

of the bibliography to researchers new to the field. Also, although the exercises are plentiful and instructive, some crucial results depend on exercises for which no solutions are given. Despite these shortcomings, this book constitutes a useful starting point for new researchers.

CAROLYN BROWN
University of Sussex

JUAN BICARREGUI *et al.*

Proof in VDM: A Practitioner's Guide. Springer-Verlag. 1994. ISBN 3 540 19813 X. 363 pp. softbound.

Proof is often cited as a technique that may be used with respect to a formal specification but not with an informal one. Using proof, we can not only discharge well-formedness and satisfiability obligations, but we can also prove validation conjectures and/or satisfaction of particular properties. However, to date, few examples of such proofs exist in the literature, and those examples that do exist provide little or no guidance to the novice on how they have been constructed. This book seeks to redress the balance by demonstrating how to go about constructing proofs in formal specifications. The style of presentation clearly illustrates how proofs evolve, from a general outline of the proof, to a detailed formal proof using axioms and inference rules. Although the text is based on specifications written in VDM, the emphasis is on the principles of formal proof, rather than on idiosyncratic properties of VDM.

The book is divided into three sections. The first section describes the logical system used in proofs and presents an axiomatization of the basic data types in VDM. In particular, this section introduces the idea of proof as a means of deriving theorems concerning data types from the axioms for these data types. Each data type is introduced by giving the appropriate axioms of formation, and, where appropriate, direct definitions. Examples of the use of such axioms, definitions and theorems are liberally scattered throughout the text, and in general are thoughtfully described. Some data types also include axioms which describe induction rules, together with corresponding examples. However, the treatment of induction is slightly weaker than the overall standard of the section—more overview of the general idea would be kinder on the novice reader.

The second section, entitled 'Proofs In Practice', presents some applications of the proof theory described in the first section. Three categories of application are considered: proof obligations arising from an arbitrary specification, proof obligations arising from reification and an illustrative case study. The first category describes, at an abstract level, how to prove well-formedness and satisfiability of specifications. The description is largely restricted to the general form that such proofs take, though a few concrete examples are given. The chapter on reification describes data reification and operation modeling, using retrieve

functions. An example is presented, which although simple, adequately illustrates the concepts described. The result is a concise, elegant description of data reification that is a pleasure to read.

The case study, however, is the real revelation of the section. It is customary, when reading tutorial texts on formal specification, to expect illustrative case studies to be so simple as to be totally unrealistic. Here, quite the opposite is true: the case study presented, though necessarily simplified to avoid excessive detail, is remarkably realistic. The case study itself concerns the specification of a tool that manages the allocation of aircraft to air-traffic controllers within an air-traffic control region. Following an analysis of the system requirements, the formal state model is presented and its corresponding theory given. Within this theory, the arising proof obligations are formulated and discharged, and similarly for the validation conditions arising from the system requirements. The top-level operations of the system are specified, then two refinement steps are given, the first of which includes all the relevant proofs. All in all, the case study perfectly fulfills its objective of drawing together the material of the preceding sections in a realistic setting.

The final section consists of a brief description of data types not already treated and a directory of theorems. The former discusses how to axiomatize such data types and where relevant discusses the resultant problems. In a sense, this chapter is not strictly necessary but merely ties up some loose ends. On the other hand, the directory of theorems is a hive of information, collecting together theorems and axioms, grouped according to data type. To anyone embarking upon a proof, such a directory is invaluable.

Throughout the text there are many exercises of varying degrees of difficulty. The solutions of the exercises are available by FTP, and mirror the thoroughness and clarity of thought manifested in the main text. However, the main text itself is rather less readable than the solutions—the small typeface and narrow margins combine to give the text a rather cramped appearance, though this is presumably to keep down the price of the book.

In the past, anyone wishing to perform proof on formal specification has had to ask the question: "where does one start?". With the publication of this book, this question has been answered totally. Not only does the book explain what proof is, when it should be used, why it should be used and how to use it, but it does so in a style that does not patronize the expert and simultaneously engenders confidence in the novice.

P. MUKHERJEE
University of Birmingham

LESLIE PACK KAEHLING

Learning in Embedded Systems. MIT Press. 1993. ISBN 0-262-11174-8. £26.95. 176 pp. hardbound.

Roboticians have, over the past few years, sought to escape the constraints imposed by procedural programming on the flexibility of their robots. This monograph, by a roboticist, describes how the introduction of learning into such systems may increase their adaptability and presents a considerable number of results obtained with algorithms which undertake trial-and-error learning.

The *embedded systems* of the title are not limited to robots—they may be any systems which interact continuously with a dynamic, external world, sensing it, acting in it and existing for considerable periods of time. Their goal is to adapt their behaviour to the environment by learning mappings between perceived states of the world and the action to take in that situation; the world gives them a reinforcement signal indicating how desirable the current state is for the system in the context of carrying out a particular task.

Kaelbling compares the performance of several algorithms from the reinforcement learning literature on a series of learning tasks in which reinforcement is generated at each time step. Her extensions to these algorithms and new algorithms address the exploration/exploitation tradeoff, the reduction of learning a single many-many mapping to learning many many-one mappings with the concomitant decrease in time complexity, learning Boolean functions from reinforcement and learning mappings with state. She also addresses the problem of delayed reinforcement, when an indication of how well a system is doing is not received until the end of a long sequence of actions. The concluding chapter goes considerably beyond her doctoral thesis (of which this book is a revised version) in discussing the merits and limitations of the approach and in pointing the way to future work.

One of the excellent things about this book is the large number of tables and graphs showing the results of performance comparisons. These, and the discussion of them, will be invaluable to anyone wanting to try out the algorithms in their own systems. It would have been nice to see more results from experiments on (real) robots in addition to those on simulations. Informal results for a robot engaged in a light-finding task are presented; however, the very real problem of applying reinforcement learning in robotics is that “it takes a long time to conduct the experiments”. Kaelbling does a good ground-clearing job in indicating what sort of algorithms are in with a fighting chance.

Good references are included, but the index is, sadly, rather sketchy. What out for binding errors: a few pages were duplicated and others missing in my review copy.

This is not an introductory book, requiring a reasonable amount of mathematical experience in the reader. I would recommend it for graduate students and anyone working in other areas of machine learning or robotics who want a crisp overview, with performance data, on action-map-learning algorithms.

GILLIAN HAYES
Edinburgh University

PAUL JOHN WERBOS

The Roots of Backpropagation. John Wiley. 1994. ISBN 0-471-598976-6 £45.40. 319 pp. hardbound.

The title *The Roots of Backpropagation* announces that the book is laying claim to chart the development of the now well-known backpropagation algorithm for Neural Network learning. Before discussing the book, I would like to place the development of this algorithm in perspective.

The backpropagation algorithm is more usually credited to recent work by Rumelhart *et al.* (1986) though there has been some dispute as to who was the first to devise this answer to the challenge laid down by Minsky and Papert (1988) in their book *Perceptrons*. Minsky and Papert argued that single neuron training is possible but lacks computational power, while networks of neurons pose an insurmountable training problem. The ‘backpropagation’ algorithm is actually a method of efficiently computing the derivative of the current error of the network with respect to its weight parameters provided the activation function is smooth. It therefore allows the parameters to be adapted by gradient descent in a direction of maximum error reduction in what should more correctly be termed the generalised delta update rule. Whatever we choose to call the Neural Network learning algorithm, it has become in one or other of its variants the de facto standard method of adapting feedforward Neural Network parameters. This is true, despite the fact that it is frequently extremely expensive computationally and can and does frequently get caught in local minima.

The first part (and 80%) of Paul Werbos’ book is a publication of his 1974 Harvard University doctoral dissertation, *Beyond Regression*. As the title suggests, the thesis develops statistical modelling techniques which move significantly beyond simple linear regression. The development is well motivated with a mixture of examples and theoretical justification, leading through the ARMA (auto-regressive moving average or Box-Jenkins) model to non-linear models.

Chapter 2 contains a careful exposition of the different models, justifying them in terms of maximizing likelihoods in particular models of noise. As part of this development a new ‘robust’ model is introduced, which assumes measurement noise only. This means that the data can be used to fit the whole curve, rather than allowing errors to affect its shape at each stage. This perhaps corresponds to the difference between fitting a polynomial (with no adjustment allowed) and a spline (adjustment of shape allowed at the nodes). In fact Werbos can mix both models to give a balance of the two approaches.

In order to compute the parameters he frequently