

Lecture Notes in Artificial Intelligence

1630

Subseries of Lecture Notes in Computer Science

Edited by J. G. Carbonell and J. Siekmann

Lecture Notes in Computer Science

Edited by G. Goos, J. Hartmanis and J. van Leeuwen

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Matthew M. Huntbach Graem A. Ringwood

Agent-Oriented Programming

From Prolog to Guarded Definite Clauses



Springer

Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA
Jörg Siekmann, University of Saarland, Saarbrücken, Germany

Authors

Matthew M. Huntbach

Graem A. Ringwood

Department of Computer Science, Queen Mary and Westfield College

Mile End Road, London E1 4NS, UK

E-mail: {mmh/gar}@dcs.qmw.ac.uk

Cataloging-in-Publication data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Huntbach, Matthew M.:

Agent-oriented programming : from prolog to guarded definite
clauses / Matthew M. Huntbach ; Graem A. Ringwood. - Berlin ;
Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ;
Paris ; Singapore ; Tokyo : Springer, 1999

(Lecture notes in computer science ; Vol. 1630 : Lecture notes in
artificial intelligence)

ISBN 3-540-66683-4

CR Subject Classification (1998): I.2, D.1.6, D.1.3, D.3, C.2.4

ISBN 3-540-66683-4 Springer-Verlag 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-Verlag. Violations are liable for prosecution under the German Copyright Law.

© Springer-Verlag Berlin Heidelberg 1999

Printed in Germany

Typesetting: Camera-ready by author

SPIN 10703359

06/3142 – 5 4 3 2 1 0

Printed on acid-free paper

Preface

A book that furnishes no quotations is, me judice, no book – it is a plaything.

TL Peacock: Crochet Castle

The paradigm presented in this book is proposed as an agent programming language. The book charts the evolution of the language from Prolog to intelligent agents. To a large extent, intelligent agents rose to prominence in the mid-1990s because of the World Wide Web and an ill-structured network of multimedia information. Agent-oriented programming was a natural progression from object-oriented programming which C++ and more recently Java popularized. Another strand of influence came from a revival of interest in robotics [Brooks, 1991a; 1991b].

The quintessence of an agent is an intelligent, willing slave. Speculation in the area of artificial slaves is far more ancient than twentieth century science fiction. One documented example is found in Aristotle's *Politics* written in the fourth century BC. Aristotle classifies the slave as "an animate article of property". He suggests that slaves or subordinates might not be necessary if "each instrument could do its own work at command or by anticipation like the statues of Daedalus and the tripods of Hephaestus". Reference to the legendary robots devised by these mythological technocrats, the former an artificer who made wings for Icarus and the latter a blacksmith god, testify that the concept of robot, if not the name, was ancient even in Aristotle's time. Aristotle concluded that even if such machines existed, human slaves would still be necessary to render the little personal services without which life would be intolerable.

The name *robot* comes from the Czech words for serf and forced labor. Its usage originates from Karel Capek's 1920s play *Rossum's Universal Robots* in which Rossum, an Englishman, mass-produced automata. The play was based on a short story by Capek's brother. The robots in the play were not mechanical but grown chemically. Capek dismissed "metal contraptions replacing human beings" as "a grave offence against life". One of the earliest film robots was the replica Maria in Fritz Lang's 1927 classic *Metropolis*. The academic turned science fiction writer Isaac Asimov (1920–1992) introduced the term *robotics* when he needed a word to describe the study of robots in *Runaround* [1942]. Asimov was one of the first authors to depart from the Frankenstein plot of mad scientist creating a monster and to consider the social implications of robots.

An example of an automaton from the dark ages is a vending machine for holy water proposed by Hero of Alexandria around 11 AD. A modern reincarnation is Hoare's choc machine [Hoare, 1985] developed to motivate the computational model CSP (Communicating Sequential Processes). The word *automaton*, often used to describe computers or other complex machines, comes from the same Greek root as automobile meaning self-mover. Modern science owes much to the Greek tradition. Analysis of the forms of argument began with Empedocles and the importance of observation stems from Hippocrates. The missing ingredients of Greek science compared with the science of today were supplied by the Age of Reason. These were

the need for deliberately contrived observation - experiments; the need for inductive argument to supplement deduction; and the use of mathematics to model observed phenomena. The most important legacy of seventeenth century science is technology, the application of science. Technology has expanded human capability, improved control over the material world, and reduced the need for human labor. Willing slaves are, perhaps, the ultimate goal of technology.

Industrial robots appeared in the late 1950s when two Americans, Devol and Engelberger, formed the company Unimation. Take-up was slow and Unimation did not make a profit for the first fourteen years. The situation changed in the mid-1980s when the automobile industry, dissatisfied with trade union disruption of production, turned to robot assembly. However, the industrial robot industry overextended as governments curtailed trade union power and the market saturated. Many firms, including Unimation, collapsed or were bought out by end product manufacturers. Today, the big producer is Japan with 400 000 installed robots compared to the US with over 70 000 and the UK with less than 10 000.

With pre-Copernican mentality, people will only freely admit that humans possess intelligence. (This, possibly, should be qualified to mean most humans on most occasions.) Humans can see, hear, talk, learn, make decisions, and solve problems. It seems reasonable that anyone attempting to reproduce a similar artificial capability would first attempt emulating the human brain. The idea that Artificial Intelligence (AI) should try to emulate the human nervous system (brain cells are nerve cells) was almost taken for granted by the twentieth century pioneers of AI. Up until the late 1960s talk of *electronic brains* was common place.

From Rossum's Universal Robots in Carel Kapek's vision to HAL in the film *2001*, intelligent machines provide some of the most potent images of the late twentieth century. The 1980s were, indeed, a good time for AI research. In the 1970s AI had become something of a backwater in governmental funding, but all that changed dramatically because of the Japanese Fifth Generation Initiative. At the beginning of the 1980s, MITI, the Japanese equivalent of the Department for Trade and Industry, announced that Japan was to concentrate on knowledge based systems as the cutting edge of industrial development. This sent tremors of commercial fear through the corridors of power of every country that had a computing industry. These governments had seen national industries such as shipbuilding, automobile manufacturing, and consumer electronics crumble under intensive Japanese competition. In what retrospectively seems to be a halfhearted attempt to target research funds to industrially relevant information technology, a few national and multinational research programs were initiated. A major beneficiary of this funding was AI. On short timescales, commercial products were supposed to spring forth fully armed from basic research.

Great advances in computer hardware were made in this decade with computing power increasing a thousandfold. A computer defeated the world backgammon champion and a computer came in joint first in an international chess tournament, beating a grandmaster along the way. This, however, did not augur the age of the intelligent machine. Genuine progress in AI has been painfully slow and industrial take-up has been mainly limited to a few well-publicized expert systems.

In the mid-1980s, it was envisaged that expert systems that contain thousands of rules would be widely available by the end of the decade. This has not happened; industrial expert systems are relatively small and narrowly focused on specific domains of knowledge, such as medical diagnosis. As researchers tried to build more extensive expert systems major problems were encountered.

There are two reasons why game playing is the only area in which AI has, as yet, achieved its goal. Though complex, chess is a highly regular, codifiable problem compared with, say, diagnosis. Further, the algorithms used by chess playing programs are not usually based on expert systems. Rather than soliciting knowledge from chess experts, successful game playing programs rely mainly on guided brute force search of all possible moves using highly powerful conventional multiprocessor machines. In reality, AI has made as much progress as other branches of software engineering. To a large extent, its dramatic changes of fortune, boom and bust, are due to fanatical proponents who promise too much. The timescale predictions of the Japanese now look very fanciful indeed. AI has been oversold more than once.

A common reaction to the early efforts in AI was that successful replication of human skills would diminish human bearers of such skills. A significant outcome of AI research is how difficult the simplest skills we take for granted are to imitate. AI is a long-term problem, a marathon, and not a sprint competition with the Japanese. Expert systems are only an early staging post on the way to developing intelligent machines.

AI pioneered many ideas that have made their way back into mainstream computer science. These include timesharing, interactive interpreters, the linked list data type, automatic storage management, some concepts of object-oriented programming, integrated program development environments, and graphical user interfaces. Whatever else it achieved, the Japanese Initiative provoked a chain of increased governmental funding for Information Technology reaction around the world from which many, including the authors, benefited.

According to Jennings et al. [1998], the fashion for agents “did not emerge from a vacuum” (who would have imagined it would?) Computer scientists of different specializations artificial intelligence, concurrent object-oriented programming languages, distributed systems, and human-computer interaction converged on similar concepts of agent. Jennings et al. [1998] state, “Object-oriented programmers fail to see anything novel or new in the idea of agents,” yet they find significant differences between agents and objects. This is because their comparison only considers (essentially) sequential object-oriented programming languages such as Java. Had they considered concurrent object-oriented programming languages they would have found fewer differences.

Three languages have been promoted for agent development: Java, Telescript, and Agent-TCL. None of these are *concurrent* object-oriented languages. Java, from SUN Microsystems, is advocated for agent development because it is platform independent and integrates well with the World Wide Web. Java does, however, follow the tradition of interpreted, AI languages but is it not sympathetic to symbolic programming. Telescript, from General Magic, was the first commercial platform

designed for the development of mobile agents. The emphasis is on mobility rather than AI applications. Agent-TCL [Gray et al., 1996] is an extension of TCL (Tool Command Language) which allows mobile code. While string based, TCL does not have a tradition of AI applications. Programs are not inductively defined, as is the case with Lisp or Prolog.

This monograph describes a concurrent, object-oriented, agent programming language that is derived from the AI tradition. A working knowledge of Prolog is necessary to fully appreciate the arguments. The monograph is divided into two parts. The first part, Chaps. 1–5, describes the evolution of the paradigm of Guarded Definite Clauses (GDC). If the paradigm is serious, and more than a fashion, then it is necessary to describe its applications. This is done in the second part of the monograph, Chaps. 6–10. To set the paradigm in context, Chap. 1 provides an irreverent survey of the issues of AI. Chap. 2 completes the background to the paradigm with a retrospective rationale for the Japanese Fifth Generation Initiative. Chap. 3 describes how the paradigm evolved from Prolog with the environment change of multiprocessor machines. Included in this chapter is a chronology of the significant developments of GDC. Chap. 4 explores the manifestations of the vital ingredient of the paradigm - event driven synchronization. Chap. 5 compares and contrasts the language evolved with actor languages. The main difference is that GDC is an actor language with the addition of inductively defined messages.

The second part of the book begins with Chap. 6, which illustrates the advantages of GDC in parallel and distributed search. Chap. 7 describes the specialization to distributed constraint solving. Chap. 8 generalizes the chapters on search to meta-interpretation. An affinity for meta-interpretation has long been a distinguishing feature of AI languages. Chap. 9 describes how the overhead of meta-interpretation can be assuaged with partial evaluation. Chap. 10 concludes with the application of GDC to robotics and multi-agent systems.

While GDC as such is not implemented, it differs only marginally from KL1C, a language developed by the Japanese Fifth Generation Computer Systems Initiative. The Institute for New Generation Computer Technology (ICOT) promoted the Fifth Generation Computer Systems project under the commitment of the Japanese Ministry of International Trade and Industry (MITI). Since April 1993, ICOT has been promoting the follow-on project, ICOT Free Software (IFS), to disseminate the research:

According to the aims of the Project, ICOT has made this software, the copyright of which does not belong to the government but to ICOT itself, available to the public in order to contribute to the world, and, moreover, has removed all restrictions on its usage that may have impeded further research and development in order that large numbers of researchers can use it freely to begin a new era of computer science.

AITEC, the Japanese Research Institute for Advanced Information Technology, took over the duties of ICOT in 1995. The sources of KL1 and a number of applications can be obtained via the AITEC home page: <http://www.icot.or.jp/AITEC>. KL1C runs

under Linux and all the GDC programs in this monograph will run with little or no modification.

Despite their best efforts, the reader will find that the authors' cynicism shows through since they, like Bernard Shaw, believe that all progress in scientific endeavor depends on unreasonable behavior. In Shaw's view the common perception of science as a rational activity, in which one confronts evidence of fact with an open mind, is a post-rationalization. Facts assume significance only within a pre-existing intellectual structure that may be based as much on intuition and prejudice as on reason. Humility and reticence are seldom much in evidence and the scientific heroes often turn out to be intellectual bullies with egos like carbuncles.

The authors are very grateful to Jean Marie Willers and Peter Landin for the onerous task of proof reading earlier drafts of this monograph. Thanks are also due to our editors at Springer-Verlag, Ingrid Beyer, Alfred Hofmann, and Andrew Ross. Each author would like to say that any serious omissions or misconceptions that remain are entirely the fault of the other author.

January 1999

Matthew M Huntbach
Graem A Ringwood

Contents

Chapter 1: The Art in Artificial Intelligence	1
1.1 Realism	1
1.2 Purism	3
1.3 Rococo	4
1.4 Classicism	7
1.5 Romanticism	10
1.6 Symbolism	14
1.7 Neo-Classicism	17
1.8 Impressionism	20
1.9 Post-Impressionism	22
1.10 Precisionism	25
1.11 New Realism	28
1.12 Baroque	29
1.13 Pre-Raphaelite Brotherhood	32
1.14 Renaissance	33
1.15 Hindsight	34
Chapter 2: Fifth Generation Architecture	37
2.1 Architecture and Design	38
2.2 Design as Evolution	40
2.3 Design as Co-evolution	43
2.4 Design as Theorem	44
2.5 Design as Premise	48
2.6 Design as Paradigm	50
2.7 Impressionist Design	54
2.8 Classical Design	59
2.9 Logic Machines	61
2.10 Hindsight	64
Chapter 3: Metamorphosis	69
3.1 Apparent Scope for Parallelism	70
3.2 Or-Parallelism	72
3.3 The Prolog Phenomenon	73
3.4 Concurrency and Operating Systems	77
3.5 Concurrency and Distributed Systems	78
3.6 Symbiosis Between Programming Language and System Engineering	80
3.7 Event Driven Synchronization	81
3.8 Earlier Manifestations of Guarded Commands	83
3.9 Condition Synchronization in AI	84
3.10 Guarded Definite Clauses	87
3.11 Simulation of Parallelism by Interleaving	90
3.12 Indeterminacy	92

3.13	The Premature Binding Problem Revisited	93
3.14	Decision Tree Compilation	96
3.15	A Brief History of Guarded Definite Clauses	97
Chapter 4: Event Driven Condition Synchronization		103
4.1	Streams for Free	104
4.2	A Picture is Worth a Thousand Words	106
4.3	Dataflow Computation	108
4.4	Dataflow Design	110
4.5	Dataflow Programming	112
4.6	Message Passing	115
4.7	Eager and Lazy Produces	118
4.8	The Client-Server Paradigm	122
4.9	Self-Balancing Merge	124
4.10	Synchronization	125
4.11	Readers and Writers	127
4.12	The Dining Philosophers	128
4.13	The Brock–Ackerman Anomaly	132
4.14	Conditional Semantics	133
4.15	Open Worlds and Abduction	135
4.16	Implementation Issues	137
Chapter 5: Actors and Agents		139
5.1	The Actor Model	140
5.2	Haggling Protocols	144
5.3	Consensus Protocols	146
5.4	Market Forces	148
5.5	Poker Faced	148
5.6	Virtual Neural Networks	149
5.7	Biological and Artificial Networks	150
5.8	Self-Replicating Neural Networks	152
5.9	Neuron Specialization	152
5.10	The Teacher Teaches and the Pupil Learns	157
5.11	Neural Simulation	160
5.12	Simulated Life	162
5.13	Life Yet in GDC	163
5.14	Cheek by Jowl	164
5.15	Distributed Implementation	165
5.16	Agent Micro-Architectures	167
5.17	Metalevel Agent Architectures	168
5.18	Actor Reconstruction of GDC	170
5.19	Inheritance Versus Delegation	172

Chapter 6: Concurrent Search	175
6.1 A Naive Prolog Solution to the 8-Puzzle	175
6.2 Speculative Parallelism	180
6.3 Non-speculative, Non-parallel Linear Search	182
6.4 A Practical Prolog Solution to the 8-Puzzle	183
6.5 A Generic Search Program	186
6.6 Layered Streams	188
6.7 Eliminating Redundant Search	190
6.8 A Direct GDC Solution Using Priorities	193
6.9 Search Anomalies	195
6.10 Branch-and-Bound Search	198
6.11 Game Tree Search	203
6.12 Minimax and Alpha-Beta Search	205
6.13 Parallel Game Tree Search	207
6.14 Parallel Search and Cooperative Distributed Solving	211
 Chapter 7: Distributed Constraint Solving	213
7.1 All-Pairs Shortest Path Problem	216
7.2 The Graph Coloring Problem	220
7.3 Minimal Spanning Trees	233
7.4 Conclusion	244
 Chapter 8: Meta-interpretation	247
8.1 Metalanguage as Language Definition and Metacircular Interpreters	248
8.2 Introspection	250
8.3 Amalgamating Language and Metalanguage in Logic Programming	251
8.4 Control Metalanguages	254
8.5 A Classification of Metalevel Systems	256
8.6 Some GDC Monolingual Interpreters	259
8.7 GDC Bilingual Interpreters	265
8.8 An Interpreter for Linda Extensions to GDC	269
8.9 Parallelization via Concurrent Meta-interpretation	274
8.10 Conclusion	277
 Chapter 9: Partial Evaluation	279
9.1 Partial Evaluation	280
9.2 Futamura Projections	283
9.3 Supercompilation	289
9.4 Partial Deduction	295
9.5 Partial Evaluation and Reactive Systems	297
9.6 An Algorithm for Partial Evaluation of GDC Programs	300
9.7 Actor Fusion	305
9.8 Actor Fusion Examples	308

9.9 Partial Evaluation of an Interpreter 310

Chapter 10: Agents and Robots 319

10.1 Reactive Agents: Robots and Softbots 319

10.2 A Simple Robot Program 321

10.3 Reaction and Intelligence 327

10.4 Objects, Actors and Agents 329

10.5 Objects in GDC 332

10.6 Agents in GDC 336

10.7 Top-Down and Bottom-Up Multi-agent Systems 339

10.8 GDC as a Coordination Language 341

10.9 Networks and Mobile Agents 345

10.10 Conclusion 348

References and Bibliography 353