

# Applying Functional Programming Theory to the Design of Workflow Engines

Peter M. Kelly

January 2011

A dissertation submitted to the School of Computer Science of  
The University of Adelaide for the degree of Doctor of Philosophy

Supervisors:

Dr. Paul D. Coddington

Dr. Andrew L. Wendelborn



# Contents

<b>Abstract</b>	<b>11</b>
<b>Declaration</b>	<b>12</b>
<b>Acknowledgements</b>	<b>13</b>
<b>Related publications</b>	<b>14</b>
<b>1 Introduction</b>	<b>17</b>
1.1 Case study: Image similarity search . . . . .	18
1.2 Language requirements . . . . .	18
1.3 Programming models for workflows . . . . .	19
1.4 Functional programming . . . . .	20
1.5 Thesis overview . . . . .	21
1.6 Thesis outline . . . . .	23
<b>2 Background and Related Work</b>	<b>25</b>
2.1 Independent tasks . . . . .	27
2.2 Fixed task structure . . . . .	30
2.3 Workflows . . . . .	31
2.3.1 Types of workflows . . . . .	33
2.3.2 Aims of scientific workflow languages . . . . .	34
2.3.3 Programming models . . . . .	35
2.3.4 Orchestration and choreography . . . . .	36
2.3.5 Error handling . . . . .	37
2.3.6 Execution management and monitoring . . . . .	37

2.3.7	Languages and systems . . . . .	38
2.3.8	Limitations . . . . .	48
2.4	Parallel functional programming . . . . .	50
2.4.1	Modes of evaluation . . . . .	51
2.4.2	Parallelism and efficiency . . . . .	52
2.4.3	Expressing parallelism . . . . .	53
2.4.4	Laziness and parallelism . . . . .	54
2.4.5	Graph reduction . . . . .	56
2.4.6	Input/Output . . . . .	60
2.4.7	Relationships between functional programming and workflows . . . . .	63
2.4.8	Discussion . . . . .	66
2.5	Web services . . . . .	67
2.5.1	Concepts . . . . .	68
2.5.2	Service oriented architecture . . . . .	69
2.5.3	Standards . . . . .	69
2.5.4	Language support . . . . .	70
2.5.5	Asynchronous RPC . . . . .	72
2.5.6	State . . . . .	73
2.5.7	REST . . . . .	74
2.5.8	Applicability to workflows . . . . .	75
2.6	XQuery . . . . .	76
2.6.1	Existing approaches to web services integration . . . . .	76
2.7	Summary . . . . .	78
<b>3</b>	<b>Workflow Model</b>	<b>81</b>
3.1	Overview of our approach . . . . .	82
3.2	Workflow language concepts . . . . .	83
3.3	Challenges addressed . . . . .	85
3.3.1	Expressiveness . . . . .	85
3.3.2	Control flow constructs . . . . .	86
3.3.3	Abstraction mechanisms . . . . .	86
3.3.4	Data manipulation . . . . .	87

3.3.5	Implicit parallelism . . . . .	88
3.3.6	Service access . . . . .	89
3.3.7	Distributed execution . . . . .	89
3.4	Orchestration and choreography . . . . .	90
3.4.1	Orchestration . . . . .	90
3.4.2	Choreography . . . . .	91
3.5	Extended lambda calculus . . . . .	93
3.5.1	Lambda calculus . . . . .	94
3.5.2	ELC . . . . .	95
3.5.3	Built-in functions . . . . .	96
3.5.4	Syntax . . . . .	99
3.5.5	Programming example . . . . .	100
3.5.6	Evaluation modes . . . . .	100
3.5.7	Parallelism . . . . .	101
3.5.8	Network access . . . . .	102
3.6	ELC as a workflow language . . . . .	104
3.6.1	Tasks . . . . .	105
3.6.2	Conditional branching and iteration . . . . .	107
3.6.3	Embedded scripting languages . . . . .	108
3.6.4	Representation . . . . .	109
3.6.5	Abstract workflows . . . . .	111
3.7	Example workflow: Image similarity search . . . . .	113
3.7.1	Workflow implementation . . . . .	114
3.7.2	Task implementation . . . . .	117
3.8	Support for higher-level languages . . . . .	118
3.9	Summary . . . . .	119
<b>4</b>	<b>The NReduce Virtual Machine</b>	<b>121</b>
4.1	Goals . . . . .	122
4.2	Execution environment . . . . .	124
4.2.1	Frames . . . . .	124
4.2.2	Abstract machine . . . . .	125

4.2.3	Data types and graph representation . . . . .	128
4.2.4	Pointers and numeric values . . . . .	129
4.3	Parallelism . . . . .	130
4.3.1	Automatic detection vs. manual specification . . . . .	131
4.3.2	Sparking . . . . .	132
4.4	Frame management . . . . .	133
4.4.1	Blocking and context switches . . . . .	135
4.5	Networking . . . . .	136
4.5.1	Abstracting state manipulation . . . . .	137
4.5.2	The I/O thread . . . . .	139
4.5.3	Asynchronous call example . . . . .	140
4.6	List representation . . . . .	141
4.6.1	Array cells . . . . .	142
4.6.2	Array references . . . . .	142
4.6.3	Sharing . . . . .	144
4.6.4	Converting between representations . . . . .	145
4.6.5	High-level list operations . . . . .	145
4.7	Distributed execution . . . . .	145
4.7.1	Architecture . . . . .	146
4.7.2	Message passing . . . . .	146
4.7.3	Distributed heap management . . . . .	148
4.7.4	Work distribution . . . . .	150
4.7.5	Frame migration . . . . .	151
4.8	Summary . . . . .	153
<b>5</b>	<b>XQuery and Web Services</b>	<b>155</b>
5.1	Benefits of XQuery as a workflow language . . . . .	156
5.2	Language overview . . . . .	157
5.2.1	Data types . . . . .	157
5.2.2	Expressions and control flow constructs . . . . .	158
5.2.3	Path expressions . . . . .	158
5.2.4	Element constructors . . . . .	159

5.2.5	Example . . . . .	159
5.2.6	Relationship to other languages . . . . .	160
5.3	Web service support . . . . .	161
5.3.1	The <code>import service</code> statement . . . . .	161
5.3.2	Mapping function calls to SOAP requests . . . . .	162
5.3.3	Stub functions . . . . .	163
5.3.4	Parallelism . . . . .	165
5.4	Usage examples . . . . .	165
5.4.1	Collating content from multiple sources . . . . .	166
5.4.2	Processing of data structures returned from services . . . . .	167
5.4.3	Two-dimensional parameter sweep with graphical plotting . . . . .	168
5.5	Implementation details . . . . .	172
5.5.1	Runtime library . . . . .	172
5.5.2	Data representation . . . . .	174
5.5.3	Static context . . . . .	175
5.5.4	Dynamic context . . . . .	176
5.5.5	Tree construction . . . . .	178
5.5.6	Compilation process . . . . .	178
5.6	Discussion . . . . .	179
5.6.1	Factors that made implementation easy . . . . .	180
5.6.2	Factors that made implementation difficult . . . . .	181
5.6.3	Performance implications . . . . .	181
5.6.4	Advantages over existing RPC mechanisms . . . . .	182
5.6.5	Advantages over existing workflow languages . . . . .	183
5.7	Summary . . . . .	184
<b>6</b>	<b>Managing Parallelism</b>	<b>185</b>
6.1	Evaluation modes and detection of parallelism . . . . .	186
6.1.1	Advantages of lazy evaluation . . . . .	187
6.1.2	Advantages of strict evaluation . . . . .	188
6.1.3	Detecting parallelism . . . . .	189
6.1.4	Comparing parallelism in both cases . . . . .	190

6.1.5	Laziness and service requests . . . . .	192
6.1.6	Summary . . . . .	194
6.2	Limiting request concurrency . . . . .	194
6.2.1	Available capacity . . . . .	195
6.2.2	TCP connection establishment . . . . .	197
6.2.3	Determining when a connection has been accepted . . . . .	198
6.2.4	Client-side connection management . . . . .	199
6.3	Spark pool management . . . . .	201
6.3.1	Costs of sparking . . . . .	202
6.3.2	Minimising direct costs . . . . .	203
6.4	Load balancing of requests during choreography . . . . .	204
6.4.1	Processing assigned sparks . . . . .	205
6.4.2	Determining when a machine is idle . . . . .	205
6.4.3	Deciding which sparks to run . . . . .	207
6.4.4	Postponing connections . . . . .	208
6.4.5	Spark addition and activation . . . . .	208
6.4.6	Example . . . . .	210
6.4.7	How many frames to distribute? . . . . .	212
6.5	Limits on connection rates . . . . .	214
6.6	Summary . . . . .	216
<b>7</b>	<b>Performance Evaluation</b>	<b>217</b>
7.1	Experimental setup . . . . .	218
7.2	Service invocation . . . . .	219
7.2.1	Number of tasks and task throughput . . . . .	219
7.2.2	Granularity . . . . .	221
7.3	Types of parallelism . . . . .	224
7.3.1	Data parallelism . . . . .	226
7.3.2	Nested data parallelism . . . . .	229
7.3.3	Divide and conquer . . . . .	232
7.3.4	Physical pipelining . . . . .	236
7.3.5	Logical pipelining . . . . .	240

7.3.6	Multi-stage data parallelism . . . . .	242
7.3.7	Summary . . . . .	247
7.4	Data transfer . . . . .	247
7.4.1	Divide and conquer . . . . .	248
7.4.2	Physical pipelining . . . . .	250
7.4.3	Logical pipelining . . . . .	251
7.4.4	Multi-stage data parallelism . . . . .	253
7.4.5	Discussion . . . . .	254
7.5	XQuery-based workflows . . . . .	255
7.5.1	Speedup and data transfer vs. number of nodes . . . . .	257
7.5.2	Speedup vs. granularity . . . . .	259
7.5.3	Data transfer vs. number of tests . . . . .	261
7.5.4	Impact of code size . . . . .	261
7.6	Internal computation . . . . .	263
7.6.1	ELC . . . . .	264
7.6.2	XQuery . . . . .	272
7.7	Summary . . . . .	275
<b>8</b>	<b>Summary and Conclusions</b>	<b>277</b>
8.1	Summary . . . . .	277
8.2	Review of case study: Image similarity search . . . . .	279
8.3	Contributions . . . . .	281
8.4	Future work . . . . .	283
8.4.1	Compiler optimisations . . . . .	283
8.4.2	XQuery . . . . .	284
8.4.3	Fault tolerance . . . . .	284
8.5	Conclusions . . . . .	285
<b>A</b>	<b>Lambda Calculus and Graph Reduction</b>	<b>287</b>
A.1	Reduction order . . . . .	288
A.2	Additional functionality . . . . .	289
A.3	Extensions . . . . .	291
A.4	Graph reduction . . . . .	291
A.5	Parallel evaluation . . . . .	292

<b>B XQuery Compilation</b>	<b>295</b>
B.1 Overview . . . . .	295
B.2 Format of compilation rules . . . . .	296
B.3 Initial compilation stages . . . . .	297
B.4 Code generation . . . . .	298
B.5 Final compilation stages . . . . .	300

## Abstract

Workflow languages are a form of high-level programming language designed for coordinating tasks implemented by different pieces of software, often executed across multiple computers using technologies such as web services. Advantages of workflow languages include automatic parallelisation, built-in support for accessing services, and simple programming models that abstract away many of the complexities associated with distributed and parallel programming. In this thesis, we focus on data-oriented workflow languages, in which all computation is free of side effects.

Despite their advantages, existing workflow languages sacrifice support for internal computation and data manipulation, in an attempt to provide programming models that are simple to understand and contain implicit parallelism. These limitations inconvenience users, who are forced to define additional components in separate scripting languages whenever they need to implement programming logic that cannot be expressed in the workflow language itself.

In this thesis, we propose the use of functional programming as a model for data-oriented workflow languages. Functional programming languages are both highly expressive and amenable to automatic parallelisation. Our approach combines the coordination facilities of workflow languages with the computation facilities of functional programming languages, allowing both aspects of a workflow to be expressed in the one language.

We have designed and implemented a small functional language called ELC, which extends lambda calculus with a minimal set of features necessary for practical implementation of workflows. ELC can either be used directly, or as a compilation target for other workflow languages. Using this approach, we developed a compiler for XQuery, extended with support for web services. XQuery's native support for XML processing makes it well-suited for manipulating the XML data produced and consumed by web services. Both languages make it easy to develop complex workflows involving arbitrary computation and data manipulation.

Our workflow engine, NReduce, uses parallel graph reduction to execute workflows. It supports both *orchestration*, where a central node coordinates all service invocation, and *choreography*, where coordination, scheduling, and data transfer are carried out in a decentralised manner across multiple nodes. The details of orchestration and choreography are abstracted away from the programmer by the workflow engine. In both cases, parallel invocation of services is managed in a completely automatic manner, without any explicit direction from the programmer.

Our study includes an in-depth analysis of performance issues of relevance to our approach. This includes a discussion of performance problems encountered during our implementation work, and an explanation of the solutions we have devised to these. Such issues are likely to be of relevance to others developing workflow engines based on a similar model. We also benchmark our system using a range of workflows, demonstrating high levels of performance and scalability.

## Declaration

This work contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution to Peter Kelly and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library catalogue, the Australasian Digital Theses Program (ADTP) and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Peter Kelly

## Acknowledgments

Firstly, I would like to express my sincere gratitude to my supervisors, Paul and Andrew, for their invaluable guidance and assistance with my work. Both have been excellent supervisors, and provided a great deal of insightful advice and feedback that has guided my project. Their unique combination of skills — in particular, Andrew’s expertise in programming language theory and parallel functional programming, and Paul’s background in high performance computing and scientific applications — has played a crucial role in the development of the key ideas behind my project. I am indebted to them for the many ways in which they have contributed to my development as a researcher.

Secondly, to those who have accompanied me on this journey in various ways — Alison, Asilah, Brad, Daniel, Donglai, Jane, Joanna, Joseph, Miranda, Nick, Nicole, Peter, Wei, Yan, and Yuval — thanks for making this whole process a more enjoyable one. I value the many opportunities we’ve had to exchange thoughts, ideas, and experiences. You have all been a great help to me in many ways, both personally and professionally.

Thirdly, many thanks go to my parents, who have done so much over the years to support and encourage me in my studies, and teach me the value of working hard to pursue my goals. To fully describe the ways in which they have helped me get to where I am today would take another three hundred pages. It is to them that I dedicate this thesis.

Finally, I am very grateful to the School of Computer Science for giving me the opportunity to pursue a PhD, and for providing me with financial support and an excellent working environment in which to carry out my studies. I would also like to acknowledge eResearch SA for providing me with access to the high performance computing facilities necessary to run my experiments.

## Related publications

1. Peter M. Kelly, Paul D. Coddington, and Andrew L. Wendelborn. Distributed, parallel web service orchestration using XSLT. In *1st IEEE International Conference on e-Science and Grid Computing*, Melbourne, Australia, December 2005.
2. Peter M. Kelly, Paul D. Coddington, and Andrew L. Wendelborn. A simplified approach to web service development. In *4th Australasian Symposium on Grid Computing and e-Research (AusGrid 2006)*, Hobart, Australia, January 2006.
3. Peter M. Kelly, Paul D. Coddington, and Andrew L. Wendelborn. Compilation of XSLT into dataflow graphs for web service composition. In *6th IEEE International Symposium on Cluster Computing and the Grid (CCGrid06)*, Singapore, May 2006.
4. Peter M. Kelly, Paul D. Coddington, and Andrew L. Wendelborn. A distributed virtual machine for parallel graph reduction. In *8th International Conference on Parallel and Distributed Computing Applications and Technologies (PDCAT '07)*, Adelaide, Australia, December 2007.
5. Peter M. Kelly, Paul D. Coddington, and Andrew L. Wendelborn. Lambda calculus as a workflow model. *Concurrency and Computation: Practice and Experience*, 21(16):1999–2017, July 2009.

All of the software developed in this project is available from the following URL, under the terms of the GNU Lesser General Public License:

<http://nreduce.sourceforge.net/>

For Mum & Dad