The Design and Construction of Deadlock-Free Concurrent Systems

Jeremy Malcolm Randolph Martin

Thesis submitted for the degree of D. Phil to the School of Sciences in the University of Buckingham 1996

Abstract
The Design and Construction of Deadlock-Free Concurrent Systems
Jeremy Martin

It is a difficult task to produce software which is guaranteed never to fail, but it is a vital goal for which to strive in many real-life situations. The problem is especially complex in the field of parallel programming, where there are extra things that can go wrong. A particularly serious problem is deadlock. Here we consider how to construct systems which are guaranteed deadlock-free by design.

Design rules, old and new, which eliminate deadlock are catalogued, and their theoretical foundation illuminated. Then the development of a software engineering tool is described which proves deadlock-freedom by verifying adherence to these methods. Use of this tool is illustrated with several case studies.

The thesis concludes with a discussion of related issues of parallel program reliability.

Acknowledgements

I am indebted to my supervisors Ian East and Sabah Jassim for their guidance, encouragement and enthusiasm for science. I have also benefited greatly from discussions with their former colleague John Rowe. My thesis is based largely on previous work and ideas of Bill Roscoe and Peter Welch, both of whom have been very helpful.

I am very grateful to my sister, Clare, who originally suggested to me the idea of studying for a doctorate and put me in touch with my supervisors. She also provided me with much useful background material. The University of Buckingham has proved a very pleasant environment where to work, with excellent facilities. I must also thank my employers, Oxford University Computing Services, for giving me time off to study.

Many thanks are due to my wife, Nathalie, who often found that although my body was present my mind was elsewhere. Thanks also to my children Adrian, Alex (who popped up half way through) and Clarisse (who popped up right at the end) for all the fun that we have had.

This thesis is dedicated, with love, to the memory of Phyllis Amy Martin, 1906 – 1994.

Contents

Introduction **1**

1 CSP and Deadlock **5**
Introduction . . . 5
1.1 The CSP Language . . . 6
1.2 The Failures-Divergences Model . . . 13
1.3 Operational Semantics . . . 19
1.4 Language Extensions . . . 24
1.5 Deadlock Analysis . . . 25

2 Design Rules for Deadlock Freedom **34**
Introduction . . . 34
2.1 Cyclic Processes . . . 35
2.2 Client-Server Protocol . . . 45
2.3 Resource Allocation Protocol . . . 55

3 A Tool for Proving Deadlock-Freedom **62**
Introduction . . . 62
3.1 Normal Form Transition Systems . . . 63
3.2 Deadlock Checker . . . 66
3.3 Checking Adherence to Design Rules . . . 70
3.4 Towards a General Purpose Algorithm . . . 92

4 Engineering Applications **106**
Introduction . . . 106
4.1 The occam Programming Language . . . 107
4.2 Case Studies . . . 108

Conclusions and Directions for Future Work **124**

References **130**

A Partial Orders **134**

B Graphs and Digraphs **136**

List of Figures

0.1 Deadlocked Dining Philosophers 2

1.1 Laws of CSP I 11
1.2 Laws of CSP II 12
1.3 Denotational Semantics for CSP I 17
1.4 Denotational Semantics for CSP II 18
1.5 State Transition Systems 19
1.6 Operational Semantics for CSP I 22
1.7 Operational Semantics for CSP II 23
1.8 Wait-for Digraphs 28

2.1 Networks of *I/O-SEQ* and *I/O-PAR* Processes 37
2.2 *LATCH*: a Composite *I/O-PAR* Process 39
2.3 Connection Digraph with Channel Labelling 44
2.4 Multi-phase Channel Labelling 46
2.5 Client-Server Digraph for *FARM* 50
2.6 Composite Client-Server Process 51
2.7 Client-Server Digraph and Exploded Client-Server Digraph 53
2.8 Adding Client-Server Connections 55
2.9 Connection Graph for Dining Philosophers 57
2.10 Arm-Wrestling Philosophers 60
2.11 Bank Database System 61

3.1 Transition System Resulting from Compilation 64
3.2 Pre-normalisation 65
3.3 Normal Form Transition System 66
3.4 Normal Form Transition Systems for Dining Philosophers 69
3.5 Normal Form Transition Systems for Two-Place Buffer 74
3.6 Normal Form Transition System for General Resource Process 76
3.7 Hasse Digraph and Normal Form Transition System for *CELL*$(0, 0)$. 81
3.8 Normal Form Transition System for *FOREMAN*(0) 84
3.9 Construction of SDD for Dining Philosophers 94
3.10 Client-Server Digraph for *CLOCK* Network 98

4.1 Labelled Connection Diagram for Laplace Solver 110
4.2 Cube Router 113
4.3 Routing Strategies for Ring and Grid 117
4.4 Connection Digraph for COMMANDER 121
4.5 Client-Server Digraph for Improved Design 123
4.6 CSP Toolkit – A Vision for the Future 129

B.1 A Graph 136

List of Tables

3.1 Machine Readable CSP 67

4.1 Relationship between occam and CSP 107

Declaration

I would like to draw attention to the following material contained within this thesis which I believe to be original.

Chapter 2: Theorems 7 and 9 are new results which generalise a theorem of A. W. Roscoe and N. Dathi and several theorems of P. H. Welch. Theorem 7 forms part of a joint publication:

J. Martin, I.East and S. Jassim *Design Rules for Deadlock-Freedom*, Transputer Communications, September 1994.

The definition of the Client-Server Protocol, and the results which follow, are a new formal adaptation of informal ideas due to Welch, G. R. R. Justo and C. J. Willcock. The Extended Resource Allocation Protocol (rule 11) is also new.

Chapter 3: Apart from the section which describes the normalisation of transition systems, this chapter is based entirely on original work.

Chapter 4: The first two case studies considered are original implementations of existing algorithms. The third is an original analysis of a published algorithm which reveals a deficiency and proposes a solution to this problem.

To the best of my knowledge, none of this material has ever previously been submitted for a degree in this or any other university.

.