to facilitate, comparative judgements. Nevertheless they will become rapidly out-of-date, and the decision to state absolute figures rather than provide some category scale as an aid to relative judgement was surely a mistake. Indeed some of the descriptive details provided are already obsolete. For instance KEE is stated to be available only on Xerox machines. Far from it: not only has it been on Symbolics and Texas machines for a while, but it will shortly, we hear, not be supported on Xerox at all.

This book left me with mixed feelings. Anything that guides managers toward beneficial usage of expert system software helps the industry advance further into what is potentially its most valuable market. When addressing this market we must be sure that our voice not only has the quality of clarity, but is pitched at a level to inspire confidence, and has behind it the authority necessary for delivering state-of-the-art distillations. This book succeeds on the first count, but fails thereafter. Given its size, structure and aims, it might have done much better as a series of articles in one of the popular management journals. In this way the author's stated aim of increasing awareness amongst managers might have been more effectively met.

ALL THE RIGHT MOVES – A VLSI ARCHITECTURE FOR CHESS

Carl Ebeling, 1987. The MIT Press, Massachusetts. Pp. 145

Reviewed by Don Beal, Queen Mary College, London University

This book is a publication of Carl Ebeling's Doctoral Thesis which merited selection as an ACM distinguished dissertation. It is remarkably well written and easy to read. Although the title includes "VLSI", hardware design details only occur in parts of the text, and the book is accessible to a wide audience.

The first two chapters survey some of the performance achieved and techniques used by chess-playing programs up to 1985. A crucial operation is that of generating the set of legal moves for a given board position. Chapter 3 discusses the details of doing this by a new approach used in the chess machine "Hitech" built by Ebeling and others at Carnegie Mellon University. The new approach uses parallelism in a different way from hitherto, replicating circuitry on a per-square basis, but in addition with each square replicating some knowledge of the global board state on a relevant-to-this-square basis. Much of the circuitry is devoted to achieving a desirable ordering of the moves.

Chapter 4 discusses position evaluation and describes the recently introduced concept of an "oracle" and a hierarchy of first-order, second-order and higher-order evaluations. The oracle is a collection of chess knowledge separated from the search mechanisms, and which serves to choose the coefficients for the evaluation function during a particular search. The other side of the coin is that the evaluation terms are deliberately simplified down to what can, with affordable hardware, be computed in a single machine cycle. This turns out to be a good trade despite the possibility that the search will encounter deep positions in which the coefficients chosen as appropriate at the root will be inappropriate. This occurs, but is sufficiently infrequent, or sufficiently minor, to be acceptable. The oracle knowledge can be as complex as desired, because it is only executed once per move computation, whereas the position evaluation is executed millions of times. This splitting of knowledge into two extreme forms has been successful enough to make Hitech currently the best chess-playing machine, having achieved a chess rating approaching 2400 against human opposition in regular human tournaments.

Chapter 5 contains hardware details of the rest of the Hitech machine. Chapter 6 contains a survey of the performance results and testing of Hitech, showing how the various hardware components, move ordering, and the hash table contribute to the overall performance. It is estimated that minor redesign and technology improvements already available could produce a machine five times faster. However, there is a section devoted to the trade-off between knowledge and speed, almost always resolved in favour of speed in the history of computer chess up to 1985, which shows that Hitech gained 150 rating points from pawn structure knowledge, and 200 points from other chess knowledge since it was first commissioned. To obtain this improvement by increased speed instead would require something of the order of a 128-fold speedup, according to experiments and estimates by the Hitech team. Ebeling suggests that these figures indicate that in Hitech at least, the path to better play lies with increased knowledge rather than increased speed.

Although designing special purpose hardware may seem irrelevant to much work in expert systems and knowledge engineering, such activities do provide an interesting case study for what happens when knowledge is "compiled down" into application specific code. Essentially, the chess hardware is the end result of massaging rules about what constitute legal moves, plus rules about what constitute desirable positions, into dedicated parallel machinery. I noticed that parallelism in this case required a certain amount of redundancy of the global knowledge, and that although the position evaluator was based on complex and global knowledge of the board, it was possible in many cases to find alternative knowledge rules that, with the aid of extra state variables, enabled the board evaluation to be incrementally computed as individual squares of the board changed.

The final chapter identified "distributed state" as being the key feature of the hardware architecture. This allows fast computation from incremental changes, and expansion without much slow-down by adding additional units. The book concluded by suggesting that this kind of architecture might also be applicable to low-level perception in vision, some aspects of speech recognition, and analysis of molecular structures.