

Reasoning About Fluids Via Molecular Collections

John W. Collins and Kenneth D. Forbus
Qualitative Reasoning Group
Department of Computer Science
University of Illinois

Abstract

Hayes has identified two distinct ontologies for reasoning about liquids. Most qualitative physics research has focused on applying and generalizing his *contained-liquid* ontology. This paper presents a technique for generating descriptions using the *molecular collection* (MC) ontology, a specialization of his alternate ontology which represents liquids in terms of little “pieces of stuff” traveling through a system. We claim that MC descriptions are *parasitic* on the Contained-Stuff ontology, and present rules for generating MC descriptions given a Qualitative Process theory model using contained stuffs. We illustrate these rules using several implemented examples and discuss how this representation can be used to draw complex conclusions.

I. Introduction

Sometimes two distinct but interrelated views of an object or system are needed to reason about the physical world. For example, sometimes an engineer must think of “the liquid in the container” as an object (the *contained-liquid* ontology) while also reasoning about a hypothetical collection of molecules traveling together through the system as an object (the *piece-of-stuff* ontology). Likewise, a river may be viewed either as a static container of water defined by its banks (i.e., the same river it was a century ago), or as a dynamic collection of little pieces of water, each of which retains its identity as it flows to the sea.

As Hayes [6, 7] notes, neither ontology alone suffices to explain commonsense reasoning about liquids. Similarly, neither ontology alone suffices for intelligent computer-aided engineering. It is easy to reason about “the pressure at a portal” in the contained-liquid ontology, but impossible to explain the details of a thermodynamic cycle without following a “piece of stuff” through the system. The *piece-of-stuff* ontology, as we shall show, makes explicit the notions of continuity of space and conservation of matter, but provides no mechanism for reasoning about the overall behavior of the system.

This paper presents a technique for generating and reasoning with descriptions of fluids as “pieces of stuff”. We introduce the *molecular collection* (MC) ontology as a

specialization of Hayes’ *piece-of-stuff* ontology. We claim that the MC ontology is *parasitic* on the Contained-Stuff ontology, in that a description of a system in terms of contained stuffs is a prerequisite to computing its description in MC terms. We show implemented rules for performing this computation, and illustrate their use with several examples. We argue that this representation provides a basis for more complex inferences, and discuss some open problems.

II. The theory of molecular collections

We begin by reviewing the original Hayes ontologies:

Contained-Liquid: Consider the liquid in a container as a single object. If the container is open then it is possible for liquid to leave the container and for new liquid to enter. Contained liquids have a continuous quantity *Amount-of* which may be influenced by various processes (e.g., flow, evaporation, condensation). They may disappear and reappear, as when a cup of coffee is emptied and refilled. In this ontology the two cups of coffee are viewed as the same object.

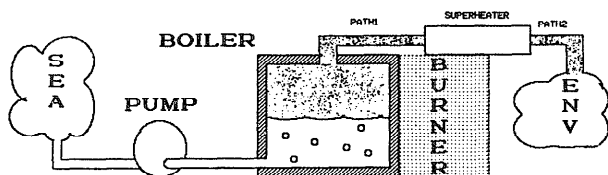
Piece-of-Stuff: Consider a particular collection of molecules as a unit traveling around inside a system. The collection of molecules will have a fixed mass and a continuous position in space, which is influenced by its velocity, which in turn is influenced by various forces acting upon the object. A piece of stuff is never created or destroyed (assuming conservation of mass), so there are fewer problems of changing existence from this ontology.

It is straightforward to generalize the Contained-Liquid ontology into a *Contained-Stuff* ontology that describes gasses and allows multiple substances as well [3]. Qualitative Process theory [3, 4, 5] can be used to generate descriptions of contained stuffs, and we build on those descriptions.

In [7] no restriction is made as to the size of a piece of stuff. We obtain the *molecular collection* (MC) ontology by stipulating that the collection be so small that we can assume it is never distributed over more than one place (we return to this later). This tiny piece of stuff is viewed as a collection of molecules — as opposed to a single molecule

Figure 1: The SWOS problem

This schematic of a Navy Propulsion plant provides an illustration of the importance of the MC ontology. A sophisticated question about this system is, "Given an increase in feedwater temperature, what happens to the steam temperature at the superheater outlet?". Understanding what happens in this situation is one of the hardest problems given at the Surface Warfare Officers' School, in New Port, R.I. The representation developed in this paper provides a basis for answering this question.



— so that it may possess such macroscopic properties as temperature and pressure. Call the arbitrary collection of molecules to be considered as a unit MC.

Any ontology must divide the world into individuals: For reasoning it is important that the number of individuals be few. The Contained-Stuff ontology partitions a fluid system into a few discrete objects using the natural boundaries provided by containment. But the Contained-Stuff ontology fails to preserve molecular identity. Considering individual molecules would be prohibitive and unnecessary, since all the billions of them act more or less alike. By considering the possible behaviors of an anonymous collection of molecules, we constrain the possibilities for the whole by considering only one individual.

We claim that the molecular collection ontology is parasitic on the Contained-Stuff ontology. No one has succeeded in saying anything coherent about establishing the conditions for reasoning with the molecular collection ontology. We believe the reason for this failure is that the MC ontology alone is insufficient. Global information is required to identify what MC is doing. In classical physics the notion of gradient provides a local method for determining such motion. But establishing the gradient requires a global view of the physical system. The Contained-Stuff ontology provides this viewpoint for the MC ontology by establishing paths and conditions for flows and state changes. The reasoning based on molecular collections starts with the results of the Contained-Stuff description, and consequently does not have to re-derive those conclusions.

Consider the system shown in Figure 1. Figuring out how MC moves requires knowing the mass properties of the fluid, viewed with respect to the components of the system. Looking solely at MC, there is no way to establish the pressure differences between system components that

Figure 2: Sample rules for generating MC movements

These rules, associated with particular processes, describe how the MC's place and state change as a consequence of that process acting. Space limitations preclude showing the entire rule set.

```
If Flow(source, destination, path)
  then if Location(MC, source)
        Transition(MC, PLACE, path)
        if Location(MC, path)
            Transition(MC, PLACE, destination)
```

```
If Boiling(substance, container)
  then Transition(MC, STATE, GAS)
```

```
If Condensation(substance, container)
  then Transition(MC, STATE, LIQUID)
```

imply the direction of flow. Although MC must play a role in the solution of the problem, the molecular collection ontology is inadequate. To determine facts like flow direction, the Contained-Stuff ontology must be used. Given a Contained-Stuff description, we can talk about pressure as a function of location rather than trying to find the pressure on an arbitrary collection of molecules. Here, a pump establishes a pressure gradient, causing a liquid flow into the boiler.

Our goal is to construct a history for MC, describing the sequence of places it is in and what is happening to it in those places. For our purposes, MC is uniquely defined by the place it is in, the substance of which it is composed,¹ and its current phase (i.e., solid, liquid or gas). The place is the container or fluid path in which MC resides.²

Constructing the MC history occurs in five steps. The first step is to feed the domain knowledge and the specific example through the Qualitative Process Engine (QPE) to generate the total environment for the given configuration. The total environment consists of all consistent situations connected by the possible transitions between them.³

Second, a single situation is selected for which the MC history is desired. In order for the history to be meaningful and interesting, the situation should involve some active processes and should last for an interval of time.

The third step finds the possible locations and states of MC and establishes how these properties can change. The critical observation is that each active process specifies a fragment of MC's history. Processes operate on objects, some of which (in fluid systems) will be contained stuffs. We can associate rules with each process to describe what, if anything, its activity implies about the location

¹ In this paper only single substance systems are considered.

² Potentially, this could be refined through the use of some coordinate system such as submerged-depth or the length along a path.

³ Each situation represents a unique set of active views and processes taken together with the signs of derivatives for all quantities.

and phase of MC. For example, the rule associated with liquid-flow implies that when MC is in liquid form in the source, it can move into the path of the flow, and end up in the destination of the flow without changing state (see Figure 2). The rule associated with boiling implies that MC will undergo a liquid to gas phase transition within the same location. By combining these partial histories, we can compute the full spatial extent of MC's travels and its associated phase transitions (if any).⁴

In the fourth step, the D_s values for MC's quantities are computed.⁵ By assumption, $D_s[\text{Amount-of}(\text{MC})] = 0$. Pressure is simply inherited from the surrounding contained stuff. Changes in Heat, Temperature, Volume and Height are determined by rules associated with processes. For example, if the temperature of the destination of a liquid flow is greater than the temperature of the source, then both Heat and Temperature of MC will be increasing in the destination. During boiling Heat is increasing and during condensation it is decreasing.

Finally, the fifth step constructs the graph defined by the relevant places and the possible movements between them. From this graph it is easy to recognize such phenomena as branching or cycles of flow. In real fluid systems these histories often branch. For example, steam coming out of a ship's boiler is often tapped off for several different purposes, such as driving the propulsion turbines, running generators to produce electricity, and powering the ship's laundry. The choice of which path to take will depend on the goal of the reasoning. Sometimes it is the properties of a specific path which are of interest. In other cases all paths must be considered. However, it will be assumed that MC retains its identity, i.e., that the molecules of MC never split themselves between two paths. This assumption is realistic if one considers MC to be a tiny subset of the liquid described in the Contained-Stuff view of the same system.

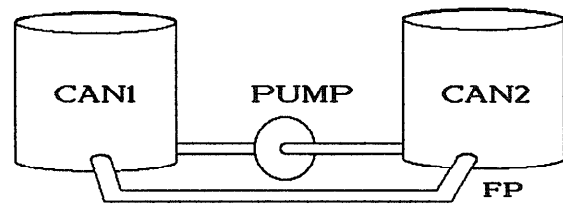
Two observations are relevant here. First, the MC-history generation algorithm is linear in the number of active process instances, making it quite fast.⁶ Second, the relationship between the episodes in the MC history and the states of the envisionment is slightly complicated. One state in the envisionment can give rise to a number of episodes in the MC history. For example, the steady flow of working fluid in a refrigeration system would typically be described as a single state in the envisionment using the Contained-Stuff ontology. But viewed from the MC level, it will give rise to episodes involving heating, liquid-gas phase transition, compression, gas-liquid phase transition, etc.

⁴More than one history can be produced if there are disconnected components in the fluid system. Each history corresponds to a different choice of subsystem for MC.

⁵The D_s value of a quantity is the sign of its derivative.

⁶Running the rules over a total envisionment takes roughly one minute: constructing an MC history for any situation afterwards takes 5-10 seconds.

Figure 3: A simple pumped-flow example



III. Examples

The MC-history generation algorithm has been tested on a number of examples of varying complexity. For concreteness we describe several of them below.

A. Pumped Flow

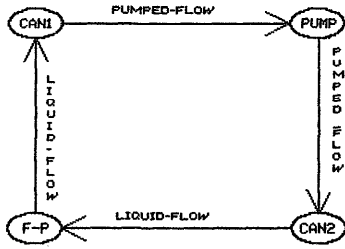
Figure 3 illustrates a scenario consisting of two open containers connected by a pump and a return fluid path. We choose from the envisionment the equilibrium situation, where liquid exists in both containers and the flow rates have equalized. Figure 4 shows the MC history for this situation. The MC history is annotated with information about the type of process responsible for the movement, as well as the derivatives and state (phase) of MC at each place in the history. This information becomes more useful for complex examples, such as in the refrigerator example below. Even though the situation is in steady state (i.e., all derivatives in the Contained-Stuff ontology are zero), the MC history shows that each little piece of stuff in the system undergoes continuous change, both in position and in pressure. However, since there is no way for MC to enter or leave the system, one can conclude that the total amount of stuff in the system is constant.

B. A Refrigerator

One of the motivations for looking at the MC ontology was to allow reasoning about complex thermodynamic cycles such as that used in a refrigerator. Figure 5 shows a simple refrigerator involving six separate processes: two heat flows, two state changes (boiling and condensation), a compressor flow and a liquid flow. As in the pumped flow example, the situation selected for the MC history is the steady state, where all flows have equalized.

Figure 6 shows the MC history. MC boils in the evaporator and then is pumped through the compressor to the condenser, where it returns to the liquid phase and is finally forced through the expansion valve back into the evaporator. This representation provides the foundation for an important class of engineering conclusions. Since MC gains heat during boiling and loses it during condensation, it

Figure 4: The MC history for the pumped-flow example



Location	Can1	Pump	Can2	F-P
Ds[Heat]	0	0	0	0
Ds[Temperature]	0	0	0	0
Ds[Pressure]	0	1	0	-1
Ds[Volume]	0	0	0	0
Ds[Height]	0	0	0	0

Figure 5: A refrigerator

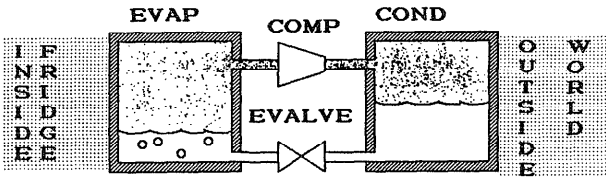
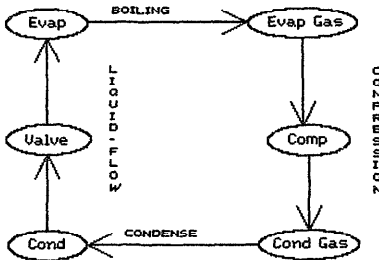
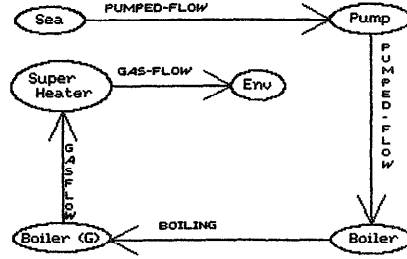


Figure 6: The refrigerator MC history



Location	Evap	Evap	Comp	Cond	Cond	EValve
State	Liq	Gas	Gas	Gas	Liq	Liq
Ds[Heat]	1	0	1	-1	0	-1
Ds[Temp]	-1	0	1	1	0	-1
Ds[Press]	0	0	1	0	0	-1
Ds[Volume]	1	0	-1	-1	0	0
Ds[Height]	1	1	0	-1	-1	0

Figure 7: The SWOS MC history



Location	Sea	Pump	B	B	P1	S-H	P2	Env
State	L	L	L	G	G	G	G	G
Ds[Heat]	0	0	1	0	-1	1	-1	0
Ds[Temp]	0	0	1	0	-1	1	-1	0
Ds[Press]	0	1	0	0	-1	0	-1	0
Ds[Volume]	0	0	1	0	1	1	1	0
Ds[Height]	0	0	1	1	0	0	0	0

must be moving more heat through the compressor than returns via the expansion valve, so there is a net heat flow from the evaporator to the condenser. Thus the refrigerator is pumping heat *uphill* to a higher temperature.

C. The SWOS problem

Here we return to the Navy propulsion plant scenario of Figure 1. Figure 7 shows the MC history. The result of an increased feedwater temperature can in principle be calculated by a differential qualitative analysis (DQ) based on this history [4]. Roughly, the increased temperature means that the boiling episode is shorter, making the steam generation rate higher. The higher steam generation rate means the steam spends less time in the superheater, hence less heat will be transferred, implying a lower temperature at the superheater outlet. Weld (in press) describes a set of DQ rules which, combined with this representation, may be powerful enough to draw this conclusion.

IV. Discussion

The ability to reason with multiple views of a situation provides significant advantages over using a single ontology. The Contained-Stuff ontology provides the conditions to determine which processes are active, and thereby determines the overall behavior of the system. The MC ontology provides the complementary ability to reason about where a piece of stuff came from and where it might go. We demonstrated that MC histories can be easily computed from QP models of fluids organized around Contained-Stuffs, and argued that this representation provides the basis for several important engineering inferences (i.e., closed-cycles, recognition of heat pumps and differential analysis).

It is unclear whether or not growing an MC history across transitions between situations in the Contained-Stuff ontology is a good idea. If one is considering a liquid system that oscillates, for instance, then this could be necessary. However, most questions that arise in engineering concerning the MC history are about steady-state behavior, i.e., a single situation in the Contained-Stuff ontology.

The MC ontology is based on infinitesimal pieces of fluid. It may be possible to generalize it to spatially extended pieces of stuff. This generalization would provide the ability to, for example, identify the spread of a contaminate through a fluid system.

We have only begun to explore the reasoning potential of the MC ontology. Currently we are implementing rules to calculate quantity space information involving MC parameters. Furthermore, we plan to augment the MC history by associating equations with each movement. These equations will be combined to yield quantitative descriptions of relevant system parameters, such as efficiency or work output per pound of working fluid. A differential qualitative (DQ) analysis could then be performed to identify how these parameters could be optimized, or in general how a change in one quantity will affect the behavior of the system.

V. Acknowledgements

Brian Falkenhainer, John Hogge, and Gordon Skorstad provided valuable comments, both in developing the theory and in writing the paper. John Hogge's ZGRAPH program has been invaluable for displaying the MC histories and total envisionment graphs. This research was supported by the National Aeronautics and Space Administration, Contract No. NASA NAG 9137, and by an IBM Faculty Development award.

References

- [1] de Kleer, J. "An Assumption-Based Truth Maintenance System", *Artificial Intelligence*, **28**, 1986
- [2] de Kleer, J. and Brown, J. "A Qualitative Physics based on Confluences", *Artificial Intelligence*, **24**, 1984
- [3] Forbus, K. "Qualitative Process Theory" *Artificial Intelligence*, **24**, 1984
- [4] Forbus, K. "Qualitative Process Theory" MIT AI Lab Technical report No. 789, July, 1984.
- [5] Forbus, K. "The Problem of Existence", Proceedings of the Cognitive Science Society, 1985.
- [6] Hayes, P. "The Naive Physics Manifesto" in *Expert systems in the Micro-Electronic Age*, D. Michie (Ed.), Edinburgh University Press, 1979
- [7] Hayes, P. "Naive Physics 1: Ontology for Liquids" in Hobbs, J. and Moore, B. (Eds.), *Formal Theories of the Commonsense World*, Ablex Publishing Corporation, 1985.
- [8] Kuipers, B. "Common Sense Causality: Deriving Behavior from Structure" *Artificial Intelligence*, **24**, 1984
- [9] Shearer, J., Murphy, A., and Richardson, H. *Introduction to System Dynamics* Addison-Wesley Publishing Company, Reading, Massachusetts, 1967
- [10] Simmons, R. "Representing and reasoning about change in geologic interpretation", MIT Artificial Intelligence Lab TR-749, December, 1983
- [11] Weld, D. "Switching Between Discrete and Continuous Process Models to Predict Genetic Activity", MIT Artificial Intelligence Lab TR-793, October, 1984
- [12] Williams, B. "Qualitative Analysis of MOS Circuits", *Artificial Intelligence*, **24**, 1984