151
Figure 7.15 Detail of BETAR3 times for the tfac* test suite.	152
Figure 7.16 BETAR3 times for the tfib* test suite on a logarithmic scale.	152
Figure 7.17 Expression $fsa2 = ((-1 0 0)(-1 (-1 0 0)(-1 1 (0 (-1 0))))))$ in tree form.	154
Figure 8.1 Organization of memory in an “active instruction” implementation of HOR. ..	165
Figure 8.2 Linear memory mapping of $(\lambda xy.xx)(\lambda uw.uw)(\lambda i.i)$ in THOR/HORSE.	166
Figure 8.3 The μ -RED low-level pure HOR machine with linear-memory.	171
Figure 8.4 Process queues under BF-HOR for five test expressions	179
Figure A.1 (EV-AP). McCarthy’s Eval-Apply as Lambda Calculus Reducer	185
Figure A.2 (SECD-V). Landin’s SECD Applicative Order Reducer.	186
Figure A.3 (HOR-NS) Normal-Form Reducer without an explicit stack	186
Figure A.4 (AP) Aiello and Prini’s RTNF/RTL Reducer	187
Figure A.5 (AP-SH-VA) RTNF/RTL Reducer with sharing, early VA and force flags.	188
Figure A.6 (W-HOR-VA) Weak HOR Reducer	188

Figure A.7 (AP-HOR) Reducer with RTNF/RTL strategy and dB/UVC machinery including sharing and early VA.	189
Figure A.8 (W-HOR-SH-VA) Sharing version of Weak HOR Reducer	189
Figure A.9 (HOR-VA) Berkling's Head Order Reducer	190
Figure A.10 (HOR-AP) Head Order Reducer	191
Figure A.11 (HOR-SH-VA) Head-Order Reducer with sharing	192
Figure A.12 (BF-HOR) Breadth-First HOR Reducer	193
Figure A.13 Code fragments defining the variations of HOR.	194
Figure C.1 Dialog for selecting a BTRD reducer and its applicable options.	206
Figure C.2 Input (top) and Output & Listener (bottom) windows of BTRD.	206
Figure C.3 Explanation of the BTRD output window and listener format.	207
Figure C.4 Reduction Unit Display for the B3/B3-SH simulator/reducer.	214
Figure C.5 Areas of the B3/B3-SH simulator/reducer.	214

List of Tables

Table 1.1 In a reduction system arguments may be ground, unbound or even absent.	6
Table 1.2 Sieve of Eratosthenes. Normal-order reduction without sharing is unacceptable in a practical system.	7
Table 1.3 (Hilton) Example of the source-transformation capabilities of a reduction system.	9
Table 1.4 University of Kiel π -RED system on the pred7 test.	13
Table 1.5 A generalized selection combinator	14
Table 1.6 Classification of reduction algorithms/systems and historical summary.	20
Table 2.1 Several depictions of the lambda expression $(\lambda xy.x)(\lambda y.y)$	31
Table 2.2 Lambda expressions in classic, de Bruijn and BTF notation.	37
Table 2.3 Normal-order reduction of $(((-2\ 0)(-2\ 1\ 0)(-2\ 1\ (1\ 0))(-2\ 1\ (1\ (1\ 0))))$	38
Table 3.1 EV-AP reduction of $(((-2\ 1\ 1)(-2\ 1\ 0)(-1\ 0))$	49
Table 3.2 Eager reduction of $(((-2\ 1\ 1)(-2\ 1\ 0)(-1\ 0))$	50
Table 3.3 Abstract machine for the eager SECD reducer.	53
Table 3.4 SECD-V reduction of $(((-2\ 1\ 1)(-2\ 1\ 0)(-1\ 0))$	53
Table 3.5 Example of transforming an expression via the in-the-large rules.	67
Table 3.6 Weak Head Order Reduction Abstract Machine.	70
Table 3.7 W-HOR reduction of $(((-2\ 1\ 1)(-2\ 1\ 0)(-1\ 0))$	71
Table 3.8 Head Order Reduction Abstract Machine.	72
Table 3.9 The syntactic elements of the pure HOR Abstract Machine.	73
Table 3.10 HOR reduction of $wads45 = (-2\ (-1\ 0\ (0\ 2))((-2\ 0\ 1\ 0))$	75
Table 4.1 Head-Normal HOR Abstract Machine.	78
Table 4.2 Additional rules of early VA lookup for W-HOR.	83
Table 4.3 Additional rules of early VA lookup for HOR.	83
Table 4.4 Head-normal reduction of e to its head normal form, h.	84
Table 4.5 Normal forms of the arguments emanating from the spine of h.	84
Table 4.6 Overall normal form of $e = ((-1\ 0\ 0)(-2\ (-1\ 0\ (1\ 0))(-1\ 0\ 1\ 0\ 1)))$	84
Table 4.7 Augmentation of the HOR Abstract Machine with a reduction quantum.	90
Table 4.8 Quantum-controlled HOR reduction of $(((-1\ 0\ 0)((-2\ 1\ 1)(-1\ 0)(-1\ 0)))$	91

Table 4.9 Quantum-controlled HOR-SH reduction of $((-1\ 0\ 0)((-2\ 1\ 1)(-1\ 0)(-1\ 0)))$	91
Table 4.10 Quantum-controlled HOR-SH-VA reduction of $((-1\ 0\ 0)((-2\ 1\ 1)(-1\ 0)(-1\ 0)))$	92
Table 4.11 Quantum-controlled reductions of $(-2\ (-1\ 0\ (0\ 2))((-2\ 0\ 1)\ 0))$ under AP.	92
Table 4.12 Equivalence-preserving lambda expression navigation.	93
Table 4.13 Definition of the effect of a move of focus within an expression.	93
Table 5.1 Substitution operators of the $\lambda\sigma$ -calculus	101
Table 5.2 (Abadi et al) The axioms of the $\lambda\sigma$ -calculus	103
Table 5.3 (Abadi et al) Untyped $\lambda\sigma$ -calculus \xrightarrow{wh}	105
Table 5.4 Head Order Reduction abstract machine.	105
Table 5.5 (Field) The axioms of ΛCCL	106
Table 5.6 Non-deterministic reduction rules for Head Order Reduction.	107
Table 5.7 HOR reduction of $(-1\ 1\ (-1\ 0)\ 0)$	107
Table 5.8 AP reduction of $(-1\ 1\ (-1\ 0)\ 0)$	107
Table 6.1 AP-SH reduction of $\text{wads45} = (-2\ (-1\ 0\ (0\ 2))((-2\ 0\ 1)\ 0))$	112
Table 6.2 W-HOR-SH reduction of $((-2\ 1\ 1)(-2\ 1\ 0)(-1\ 0))$	114
Table 6.3 W-HOR-SH-VA reduction of $((-1\ 0\ 0)((-2\ 1\ 1)(-1\ 0)(-1\ 0)))$	114
Table 6.4 Rules which add sharing in the HOR Abstract Machine	116
Table 6.5 Using the HOR machinery for correcting artificially closed terms.	118
Table 6.6 Equivalent literal substitutions under HOR-SH.	119
Table 6.7 Reduction of $\text{wads34} = (-2\ (-1\ 0\ 2\ 0)\ (-1\ 0\ ((-1\ 0)\ 1)))$ under HOR-SH-VA.	120
Table 6.8 Description of possible variations of HOR.	123
Table 6.9 Incompleteness of AOR and HOR-SH strategies.	125
Table 7.1 Results of whnf reduction for pred7	129
Table 7.2 Operational whnf results of the pred2 test.	130
Table 7.3 Results of nf reduction with sharing for pred7	130
Table 7.4 Results of nf reduction without sharing for pred7	130
Table 7.5 Results of nf reduction with sharing for pred14	131
Table 7.6 Total reduction counts for all* with and without early VA.	133
Table 7.7 Reduction counts for pred1-16	143
Table 7.8 Definition of variations of HOR.	148
Table 7.9 Cost statistics of fourteen combinations of HOR on the all* suite.	148
Table 7.10 Definition of the fsa^* expression family.	154
Table 7.11 Definition and run of cheq , a lambda-defined equality test for Church Numerals.	155
Table 7.12 π -REDs/HORSE. Mutual application of Church Numerals.	155
Table 7.13 π -REDs/HORSE. Predecessor function on Church Numerals.	156

Table 7.14 Family with infinitely increasing μ_β/μ_π of Frandsen & Sturtivant.	156
Table 7.15 π -RED*, AP-SH and HOR-SH. Family with infinitely increasing μ_β/μ_π	157
Table 7.16 Comparison of HORSE with the abstract reducers on the equality test on Church Numerals.	157
Table 7.17 Effect of instrumentation on the performance of the experimental reducers. ...	159
Table 8.1 Description of the registers of the μ RED machine.	170
Table 8.2 Lambda Calculus definitions	172
Table 8.3 Lists can be encoded as nested head normal forms.	173
Table 8.4 $(b + \sqrt{b^2 - 4ac})/2a$ under generic and in-line instruction streams.	175
Table 8.5 Tree and linear graph of fact $\equiv \lambda n. \text{if } (= n 0) 1 (* n (\text{fact } (- n 1)))$	176
Table 8.6 THOR definitions of four versions of the factorial function.	176
Table 8.7 HORSE statistics of fact 3000 under the four definitions of factorial.	176
Table B.1 Annotated reduction sequence for pred2	195
Table B.2 Test expression #1 under HOR-VA	196
Table B.3 W-HOR-VA reduction of $((-1 0 0)((-2 1 1)(-1 0)(-1 0)))$	196
Table B.4 Test expression #1 under HOR-SH reduction without Early VA.	197
Table B.5 AP reduction of wads45 = $(-2 (-1 0 (0 2))((-2 0 1) 0))$	198
Table B.6 HOR-SH (no VA) reduction of wads45 = $(-2 (-1 0 (0 2))((-2 0 1) 0))$	199
Table B.7 AP-SH-VA reduction of wads34 = $(-2 ((-1 0 2 0)(-1 0 ((-1 0) 1))))$	200
Table B.8 Graph EV-AP reduction of $((-2 1 1)(-2 1 0)(-1 0))$	201
Table B.9 HN-HOR trace for the example of section 4.1.	202
Table B.10 BF-HOR trace of tw10, section 4.4, in raw graph form.	203
Table B.11 Graph, BTF and Left Spine depiction of the tw7-2 ohnf.	204
Table B.12 Graph and BTF depictions of the ohnf of tw7-2 with late and early VA lookups.	204
Table C.1 Runtime statistics fields of the B3 simulator of BTRD.	207
Table C.2 Definition of a parameterized selector in the syntax of BETAR3.	208
Table C.3 Possible entries in the B3 environment arrays E1 and E2.	215
Table C.4 HOR-VA trace for the expression hz1 = $(-3 (-1 (-1 (-1 1 4) 2)(-1 1 0)) 2)$	215
Table D.1 The general-purpose test expression set all*	219
Table D.2 HOR-SH-VA reduction of the all* test set.	221

Chapter 1

Lambda Calculus and Reduction Computing

1.1 INTRODUCTION - MOTIVATION

An intriguing direction in symbolic computing is to base an architecture (hardware or software) on equivalence preserving transformations of expressions. In this dissertation we show that this viewpoint can be made practical. The computational model is the Lambda Calculus—a simple model which views functions as rules. Our motivation is the belief that direct Lambda Calculus reduction can be done efficiently. Beta reduction in the scope of relative free variables is not difficult with proper representations. Strong normalization¹ opens up new applications by not relegating abstract interpretation to a meta level.

Berklings Head Order Reduction (HOR) is the main tool of this approach. It is a delayed substitution algorithm with emphasis on the extended left spine. It uses the *de Bruijn* coding for variables, *unbound variable counts* for “neutralizing” relatively free variables and proceeds to an operational variant of the natural midway point of reduction, the *head normal form*. The machinery of Head Order Reduction appears to be a direct and minimal match to the requirements of strong reduction. Among its operational properties are extremely low space requirements and independence from pointer representations for variables. Michael Hilton’s *THOR/HORSE*, an applied derivative system, is a good example of the overall direction and goals.

¹ In this thesis this term simply means that reduction proceeds to full β -normal form.

THOR/HORSE has demonstrated that systems based on the preceding principles are not only feasible but can also attain levels of efficiency previously thought unreachable. Hence, we downplay transformation techniques like combinators and strong typing that are commonly employed solely for the purpose of compilation to von Neumann code. It is our belief that specialized hardware or a low-level software “interpreter” for a proper operational casting of the Lambda Calculus can serve large classes of problems well, both in terms of expressiveness and efficiency.

But a pure normal-order strategy like that of *THOR/HORSE* is unacceptable for practical work if it is not augmented with sharing. A large part of our efforts concentrates on augmenting HOR with sharing. Initially, the task seemed intractable; early reducers were prone to subtle binding errors. As our understanding improved, we realized that no single technique is a complete solution to the problem and free of disadvantages. But progress has been made. By utilizing combinations of essential techniques provided by Berkling’s earlier work the state of affairs described in this dissertation was reached. The initial goal of adding effective full sharing in a system having the scope of *THOR/HORSE* is unfulfilled.

The essentials of Head Order Reduction were implicitly present as early as 1972 in de Bruijn’s *Indagationes Mathematicae* publication of a Lambda Calculus with integer “Nameless Dummies” instead of the usual named variable identifiers. Retrospectively, one can state that the logical progression that leads to the operational HOR is the addition of delayed substitutions via an environment and the emphasis on head normal forms. Historically, the development took a different course via Berkling’s “in-the-large” reduction rules presented for the first time at the *Santa Fe Graph Reduction Workshop* in 1986.

The quest for grasping the interrelationships and intricacies of reduction methods led us to review several of the well-known weak and strong algorithms. Our attention was limited to the pure, untyped Lambda Calculus to keep complications to a minimum. It turned out that the weak form of HOR is identical in strategy with the so called “RTL” strategy of Wadsworth and Aiello–Prini and other less well known algorithms. Detailed comparisons of all the popular algorithms demonstrate the directness and efficiency of HOR.

Our work represents a shift in the point of view of traditional research on functional programming. It is based on the realization that the Lambda Calculus, which is commonly regarded as one of the foundations of programming, is not intrinsically oriented towards producing a value. Therefore, not all languages and systems should be limited in this regard. We believe that much may be gained in the realm of software technology if the first class objects are full expressions rather than ground values of some output domain.

We attempt to derive architectural principles while remaining close to the inherent mechanisms suggested by the underlying theory. The sort of computations that benefit from the existence of a reduction system are those in which one is not only interested in getting a

final ground result but also insists that the steps taken in arriving at such result are as transparent and direct as possible. If this line of thinking is followed, then the Lambda Calculus appears indeed to be a solid choice because most of the thorny issues surrounding declarative computation of the functional kind surface here in their purest form. How practical or influential this path may be is something that will be determined eventually but for the moment it seems exciting work to pursue.

Combinator techniques are sidestepped because we regard them as indirect. With combinators one forgoes the random reference ability of variables and replaces it by argument propagation systems. From an operational standpoint this is cumbersome and circuitous. In addition, such techniques make it very difficult to do *incremental* reduction and to produce recognizable intermediate results. The efficiency of reducers like HOR may be largely attributed to using binding indices for variables together with a correction mechanism for dynamically-created abstractions that operates in an incremental and local manner. This way one can avoid the problems of scope and name capturing of variables and thus render the detour via combinators much less attractive.

Our overall emphasis is, by necessity, the invention of reduction schemes which are economical from a practical rather than a theoretical standpoint. Hence, we reject the usual notion of optimality—which equates excess work to excess reductions—as too limiting and not very helpful for reduction system design. The Lambda Calculus can model situations of arbitrary complexity. A reduction strategy which always avoids excess reductions requires bookkeeping which is more substantial than the effort required for performing “excess” reductions. A more reasonable goal is to attempt to match or exceed the natural savings of applicative order while avoiding, if possible, its incompleteness. Opportunities for parallelism are plentiful, but again, the orchestration overhead can outweigh the benefits derived over simpler sequential tactics. Source program transformations can help in exposing the inherent parallelism and research in this area seems to us worthwhile.

The rest of this chapter highlights the reduction-based viewpoint and pinpoints the challenges that are present. Results are summarized in section 1.7.

1.2 OPERATIONAL LAMBDA CALCULUS AS COMPUTATIONAL/MACHINE MODEL

Attempting to derive architectural principles from the operational characteristics of Lambda Calculus reduction is a slightly unconventional line of research. It could be likened to an aspiration to advance the techniques of modern microprocessor design by examining closely the workings of a Turing Machine. But the parallel breaks down—and hence the former does not seem as unconvincing as the latter—if one considers that a Lambda Calculus expression is a higher level of expression than, let us improvise, the transitions rules of a

Turing Machine and that it should impart more structure to a mechanization procedure. Additionally, the declarative nature of the Lambda Calculus is in sharp contrast with the sequential instructions and state-oriented nature of the von Neumann model. Consequently, since the semantic gap to be covered is larger, the journey could be more rewarding.

Lambda Calculus as Basis of a Reduction System

Our interest in the Lambda Calculus stems from two basic premises: First, research efforts on making high-order declarative languages work better and serve larger classes of practical computing tasks should intensify². Second, a significant portion of our current difficulties with the building of software is attributed at to the low-level nature of the popular programming language and environments and the batch style of most of real world development tools. Higher-order programming, symbolic computation, partial evaluation and incremental development hold serious promises in advancing common computing practice³. Programming environments could become more fluid and the so called “Computer-Aided Software Engineering” tools would be able to focus more on the actual artifact.

One could regard computation as proceeding under two rather distinct flavors. First, functions are applied to arguments in a pure sense by fusing them together which means propagating arguments to the proper places at the appropriate time. Second, ground data objects, like encodings for numbers, characters or graphics primitives, are transformed in well-defined and predesigned ways; for example, multiplying two integers together or shifting the value of a register by a fixed amount. Symbolic computation often has a stronger first component while “number crunching” is nearly a pure form of the second. The Lambda Calculus is one of the simplest settings where one can experience these two flavors in a distilled manner and offers the opportunity of discovering techniques that may be useful in more general settings. The considered application areas are reduction languages, semantics-directed machine architecture and high-level software development tools.

One cannot avoid touching upon the potential criticism that the Lambda Calculus is inherently a model of *sequential* computing and hence efforts based on it may not be as applicable to future systems⁴. Exploring concurrent computation in a similar spirit, starting

² We are particularly interested in function-based languages. But efficient Lambda Calculus reduction can have implications for languages with side effects due to the generality of the concepts of abstraction and application.

³ The literature has a wealth of research on each of these topics and it is outside of the scope of this thesis to attempt further explications. For high-order functional programming see, for instance, the paper by J. Hughes [Hughes 90] and the proposal by J. Backus et al [Backus 90].

⁴ The major difficulty is that from a foundational viewpoint the pure calculus can not “escape” an undefined argument of a disjunction; i.e., there is no mechanism for expressing

maybe with Milner's *Calculus of Communicating Systems (CCS)* [Milner 80] or Hoare's *Communicating Sequential Processes (CSP)* [Hoare 85] or some other concurrent process modeling theory, would have been closer to recent trends. The author feels though, that progress remains to be made at the level of sequential expression. As it will become evident in section 8.4, the core of the Head-Order strategy has important implications for automatic *run-time* parallelism given the structural invariance of the operational head normal forms.

An Example of Reduction Semantics

A reduction system is envisioned as a delivery platform for exploratory programming environments of declarative or well-structured imperative nature. Solid representation choices and direct support from a specialized instruction set can result in a high-level organization, compared to the von Neumann one. Compilation is limited to optimizing but equivalence-preserving transformations on portions of a program that consist of arithmetic, logical and other low-level hardware-supported operations.

To demonstrate this sort of system let us visit a well-known high-order function like *reduce*. In the syntax of *THOR*, *reduce* has a definition identical to that of a *Scheme* function. It also behaves identically to a *Scheme* function as seen in the first and second examples of Table 1.1. But a vital, and commonly considered prohibitively expensive, change of semantics is evident in the third and fourth runs. When *reduce* is supplied with fewer arguments than those implied by its definition or with enough but unbound ones, *HORSE* (the programming environment for *THOR*) is perfectly content to perform the available reductions without any of the usual error messages. This is akin to abstract interpretation but it is accomplished with nothing more than the ordinary reduction mechanisms of the underlying calculus; there is nothing "abstract" or "meta" about such behavior—it is simply that the rules of the underlying calculus are followed to a greater extent.

A major theme of our work is that the cost of such fluidity is not as high as commonly thought; the proviso is, of course, that efficient reduction algorithms and systems are available. Fundamentally, not much is new or radical with this approach. The novelty is the recognition of the importance of *expression-based* computation which advances beyond weak normalization.

the meaning of a "parallel or". Applied variants of the pure calculus like *THOR* can and do include such mechanisms by treating, for instance, truth values specially. At this level there is no such constraint; in fact, the side-effect free nature of the calculus makes it relatively easy to add parallelism. Graph transformations, the natural mechanism for reduction, do not impose any sequentiality by themselves.