Review of

Qualitative Reasoning

by Benjamin Kuipers

MIT Press, 1994.

Ernest Davis, New York University

Qualitative simulation is the process of deriving a partial characterization of the behavior of a dynamic system given only weak information about it. For example, if one observes an irregularly shaped can full of water with a hole in the bottom, one can predict that the water will drain out with a steadily decreasing rate of flow, but without more precise information it is impossible to state what the rate of flow will be, or how long it will take for the can to empty.

The AI study of qualitative simulation began with Johan de Kleer's NEWTON program of 1977 [2]. The field rapidly gathered momentum, and in 1984, a special issue of *Artificial Intelligence* [1] was devoted to the subject. Particularly notable in that issue were descriptions of three systems that have become standards in the field: de Kleer and Brown's ENVISION [3], Forbus' QP [4] and Kuipers' QSIM [5]. Since that time, these programs, particularly QP and QSIM, have been widely studied, applied, and extended, both by their original creators and by researchers worldwide.

Qualitative Reasoning is a description of QSIM and its many extensions: their theoretical foundations, implementations, and domain applications. The book is written in textbook style, with careful explanations, a wealth of examples worked through in detail, and exercises ranging from simple homework problems to doctoral-level research projects. The style is consistently clear, precise, and engaging. A student with a minimal background in differential equations will find it a pleasant and easy introduction to the field of qualitative simulation. The student will gain both a solid knowledge of the techniques involved and a gut feeling for the kinds of problems that are reasonably addressed by these techniques and the kinds of information that they provide.

Chapters 1-7 present the basic QSIM program, as described in the original AIJ article. This is well-worn ground, but the account here is certainly the most easily accessible in the literature: patient, systematic, clear, detailed, and full of real-world examples, both from physics and from physiology. Myself, I particularly enjoy the physiological examples, for a number of reasons: they teach me something new; they provide a welcome break from the weary stream of hydraulic, electronic and mechanical systems that fill the AI physical reasoning literature; and they constitute an excellent answer to the eternal question of students and skeptics, "Why not just use the exact equations?" Kuipers originally came to the study of qualitative simulation from work on reasoning for medical applications, so these physiological problems were genuinely motivating examples for qualitative simulation, rather than constructed *post hoc* to fit the theory.

Chapters 8 through 12 present a variety of extensions to QSIM: dealing with region transitions, integrating partial quantitative information, using constraints on higher-order derivatives, using generalized energy functions and phase space analysis, and reasoning across multiple time scales. These techniques can be used to eliminate some of the ambiguities in QSIM's output and to provide richer types of information about the behavior of the dynamical system. The formal treatment here is necessarily faster than in chapters 1-7, and the mathematics somewhat more demanding, but the exposition is still very clear.

Chapters 13 and 14 discuss other work in qualitative physical reasoning, particularly the component model [3], QP theory [4], and recent work on model construction [7,9], and show how these can be connected to QSIM. Quite new to me here, and very interesting, is a survey (section 14.7) of a number of projects developing large knowledge bases for reasoning in such domains as thermodynamics, botany, and chemical engineering.

Finally, the book ends with three appendices, containing a glossary of technical terms, an index of QSIM functions, and a brief discussion of creating and debugging QSIM models. This last is one of the most interesting sections of the book; I wish that there were a whole chapter on the subject.

There are some significant flaws and gaps in the book. First, the many figures that display the output of QSIM as graphs of behavior traces are consistently very difficult to read, and occasionally (e.g. figure 10.10) impossible to read. They should have been redrawn and reproduced at twice the size. This production flaw mars the book and obstructs the reader.

Second, two important extensions to QSIM are omitted: Dan Weld's work on comparative dynamic analysis [8] (not even cited in the bibliography) and the recent paper by Kuipers and Shults [6] on using a query posed in a modal temporal logic to prune the construction of the behavior tree. Presumably, the first was a victim of the "Not invented here" syndrome, coupled with a simple oversight, while the second appeared too late for inclusion.

The third problem is the most serious, as it is a problem, not with the book, but with the state of the art. The AI study of qualitative simulation has not been analyzed in the standard modes of computer science. Next to nothing, for instance, is known about the computational complexity of the tasks involved. Practically the only result of this kind I know that it is NP-complete to determine whether a set of confluences on signs has a solution. I have the impression that the hybrid-systems folks have some more results on computational complexity of solving differential equations. I have not taken the time to look these up and to figure out whether they are relevant to QDE's, but I wish that Kuipers had. Kuipers does not mention the issue, even as a research problem.

Ignoring formal computational complexity can perhaps be excused: these results are hard to derive and it can be debated how much they mean in practice. There is no such excuse, however, for ignoring empirical study of the behavior of qualitative simulators. How does the running time of QSIM grow with the size of the problem? or with the size of the output? What features of a QSIM problem have the most impact on running time? How do the various extensions affect the running time? I have seen in the literature some comparisons of the running time of different systems, but I don't remember ever having seen a study of the dependence of running time on problem characteristics. Certainly, there is no discussion of the issue in this book. In the last few years, the planning community has made great strides in the formal and empirical evaluation of planning programs. It is high time that the qualitative simulation community did likewise.

Overall, however, this book is a major contribution to the field. For the student, it is the best introduction to the field that has been written. For the researcher, it summarizes the many accomplishments of more than ten years of work of a very successful and productive research project. If you are putting together a library of books on AI physical reasoning, then the first book to buy is still Weld and de Kleer's collection of readings [10]; but this is the second book.

References

[1] D.G. Bobrow (ed.) *Artificial Intelligence*, special issue on Qualitative Reasoning about Physical Systems, vol. 24, 1984. Reprinted as *Qualitative Reasoning about Physical Systems,* M.I.T. Press, Cambridge, MA, 1985.

[2] J. de Kleer, "Multiple representations of knowledge in a mechanics problem solver," *Proc. 5th Int. Joint Conf. on Artificial Intelligence,* 1977, pp. 299-304.

[3] J. de Kleer and J.S. Brown, "A qualitative physics based on confluences," in D.G. Bobrow (ed.) *Qualitative Reasoning about Physical Systems,* M.I.T. Press, Cambridge, MA, 1985, pp. 7-83.

[4] K. Forbus, "Qualitative process theory" in in D.G. Bobrow (ed.) *Qualitative Reasoning about Physical Systems,* M.I.T. Press, Cambridge, MA, 1985, pp. 85-168.

[5] B. Kuipers, "Commonsense reasoning about causality: Deriving behavior from structure," in D.G. Bobrow (ed.) *Qualitative Reasoning about Physical Systems,* M.I.T. Press, Cambridge, MA, 1985, pp. 169-204.

[6] B. Kuipers and B. Shults, "Reasoning in logic about continuous systems," *Proc. 4th Int. Conf. on Principles of Knowledge Representation and Reasoning,* Morgan Kaufmann, 1994.

[7] P.P. Nayak, "Causal approximations," *Proc. 10th National Conf. on Artificial Intelligence,* pp. 703-709.

[8] D. Weld, *Theories of Comparative Analysis* MIT Press, 1990.

[9] D. Weld, "Approximation reformulations," *Proc. 8th National Conf. on Artificial Intelligence,* pp. 407-412.

[10] D. Weld and J. de Kleer, *Readings in Qualitative Reasoning about Physical Systems,* Morgan Kaufmann, 1990.