

NARROW VIEWS, OLD TASKS, AND NEW BEGINNINGS

BRIAN C. WILLIAMS

MARK SHIRLEY

OLIVIER RAIMAN

BRIAN FALKENHAINER

JOHAN DE KLEER

DANIEL G. BOBROW

*Scientific and Engineering Reasoning Area
Xerox Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto CA 94304, U.S.A.*

One can respond to "Prolegomena to any Future Qualitative Physics" at several levels. As practitioners in the field, it is easy to become incensed and respond to every minute misrepresentation and error. This would unfortunately miss the forest for the trees. Unfortunate, because there is a much simpler story to be told here. Sacks and Doyle make two high-level claims as the basic theme of their paper. First, they state that the field of qualitative reasoning is roughly equivalent to SPQR. In other words, the principal goal of the qualitative reasoning enterprise is to (1) produce abstract state transition diagrams of system behavior (i.e., simulations) which are (2) solely for the purpose of predicting behavior. Second, in reasoning about such systems' behavior, "experts" are predominantly concerned with their asymptotic, dynamical properties.

As stated, we view both points as incorrect, misleading, and quite surprising coming from Sacks, who, as a frequent attendee of the qualitative reasoning workshops, should have a better understanding of the field. However, there is another, less controversial interpretation of the paper which we obtained from personal conversations with Sacks at the most recent workshop. First, their argument was intended to address work on the SPQR subset of the field rather than suggest that all of qualitative reasoning is SPQR. Second, their reference to the concerns of an expert was referring to an expert academic dynamicist, and was a suggestion about how to compare the relative merits of SPQR and dynamical analysis.¹

Given this interpretation shift, Sacks and Doyle's paper raises an interesting, but much narrower, technical issue related to acausal simulation, a small portion of the work going on in the field of qualitative reasoning. One argument in favor of qualitative simulation (SPQR) is that qualitative differential equations enable reasoning about entire classes of functions and is thus a more efficient way to perform certain tasks, such as the preliminary stages of design. Alternatively, Sacks and Doyle argue that it is better to pick a specific differential equation, analyze its solutions using conventional techniques, and if these are undesirable (for the task, say diagnosis or design) to pick another one. We believe that neither has demonstrated clear superiority for the kinds of tasks qualitative simulation is intended to address and we think it is perfectly reasonable that researchers in qualitative simulation explore and evaluate both approaches.

Rather than focus on that debate in our response, we instead address the broader issues that have arisen while discussing this paper. We try to clarify what we see as the goals of the field and describe some of the major research efforts.

¹We were told that they could not make these clarifying changes because some people had already invested significant effort responding to the original paper.

1. QUALITATIVE REASONING'S OTHER GOALS

Although the relative merits of SPQR and dynamical analysis is a very important question, SPQR is in no way central to the goals of qualitative reasoning. To us the heart of the qualitative reasoning enterprise is and always has been to develop computational theories of the core skills of engineers, scientists, and just plain folk's who need to answer questions like:

- What does it do?
- Why does it work?
- What is it like?
- What is most relevant?
- How do we create it?
- Does it work correctly?
- Where is it broken?
- How do we fix it?

as they test, predict, create, optimize, diagnose and debug physical mechanisms (Williams 1984; de Kleer 1990; Williams and de Kleer 1991). This perspective both gives us the means to evaluate new theories (do they address significant components of or yield significant improvements on these tasks?) and suggests places to look for new ideas. Engineering practice and the study of performance tasks, such as diagnosis and design have led to the development of core skills such as qualitative symbolic algebra, causal explanation, terminological reasoning, candidate generation, minimum entropy and order of magnitude reasoning. They have not replaced traditional mathematics techniques but have been integrated with them (see Williams and de Kleer 1991 for numerous examples). Qualitative reasoning has properly used the tasks of diagnosis, design, theory formation and tutorial explanation to elucidate and focus on the most significant skills and their interconnections.

2. RE-VIEWING THEIR OBSERVATIONS

Sacks and Doyle assert in their conclusion what they see as the five principal limitations of SPQR. Here, we briefly examine these points, not in the narrow context of SPQR, but qualitative reasoning research in the large. We use them as a way to describe what we see as some of the major research endeavors of qualitative reasoning and to give references to some of the relevant work.

1. "The accepted approach, which we have called SPQR, has analyzed only the simplest systems, while routine expert methods succeed on far more complicated systems."

Among others, qualitative reasoning techniques have successfully diagnosed large transistor circuits (Dague *et al.* 1987) explained thermal power plants (Falkenhainer and Forbus 1988; Skorstad and Forbus 1989) and interpreted geological structures (Simmons 1988). Why do Sacks and Doyle say that these are only the "simplest systems?" What routine expert systems diagnose, explain, or interpret far more complicated systems?

We can only make sense of this if we recall that Sacks and Doyle are only talking about the dynamical aspects of systems, and that their analysis is specifically concerned with the asymptotic properties of these systems. Simply stated, Sacks and Doyle are tacitly examining everything through the eyeglasses of an academic dynamicist. Indeed we readily admit that the dynamical properties of many qualitative reasoning examples seem relatively

simple. But are they the right glasses in which to examine the work of qualitative reasoning? Academic dynamicists are probably not the ones who should define what conceptual design, diagnosis and explanation problems are simple or complex.

2. "Virtually all the systems reduce to three equations."

Consider one of the larger bipolar examples diagnosed by Dedale using order of magnitude reasoning (Raiman 1986). It has 17 bipolar transistors, 27 resistors and 3 diodes, and is modeled with roughly 180 equations. Of course, Sacks and Doyle didn't include this example since the Dedale AAAI paper didn't have enough room to include these 180 equations. Relatively complicated circuits have been widely explored in qualitative reasoning (Roynance 1980; Brown *et al.* 1982; de Kleer 1984; Williams 1984; Dague *et al.* 1987). But no circuit of two or more components can be described by only three equations.

So what do Sacks and Doyle mean by "three equations"? It would appear that they are performing arbitrary amounts of algebra on the equations, and are interested in the behavior of only one dependent variable. Thus, if we restrict each transistor to a single operating mode we can reduce the preceding Dedale example to a single equation. But it would be very difficult to diagnose, redesign or explain the circuit's behavior based on that one equation. Again, Sacks and Doyle are only beginning their academic dynamicist's eyeglasses to bear.

The above example is not atypical. The fact is that engineers frequently design devices so that their dynamics are quite simple. Thus, rather than concluding that most engineering examples are toys, the other interpretation is that many problems do not have sufficiently complex dynamics to warrant heavy use of the academic dynamicists' tools. This is in fact why many qualitative reasoning researchers, although trained in the tools of dynamics, have chosen to place their energy elsewhere.

3. "SPQR equations are far too general for practical use. Experts instead hypothesize and revise specific equations, until they obtain equations of adequate accuracy."

No practical use of any kind? In Section 4.2 Sacks and Doyle include order of magnitude equations (Raiman 1986) as part of SPQR, yet Dedale was quite successful at bipolar circuit diagnosis. Likewise, preliminary research in modeling Xerox copier document handlers (consisting of roughly 160 components and 50 component types) suggests that field service diagnosis of copiers can be done using simple behavioral descriptions like "moving," "not moving," "paper on-time," or "late." Significantly, more detailed models, e.g., the differential equations of paper motion, seem both unnecessary and inappropriate for this application.

One might be tempted to conclude that research for engineering related domains should *only use* traditional engineering representations. However, the question for research should always be "what is the best tool to solve the problem?" not "what is the traditional tool"? This is one of the most basic tenets of AI for what constitutes a good representation (Winston 1984).² The experts we talk to, who diagnose our artifacts and design our machines, rarely hypothesize equations let alone revise them. The important task is not to be able to build superaccurate models, but to fix or design the darn things.

While we cannot go into detail here about why each of the representations we proposed are appropriate, in reports like de Kleer 1979; Williams 1984, 1991 we have put considerable effort into justifying them through analyses of tasks and protocols. Likewise, for the

²As paraphrased from Winston (1984): Good representations facilitate problem solving. They make the important things explicit and expose natural constraints. They are complete, concise, transparent, facilitate computation, and suppress detail.

qualitative reasoning community to fully appreciate the importance of the computational dynamics work we need to see descriptions of what tasks dynamicists perform and rationalizations for why the proposed tools provide exactly the right information. Unfortunately in Sacks (1990*a,b*) we find these missing. Sacks (1991) begins to shed some light on this question, but a clearer explication is still needed.

4. "Experts focus on asymptotic behavior, while SPQR focuses on transient behavior."

But what about transient analysis programs like SPICE and timing simulators of all forms which have been proven useful by widespread use for many years? For example, the primary concern when designing high performance n-MOS circuits is switching speed and peak power dissipation. Transient behavior is exactly what a designer focuses on in order to understand and modify these types of circuits (Glasser and Dobberpuhl 1985) and is precisely what motivated the development of explanation systems like temporal qualitative analysis (Williams 1984).

This is not to say that asymptotic behavior is unimportant. For example, it is indeed used when evaluating the stability of a bipolar amplifier, when designing the amplifier's compensation. But the analysis performed is typically quite simple (Gray and Meyer 1984), and is performed during the late stages of design, not its initial conception. More generally, while naturally occurring systems may have complicated dynamics, engineered devices are designed so that their dynamics is as simple as possible (e.g., nonlinear systems are often made to behave near-linear). Now Sacks might tell you, as he did us, that recent advances in dynamics will change all this, as evidenced by Gleick's popular book, *Chaos* (1987), and that in the future all engineers, scientists (and even car mechanics?) will embrace the mathematical tools of dynamics. If so, this is quite exciting. But it hasn't happened yet, and expert dynamicists we have spoken to greet this claim with considerable skepticism.

5. "Experts derive the behavior of dynamical systems with deep mathematics and extensive numerical analysis, whereas SPQR uses little of either."

Qualitative reasoning research began with the intuition that an engineer's ability to explain *why* a device behaves as it does is a key to being able to perform a diverse set of tasks, like diagnosis and design.³ Traditional computational tools like SPICE (Vladimiresco *et al.* 1981) and MACSYMA (Moses 1971) are invaluable to an engineer, but only provide a partial answer to automating the explanation skill, since they tell us *what* a device does, but not *why*. Thus the essential goal of work like ENVISION (de Kleer 1984; de Kleer and Brown 1984), temporal qualitative analysis (Williams 1984), qualitative process theory (Forbus 1984), and temporal constraint propagation (Roylance 1980) was causal explanation, not, as Sacks and Doyle claim, SPQR.

Researchers in qualitative reasoning are interested in any mathematical tool that has proven to be right for the job. However, we very much doubt that the academic dynamist's "deep" mathematics is nearly as all encompassing as Sacks and Doyle would like us to think. For example, when we asked around Xerox we found that our diagnostic technicians use little other than algebra and arithmetic, and our designers mainly use algebra, finite element analysis, and some transient analysis. These are the types of algebraic reasoning being captured in research on hybrid qualitative/quantitative (Williams 1991) symbolic algebra and order of magnitude reasoning (McAllester 1981; Raiman 1986; Mavrouniotis and Stephanopolous 1988), and the types of numerical tools being incorporated

³At the time we began our research, causal explanation was considered central to tasks like medical diagnosis, thus lending support to the belief that it is a key skill, and that AI had something to contribute.

into systems like SIMGEN (Forbus and Falkenhainer 1990) and Q3 (Berleant and Kuipers 1990). Even the experimental physicists we talk to mainly use algebra and least-square fit. In fact, very few of the experts who perform the tasks we are seeking to model have the slightest understanding of theoretical dynamics.

Given the questions Sacks and Doyle's points raise, future discussion critically depends upon the research being conducted in a style that facilitates direct and objective comparison. Returning to the narrower technical issue raised by their paper, one way to proceed might be to examine SPQR and dynamical analysis in the context of a common example involving a concrete design, diagnosis or other task. In particular, the goals of the task should be stated clearly and specifically enough that comparison is meaningful.

3. PROLEGOMENAS TO PROLEGOMENAS

So what is the "real prolegomena" for any future qualitative reasoning? It isn't that qualitative reasoning eschews mathematics—we didn't and we don't. It's not that our examples are toys, as in any young field some are, but many aren't. Sure we should each strive more to exploit traditional mathematical tools *when* appropriate. Sure we should carefully select examples (big or small) that best drive our research toward significant issues, carefully tempered by the fact that what is doable varies depending on the maturity of the research vector pursued. However, let's not fool ourselves that movement in this direction is the result of anything that any of us have said here; the game is already afoot and is doing quite nicely. The prolegomena is not that Sacks' work is far superior for theoretical dynamics, or for example our work is far superior for copier diagnosis. Independent of all the assertions floating around, neither has been argued effectively.

What should truly concern all of us is that our dreams are of artificial engineers and scientists, yet the issue that 20 of us have been asked to comment on in front of the international AI community is whether we think numerical phase portrait analyzers or QSIM-style qualitative simulation is a more effective tool for academic dynamicists! Spending the next 10 years polishing either approach won't get us any closer to our dreams. We've fallen into a black hole of subgoals and will fall only deeper if we expend all our energy blindly adding deep, deep mathematics and modeling subsystems in service of these two masters. If there is a prolegomena to any future qualitative physics (or scientific and engineering reasoning), then it's that our dreams must be attacked directly. This only happens as we grapple with top-level tasks.

REFERENCES

- BERLEANT, D., and B. KUIPERS. 1990. Combined qualitative and numerical simulation with Q3. *In Proceedings of the Workshop on Qualitative Physics*, Lugano, Switzerland. pp. 140–152.
- BROWN, J. S., R. R. BURTON, and J. DE KLEER. 1982. Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III. *In Intelligent tutoring systems*. Edited by D. Sleeman and J. S. Brown. Academic Press, New York. pp. 227–282.
- DAGUE, P., O. RAIMAN, and P. DEVES. 1987. Troubleshooting: when modeling is the trouble. *In Proceedings AAAI-87*, Seattle, WA. pp. 590–595. Also in Weld and de Kleer (1990).
- DE KLEER, J. 1979. Causal and teleological reasoning in circuit recognition. AI-TR-529, Artificial Intelligence Lab., MIT, Cambridge, MA.
- DE KLEER, J. 1984. How circuits work. *Artificial Intelligence*, 24:205–280.
- DE KLEER, J. 1990. Qualitative physics: a personal view. *In Readings in qualitative reasoning about physical systems*. Edited by D. S. Weld and J. de Kleer Morgan Kaufmann.

- DE KLEER, J., and J. S. BROWN. 1984. A qualitative physics based on confluences. *Artificial Intelligence*, **24**:7–84. Also in Weld and de Kleer (1990).
- DE KLEER, J., and B. C. WILLIAMS. 1991. Qualitative reasoning about physical systems II. Special Issue of the *Artificial Intelligence Journal*, Cambridge: MIT Press.
- FALKENHAINER, B., and K. D. FORBUS. 1988. Setting up large-scale qualitative models. *In Proceedings of AAAI-88*, St. Paul. pp. 301–306. Also in Weld and de Kleer (1990).
- FORBUS, K. D. 1984. Qualitative process theory. *Artificial Intelligence*, **24**:85–168. Also in Weld and de Kleer (1990).
- FORBUS, K., and B. FALKENHAINER. 1990. Self-explanatory simulations: an integration of qualitative and quantitative knowledge. *In Proceedings of AAAI-90*, Boston. pp. 380–387.
- GLASSER, L. A., and D. W. DOBBERPUHL. 1985. *The design and analysis of VLSI circuits*. Addison-Wesley.
- GLEICK, J. 1987. *Chaos: making a new science*. Sphere.
- GRAY, P., and R. MEYER. 1984. *Analysis and design of analog integrated circuits*. Wiley.
- KUIPERS, B. 1986. Qualitative simulation. *Artificial Intelligence* **29**:289–338.
- MAVROVOUNIOTIS, M., and G. STEPHANOPOULOS. 1988. Formal order-of-magnitude reasoning in process engineering. *Comput. Chem. Engng.*, **12**(9/10):867–881. Also in Weld and de Kleer (1990).
- MCALLESTER, D. 1981. Algebraic approximation. *In Proceedings IJCAI-81*, Vancouver, BC. pp. 624–628.
- MOSES, J. 1971. Algebraic simplification: a guide for the perplexed. *Comm. ACM* **14**(8):527–537.
- RAIMAN, O. 1986. Order of magnitude reasoning. *In Proceedings AAAI-86*, Philadelphia, PA. pp. 100–104.
- ROYLANCE, G. 1980. A simple model of circuit design. AI-TR-703, Artificial Intelligence Lab., MIT, Cambridge, MA.
- SACKS, E. 1990a. Automatic qualitative analysis of dynamic systems using piecewise linear approximations. *Artificial Intelligence*, **41**:313–364.
- SACKS, E. 1990b. A dynamic systems perspective on qualitative simulation. *Artificial Intelligence*, **42**:349–362.
- SACKS, E. 1991. Automatic analysis of one-parameter planar ordinary differential equations by intelligent numeric simulation. *Artificial Intelligence*, **48**:27–56.
- SIMMONS, R. G. 1988. Combining associational and causal reasoning to solve interpretation and planning problems. AI-TR-1048, Artificial Intelligence Lab., MIT, Cambridge, MA.
- SKORSTAD, G., and K. D. FORBUS. 1989. Qualitative and quantitative reasoning about thermodynamics. *In Proceedings of the eleventh annual conference of the cognitive science society*, Ann Arbor, MI. pp. 892–899.
- VLADIMIRESCU, A., K. ZHANG, A. R. NEWTON, D. O. PEDERSON, and A. SANGIOVANNI-VENCEN-TELLI. 1981. *SPICE version 2G user's guide*, University of California, Berkeley.
- WELD, D. S., and J. DE KLEER. 1990. *Readings in qualitative reasoning about physical systems*. Morgan Kaufmann.
- WILLIAMS, B. C. 1984. Qualitative analysis of MOS circuits. *Artificial Intelligence*, **24**:281–346. Also a revised version in Weld and de Kleer (1990) as Temporal qualitative analysis: explaining how physical systems work.
- WILLIAMS, B. C. 1986. Doing time: putting qualitative reasoning on firmer ground. *In Proceedings AAAI-86*, Philadelphia, PA. pp. 105–113. Also in Weld and de Kleer (1990).
- WILLIAMS, B. C. 1991. A theory of interactions: unifying qualitative and quantitative algebraic reasoning. *Qualitative reasoning about physical systems II*, Special Issue of the *Artificial Intelligence Journal*, Cambridge: MIT Press.
- WILLIAMS, B. C., and J. DE KLEER. 1991. Qualitative reasoning about physical systems: a return to roots. *Qualitative reasoning about physical systems II*, Special Issue of the *Artificial Intelligence Journal*, Cambridge: MIT Press.
- WINSTON, P. 1984. *Artificial intelligence*. 2nd ed. Addison-Wesley, Reading, MA.