BIBLIOGRAPHY

D. G. Bobrow and P. J. Hayes, eds., *Artif. Intell.* 24(1–3) (1984).

J. de Kleer and J. S. Brown, "Assumptions and Ambiguities in Mechanistic Mental Models," in Gentner and Stevens, 1983, pp. 155–190.

J. de Kleer and J. S. Brown "A Qualitative Physics Based on Confluences," in Bobrow and Hayes, 1984, pp. 7–83.

K. D. Forbus, "Qualitative Reasoning about Space and Motion," in Gentner and Stevens, 1983, pp. 53–73.

K. D. Forbus, "Qualitative Process Theory," in Bobrow and Hayes, 1984, pp. 87–168.

D. Gentner, "Flowing Waters or Teaming Crowds: Mental Models of Electricity," in Gentner and Stevens, 1983, pp. 99–129.

D. Gentner and A. L. Stevens, eds., *Mental Models,* Lawrence Erlbaum, Hillsdale, N.J., 1983.

P. Hayes, "The Naive Physics Manifesto," in Hobbs and Moore, 1985, pp. 1–36.

J. R. Hobbs and R. C. Moore, eds., *Formal Theories of the Commonsense World,* Ablex, Norwood, N.J., 1985.

B. Kuipers, "Commonsense Reasoning about Causality: Deriving Behavior from Structure," in Bobrow and Hayes, 1984, pp. 169–203.

S. L. HARDT
Bellcore

PHYSICS, QUALITATIVE

Qualitative physics, like conventional physics, provides an account of behavior in the physical world. The vision of qualitative physics is to provide, in conjunction with conventional physics, a broad integrated and formal account of behavior—an account rich enough to enable intelligent systems to perform tasks such as design, diagnosis, analysis, explanation and simulation (the emphasis of this article). However, unlike conventional physics, qualitative physics predicts and explains behavior in qualitative terms.

The behavior of a physical system can be described by the exact quantitative values of its variables (forces, velocities, positions, pressures, etc) at each time instant. Such a description, although complete, fails to provide much insight into how the system functions. The insightful concepts and distinctions are usually qualitative, but they are embedded within the much more complex framework established by continuous real-valued variables and differential equations. Qualitative physics is an alternative physics in which these concepts are defined within a far simpler, but nevertheless formal, symbolic qualitative basis.

It is important to note that qualitative physics yields qualitative descriptions of behavior based on qualitative descriptions of the physical situation and physical laws. The key contribution that makes qualitative physics useful and possible is that moving to the qualitative level preserves many of the important behavioral distinctions. For example, important concepts and distinctions underlying behavior include state, cause, law, equilibrium, oscillation, momentum, quasistatic approximation, contact force, and feedback. These terms are qualitative and can be intuitively understood.

Qualitative physics is, perhaps more than anything, a long-term research program. A great deal of further research is required before qualitative physics can even come close to explaining the range of physical phenomena accounted for by conventional physics. Much of the research on qualitative physics recapitulates fundamental physical and mathematical investigations that took place centuries ago. Driven by the necessity to formalize common sense and enabled by the idea of computation and the modeling techniques of artificial intelligence, these new theories characterize what these investigations took for granted.

Qualitative physics is an active area of artificial intelligence research. Although this research is unified in its goal to account for physical behavior qualitatively, this goal is addressed with a great deal of technical, notational and methodological diversity. For conciseness this article adopts the point of view of de Kleer and Brown (1984) and explains the alternative proposals in its terms.

QUANTITATIVE VS QUALITATIVE

Figure 1 illustrates the approach of qualitative physics contrasted with conventional physics. Both start by modeling physical situation, both end with a qualitative commonsense description of the behavior. The first step in the conventional approach is to formulate and solve the differential equations to obtain a solution. The second step is to interpret this solution to obtain a commonsense description of the behavior. The qualitative analysis begins by formulating qualitative differential equations and then solves these. The result is a similar commonsense description of the behavior that is obtained more simply, and a causal explanation for that behavior.

Unlike quantitative variables, qualitative variables can only take on one of a small number of values. Each qualitative value corresponds to some (disjoint) interval on the real number line. *Landmark values* demarcate the boundaries between qualitatively distinguishable intervals. The most common landmark is zero, in which case the three important qualitative values are positive, negative, and zero or, for a derivative, whether a quantity is increasing, decreasing, or constant. This simple, but most important, quantity space with respect to the landmark zero consists of only three values: +, −, and 0.

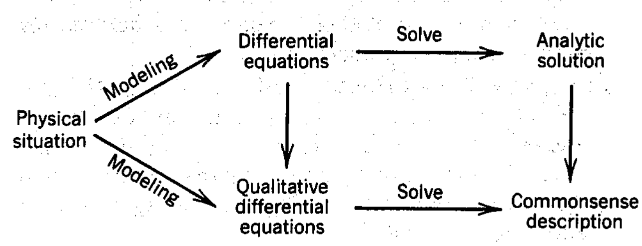

Figure 1. Qualitative vs quantitative.

ing program erroneously introduces assumptions of the function of a device into its constituent components, then the program will miss faults and symptoms or incorrectly identify components as being faulted.

This view raises a difficult question: where do the laws and the descriptions of the device being studied come from? Unless some conditions are placed on the laws and the descriptions, the inferences that can be made may be (implicitly) preencoded in the structural description or the component model library.

The no-function-in-structure principle is central: the laws of the parts of the device may not presume the functioning of the whole. Take as a simple example a light switch. The model of a switch that states, "if the switch is off, no current flows; and if the switch is on, current flows," violates the no-function-in-structure principle. Although this model correctly describes the behavior of the switches in some offices, it is false because there are many closed switches through which current does not necessarily flow (such as two switches in series). Current flows in the switch only if it is closed and there is a potential for current flow.

Without this principle, qualitative physics would be an architecture for building handcrafted (and thus *ad hoc*) theories for specific devices and situations. It would provide no systematic way of ensuring that a particular class of laws did not already have built into them the answers to the questions that the laws were intended to answer. Therefore, there would be no guarantee that the same models would accurately describe the behavior the same constituents in new systems.

QUALITATIVE PHYSICS ANALYSIS

Structure

To do a qualitative physics analysis a program must be provided with a description of the physical situation about which it must draw inferences. There are three main approaches: constraint based, component based, and process based.

The constraint-based approach (Kuipers, 1984) describes the physical situation directly in terms of a set of variables and constraints relating those variables. For example, Figure 2 presents the constraint structure description for a simple heat-flow system. The description contains three constraints: a qualitative adder, a qualitative proportionality, and a qualitative differential. These constraints relate three variables: T the temperature of the material, T_s the temperature of the source of heat, ΔT the temperature difference, and inflow the resulting rate of heat flow into the material.

The process-based approach (Forbus, 1984) describes the physical situation in terms of the physical processes that are potentially present. Intuitively, a process is something that causes changes in objects over time. For example, flowing, bending, heating, cooling, stretching, compressing, and boiling are all processes in qualitative process theory.

The component-based approach (de Kleer and Brown, 1984; Williams, 1984a; Weld, 1986) is reductionist: the behavior of a physical structure is accounted for by the behaviors of its physical constituents. The situation has been described in terms of the organic molecules present (Weld, 1986). Other authors distinguish three kinds of constituents: materials, components, and conduits (de Kleer and Brown, 1984; Williams, 1984a). Physical behavior is accomplished by operating on and transporting materials such as water, air, electrons, etc. The components are constituents that can change the form and characteristics of material. Conduits are simple constituents that transport material from one component to another and cannot change any aspect of the material within them. Some examples of conduits are pipes, wires, and cables. The physical structure of a device can be represented by a topology in which nodes represent components and edges represent conduits. Figure 3a illustrates the device topology of the pressure regulator diagramed in Figure 3b. In this device topology, the conduits IN and OUT transport material of type fluid and conduit FP transports material of type force. The control on valve VV adjusts the area available for flow, and SNS senses the output pressure to control the value inversely (so that if pressure rises, area available for flow decreases).

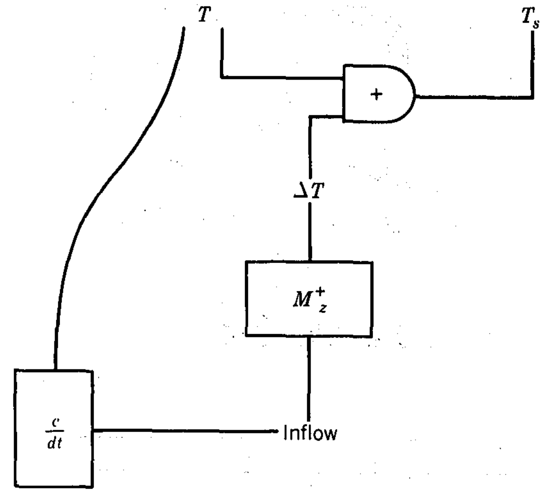

Figure 2. A qualitative constraint structure.

Each modeling approach has its advantages and disadvantages. The constraint-based approach makes it possible to sidestep the modeling constituents at the cost of ignoring the no-function-in-structure principle. The process-based approach is extremely general, including the capability for creating and destroying constituents as well as rearranging objects. This capability introduces a large amount of inefficiency, particularly for systems with large numbers of interconnections (eg, electrical and fluid circuits). The component-based approach is efficient, more easily obeys no-function-in-structure, and is ideally suited for systems with a fixed interconnect topology.

The naive physics approach (Hayes, 1979) does not really fit into the classifications of this article. It advocates the large-scale axiomatizations of commonsense domains in predicate calculus. Although the spirit of this proposal guides much of the research, most qualitative

Process heat-flow

Individuals:
src an object, Has-Quantity(src, heat)
dst an object, Has-Quantity(dst, heat)
path a Heat-Path, Heat-Connection(path, src, dst)

Preconditions:
Heat-Aligned(path)

Quantity conditions:
A[temperature(src)] > A[temperature(dst)]

Relations:
Let flow-rate be a quantity
A[flow-rate] > ZERO
flow-rate $\propto_{Q^-}$ (temperature(src) − temperature(dst))

Influences:
1−(heat(src), A[flow-rate])
1+(heat(dst), A[flow-rate])

Figure 4. Heat flow process.

M. For example, the inflow $= M_z^+(\Delta T)$ of Figure 2 indicates that inflow is a strictly monotonically increasing function (indicated by the superscript +) of ΔT and that this function passes through zero (indicated by the subscript z). This proportionality corresponds to [inflow] = $[\Delta T]$ and ∂inflow $= \partial\Delta T$ in confluence notation. The analogous operator of Forbus (1984) is $\propto_{Q^+}$ (Fig. 4).

Quantity Spaces

The +, −, 0 value space is insufficient for many applications. Often some variable has multiple landmarks above or below which a component's behavior is governed by differing confluences. For example, water temperature has two landmarks 0°C and 100°C. Water is a liquid between these two landmarks, a solid below them, and a vapor above them. Thus a network of inequalities may exist among the variables.

The choice of landmarks for variables is a serious problem. Clearly, zero is an important landmark for derivatives, but what is the origin of the landmarks for other variables? The phase transition temperatures for water are obvious landmarks. Landmarks should either be defined by the component models themselves (as is the case with water temperature) or inferred by the qualitative physics analysis. The seductive scheme of choosing simple symbolic vocabularies (eg, hot, very hot, etc) is problematic. Such arbitrarily chosen schemes are based on the particular situation being analyzed. It is always possible to choose just the right symbolic vocabulary for each variable *after* analyzing the situation. What is very hot for some task may be just hot for another. Choosing a vocabulary arbitrarily produces a model, a solution, and an interpretation that have the appearance of cogency, but are, in fact, vacuous because the appearance of success depends on knowing the answer beforehand. Most implementations require extensive computation with inequalities. Forbus (1984) utilizes a general quantity space representation to reason with inequalities in qualitative process theory.

Ambiguities

Qualitative reasoning is inherently ambiguous and thus may make multiple behavioral predictions (called interpretations). At a minimum for a prediction to be correct, one of the interpretations must correspond to the actual behavior of the real system. A stronger criterion follows from observing that the structural description of a particular device implicitly characterizes a wide class of physically realizable devices with the same topology. Ideally, the behavior of each device in the class is described by one of the interpretations, and every interpretation describes the behavior of some device in the class. For example, every possible interpretation of the pressure regulator (Fig. 3**a**) corresponds to some physically possible regulator behavior, and every physically possible regulator behavior is described by some interpretation.

Unfortunately, this criterion is too strong (Kuipers, 1986). There are devices that have interpretations that are not physically realizable. For example, if the quantitative model was $x + y + z = 0$ and $y + z = 0$, the confluences would be $[x] + [y] + [z] = 0$ and $[y] + [z] = 0$. The qualitative operations of Tables 1 and 2 place no constraints on the value of $[x]$. The original quantitative model indicates that $x = 0$, but this cannot be inferred from the confluences alone. The same two confluences also describe the quantitative equations $x + 2y + z = 0$ and $y + z = 0$ from which it is not possible to infer $x = 0$.

Modeling

In the process- or component-based approach qualitative analysis must construct a qualitative model consisting of constraints from the structural description. (This step is avoided in constraint-based approaches in which the physical situation is initially described in terms of constraints.) In the component-based approach each type of physical constituent has a distinct model. The component model characterizes all the potential behaviors that the component can manifest. The lawful behavior of a component is expressed by a set of possible states and their associated specifications and confluences. For example, a valve could be modeled by A the area available for flow, P the pressure across the valve, and Q the flow through the valve. In state OPEN, the valve functions as a simple conduit, there is no pressure drop across it, and the flow through it is unconstrained. Neither can the pressure across it change; that can only be caused by a change in position of the valve. The state CLOSED is the dual to state OPEN. In it the valve completely disconnects the input from the output. There is no flow through the valve and the pressure across it is unconstrained. The flow through it cannot change without changing the area available for flow. In the WORKING state, the valve acts like a fluid resistance; its resistance controlled by A. For example, if $\partial A = 0$, then $\partial P = \partial Q$. This can be encoded as:

$$\text{OPEN: } [A = A_{\max}],\ [P] = 0,\ \partial P = 0$$

$$\text{WORKING: } [0 < A < A_{\max}],\ [P] = [Q],\ \partial P + [P]\partial A - \partial Q = 0$$

$$\text{CLOSED: } [A = 0],\ [Q] = 0,\ \partial Q = 0.$$

one component suffices in the component-based approach. For example, an additional instance of the heat flow process is required if heat could flow in either direction, while the same component set models heat flow in both directions.

Constructing models for larger systems can be quite difficult. There may be multiple models available for system components each making its own approximations. In this case the goal of the modeler is to construct a composite model of the device that is adequate to correctly solve the task at hand but yet simple enough to be solvable with the information and techniques available. These issues have been examined (Addanki and co-workers, 1991; Falkenhainer and Forbus, 1991).

Qualitative Calculus

As time passes, variables change toward landmarks causing transitions to other operating regions and component states (equivalently, processes become active and inactive). At any given time, many variables may be approaching their respective landmarks, and the analysis must determine which variable(s) reach their landmarks first. This process is sometimes called limit analysis (Forbus, 1984) or transition analysis (Williams, 1984a).

A theory of change presumes a theory of time. Roughly speaking, three models of time have been used in qualitative physics. Forbus (1984) and Weld (1986) model time as a sequence of intervals as suggested by Allen (1983). de Kleer and Brown (1984) model time by intervals separated by any number of instants. Williams (1984a), Kuipers (1984), and de Kleer and Bobrow (1984) model time as intervals separated by single instants. This latter approach is used in the remainder of the article.

The rules necessary to reason about change over time derive directly from the conventional calculus (Williams, 1984a, 1984b), particularly the mean value theorem. These rules, in essence, solve the qualitative differential equations constructed by the modeling. If variables are differentiable, then the rules for determining time behavior are remarkably simple, The following five rules apply to all derivative orders:

A. *Instant to Interval.* Any nonzero quantity, must remain nonzero during the following interval; if a quantity $[x]$ is zero at the instant, then on the following interval it must obey the integration constraint $[x] = \partial x$.

B. *Interval to Instant.* A nonzero quantity may become zero on the following interval if and only if $[x] = -\partial x \neq 0$.

C. *Continuity.* Variables change continuously; continuous changes are between 0 and + or − (in either direction), but not between + and −.

D. *Contradiction Avoidance.* A transition is only possible if the resulting state satisfies the qualitative equations for that state.

E. *Analyticity.* A quantity that is zero for any interval, is zero for all time; conversely, a quantity that is nonzero at some time cannot become identically zero.

Envisioning

Envisioning is a reasoning process that uses the modeled device to produce a description of the behaviors of the device over time. An attainable envisionment describes all possible behaviors given some initial state. A total envisionment describes all possible behaviors for each possible initial state.

The envisioning process can be illustrated by examining the diaphragm–spring–stem fragment of the pressure regulator. If the input pressure increases, the output pressure increases, producing a force on the diaphragm. This force acts against the spring force and friction. The valve slowly gains velocity as it closes; however, by the time it reaches the position where the force exerted by the pressure balances the restoring force of the spring, the valve has built up a momentum causing it to move past its equilibrium position, thus reducing the pressure below what it should be. As it has overshot its equilibrium the spring pushes it back, but by the same reasoning, the valve overshoots again, thereby producing ringing or oscillation. Figure 5 illustrates the essential details: a mass situated on a spring and shock absorber (ie, friction).

The envisioning process of (de Kleer and Bobrow, 1984) is based on the five rules of the qualitative differential calculus:

1. Start with some initial state(s) (usually at an instant).
2. Identify those quantities that are moving to their landmarks.
3. Construct partial descriptions of tentative next states using rule A or B.
4. Using rules C, D and E expand and prune the possible next states.
5. For each state not yet analyzed, go to step 2.

Note that envisioning can proceed in parallel. The resultant state diagram is the envisionment for the system.

The behavior of the mass is described by Newton's law $F_m = ma$, or qualitatively $[F_m] = \partial v$ (more generally, $\partial^n F_m = \partial^{n+1} v$). Hooke's law for the spring $F_s = -kx$ becomes $\partial^{n+1} F_s = -\partial^n v$. The resistance of the shock absorber is modeled by $\partial^n F_f = -\partial^n v$. For simplicity sake, define $x = 0$ as the mass position with the spring at equilibrium, and $x > 0$ to be to the right. The net force on the mass is provided by the spring and shock absorber: $F_m = F_s + F_f$, or qualitatively $\partial^n F_m = \partial^n F_s + \partial^n F_f$.

Suppose the system is started by stretching the mass to the right. At this instant the velocity is zero, but the mass

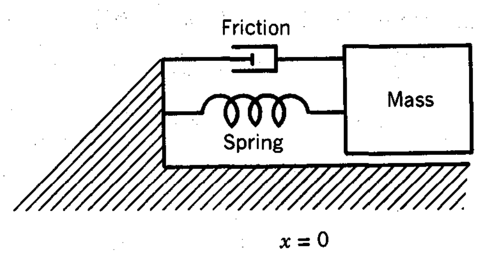

Figure 5. Mass–spring–friction system.

depends on every other, there is no simple dependence among these quantities. Some researchers (Simon, 1977; Iwasaki and Simon, 1986) argue that as a consequence there is no causality in feedback loops. Others (de Kleer and Brown, 1984; Williams, 1984a) introduce additional heuristics that identify the causality in feedback loops approximating conventional engineering intuitions (ie, causality goes through the signal path from input to output and then back through the error path).

Causality is an ontological commitment. There is an, often implicit, presupposition in much of qualitative physics research: by developing a notion of causality that is in close alignment with how the physical system actually functions, the resulting explanations will be more powerful. However, this is an unnecessarily extreme position. It is sufficient that present society embodies a notion called *causality*. By formalizing that notion, qualitative physical analysis can describe device functioning in familiar terms and access society's accumulated knowledge of device behavior.

CURRENT AND FUTURE RESEARCH

Compared to conventional physics, qualitative physics represents a research direction that has only just started. Although significant progress has been made, relatively few physical phenomena have been accounted for. The following are a few of the current directions of qualitative physics research.

Most of the qualitative physics research presumes that the physical structure can be modeled by a relatively small number of distinct objects and interactions. However, many physical situations are, in fact, modeled by a large number of identical objects: water flow in a river, heat flow along a slab, molecules in a pipe, etc. Such systems cannot be accounted for in current qualitative physics. More formally, qualitative physics has [with few exceptions (Weld, 1986)] thus far focused on lumped-parameter systems that are described by ordinary differential equations. Little progress has been made on distributed systems whose behavior is described by partial differential equations with respect to spatial variables.

Qualitative physics has tended to focus on dynamics. Reasoning about spatial movement of objects and shapes of objects and their interactions is extremely difficult and less progress has been made (de Kleer, 1975; Forbus, 1983). Examples of recent progress (Nielsen, 1983; Gelsey, 1987) and discussion of some of the difficulties inherent in spatial reasoning (Davis, 1987) have been published.

A much more sophisticated notion of time is required. Thus far time has been an implicit, not explicit, variable, making it difficult to express certain laws (delays) and to reason about the consequences of the events without analyzing every intermediate event as well. Systems oscillate, approach asymptotes, and become unstable. Most qualitative physics cannot account for such long-term behavioral effects. Most qualitative physics has assumed that the underlying functions are well behaved, continuous, and differentiable. Often systems go through discontinuous transitions that are produced or produce impulses. A qualitative theory of such generalized functions is required and a start has been made (Nishida and Doshita, 1987).

Humans must learn commonsense physics from interactions with the world. Some suggestions about learning qualitative laws have been published (Forbus and Genter, 1983). Complex system can be described at many levels of abstraction. Currently, qualitative physics does not include any notion of hierarchy or when it is necessary to move to other levels of abstraction.

BIBLIOGRAPHY

S. Addanki, R. Cremonini, and J. S. Penberthy, "Graph of Models," in de Kleer and Williams, 1991.

A. Collins and D. Gentner, "Multiple Models of Evaporation Processes," in *Proceedings of the Fifth Annual Meeting of the Cognitive Science Society*, Rochester, N.Y., 1983.

E. Davis, "A Framework for Qualitative Reasoning about Solid Objects," in G. Rodriquez, ed., *Proceedings of the Workshop on Space Telerobotics*, NASA and the Jet Propulsion Laboratory, 1987, pp. 369–375.

J. de Kleer, *Qualitative and Quantitative Knowledge in Classical Mechanics*, TR-352, AI Laboratory, MIT, Cambridge, Mass., 1975.

J. de Kleer, "How Circuits Work," in D. G. Bobrow, ed., *Qualitative Reasoning about Physical Systems*, MIT Press, Cambridge, Mass., 1985.

J. de Kleer and D. G. Bobrow, "Higher-Order Qualitative Derivatives," in *Proceedings of the Fourth National Conference on Artificial Intelligence*, Austin, Tex., AAAI, Menlo Park, Calif., 1984, pp. 86–91.

J. de Kleer and J. S. Brown, "The Origin, Form, and Logic of Qualitative Physical Laws," in *Proceedings of the Eighth IJCAI*, Karlsruhe, FRG, Morgan-Kaufmann, San Mateo, Calif., 1983a, pp. 1158–1169 [ambiguity].

J. de Kleer and J. S. Brown, "Assumptions and Ambiguities in Mechanistic Mental Models," in D. Gentner and A. L. Stevens, eds., *Mental Models*, Lawrence Erlbaum, Hillsdale, N.J., 1983b, pp. 155–190.

J. de Kleer and J. S. Brown, "A Qualitative Physics Based on Confluences," *Artif. Intel.* **24**, 7–83 (1984).

J. de Kleer and J. S. Brown, "Theories of Causal Ordering," *Artif. Intell.* **29**, 33–62 (1986).

J. de Kleer and B. C. Williams, eds., *Qualitative Reasoning about Physical Systems II, Special Issue, Artificial Intelligence* **51**, MIT Press, Cambridge, Mass., (1991).

J. Dormoy and O. Raiman, "Assembling a Device," in *Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, Minn., AAAI, Menlo Park, Calif., 1988.

B. Falkenhainer and K. D. Forbus, "Compositional Modeling," in de Kleer and Williams, 1991.

B. Faltings and P. Struss, *Recent Advances in Qualitative Physics*, MIT Press, Cambridge, Mass., 1991.

K. D. Forbus, "Qualitative Reasoning about Space and Motion," in D. Gentner and A. L. Stevens, eds., *Mental Models*, Lawrence Erlbaum, Hillsdale, N.J., 1983, pp. 53–73.

K. D. Forbus, "Qualitative Process Theory," *Artif. Intell.* **24**, 85–168 (1984).

K. D. Forbus, "The Qualitative Process Engine," in Weld and de Kleer, 1990.

AAAI, Menlo Park, Calif., 1990, pp. 349–356 [relation to conventional calculus; transition ordering].

M. K. Yip, "Understanding Complex Dynamics by Visual and Symbolic Reasoning," in de Kleer and Williams, 1991 [describes the KAM system].

JOHAN DE KLEER
Xerox Palo Alto Research Center

PLANES

PLANES was a natural language front-end developed at the University of Illinois for accessing a relational database of Navy flight and maintenance data. PLANES used a semantic grammar (see GRAMMAR, SEMANTIC); it matched processed inputs against a set of framelike templates (corresponding to well-formed queries) and filled in the templates to generate database queries. PLANES could handle some nongrammatical inputs. [See D. Waltz, "An English Language Question Answering System for a Large Relational Database," *CACM 21* (7), 526–539 (1978), for an overview of the system; and H. Tennant, "Experience With the Evaluation of Natural Language Question Answerers," *Proceedings of the Sixth IJCAI*, Tokyo, Aug. 1979, pp. 874–876, for user evaluations of PLANES.]

DAVID WALTZ
Thinking Machines Corp. and Brandeis University

PLANLOG

A proposal for a logic programming language that incorporates "procedural" features to cope with destructive change (see B. Fronhöfer, "PLANLOG: A Language Framework for the Integration of Logical and Procedural Programming," in J. McDermott, ed., *Proceedings of the 10th International Joint Conference on Artificial Intelligence,* Milan, Aug. 1987, Morgan-Kaufmann Inc., San Mateo, Calif., 1987, pp. 15–17). It is based on a logical plan generation calculus called linear proofs which avoids the frame problem (see W. Bibel, "A Deductive Solution for Plan Generation," *New Generation Computing* **4**, 115–132 (1986)).

BERTRAM FRONHÖFER
Technical University of Munich

PLANNER

A LISP-based AI programming language for inference control, PLANNER was designed in 1972 at the MIT AI Lab by Hewitt and extensively demonstrated by Winograd in his SHRDLU (qv) project [see G. Sussman, T. Winograd, and E. Charniak, "MICRO-PLANNER Reference Manual," AI Memo 203, AI Laboratory, MIT, Cambridge, Mass., 1970, and C. Hewitt, "Description and Theoretical Analysis (Using Schemata) of PLANNER, a Language for Providing Theorems and Manipulating Models in a Robot," Report No. TR-258, AI Laboratory, MIT, Cambridge, Mass., 1972].

A. HANYONG YUHAN
AT&T Bell Laboratories

PLANNING

Planning is the generation of an action sequence or action program for an agent, such as a robot, that can change its environment. The purpose of a plan is to achieve one or more explicit goals. The essential inputs for planning are an initial world state, a repertoire of actions for changing that world, and a set of goals. The form of the plan is commonly just a linear sequence or acyclic directed graph of actions, although the full range of programming control structures are potentially relevant. For planning to be effective, the environment in which the plan will be executed must be largely predictable, but need not be completely deterministic. However, planning will be ineffective for chaotic domains, and an agent can only react to events.

Planning is probably the most reliable method for controlling the behavior of living and artificial agents, but it is not the only method, and not necessarily the fastest. A person certainly does not synthesize a plan to remove his hand from a hot dish. It seems a reasonable hypothesis that people normally plan their behavior only in novel or critical situations. In familiar situations, people probably merely retrieve and apply stored behavior programs. Artificial systems can adopt the same approach. However, application of stored routines, while faster than planning, can lead to inappropriate or unsuccessful behavior. An illustration is the situation where a person enters a darkened room during a power failure and automatically turns on a wall switch even when he knows perfectly well that the power is off.

Planning may require a search through an enormous space of possible plans, and so search control is an important consideration. Other key planning topics are the representation of actions, goal protection, management and modeling of time (see REASONING, TEMPORAL), and the ordering of goals and subgoals for achievement. Dynamic situations present additional challenges, including the detection of plan violations during execution of a plan, replanning in response to changing goals or unexpected events, actions and events with uncertain consequences, and the presence of other agents or physical processes in the environment. Areas of application for planning include automated manufacturing (qv), autonomous vehicles, the control of unmanned spacecraft, and robotics (qv). Planners are large, complex AI artifacts. A state-of-the-art planner requires many years of work to perfect, although a simple planner with numerous limitations can be written in a couple of months. This article emphasizes task planning with logical world models, rather than the