

This article was downloaded by: [Northwestern University]

On: 09 September 2014, At: 11:18

Publisher: Routledge

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Interactive Learning Environments

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/nile20>

Using Qualitative Physics to Build Articulate Software for Thermodynamics Education: A Preliminary Report

Kenneth D. Forbus^a & Peter B. Whalley^b

^a Qualitative Reasoning Group, The Institute for the Learning Sciences, Northwestern University, 1890 Maple Avenue, Evanston, IL, 60201, USA E-mail:

^b Department of Engineering Science, Oxford University, Parks Road, Oxford, OX13PJ, UK E-mail:
Published online: 28 Jul 2006.

To cite this article: Kenneth D. Forbus & Peter B. Whalley (1998) Using Qualitative Physics to Build Articulate Software for Thermodynamics Education: A Preliminary Report, *Interactive Learning Environments*, 5:1, 19-32, DOI: [10.1080/1049482980050102](https://doi.org/10.1080/1049482980050102)

To link to this article: <http://dx.doi.org/10.1080/1049482980050102>

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or

indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms & Conditions of access and use can be found at <http://www.tandfonline.com/page/terms-and-conditions>

Using Qualitative Physics to Build Articulate Software for Thermodynamics Education: A Preliminary Report

KENNETH D. FORBUS
PETER B. WHALLEY

One of the original motivations for research in qualitative physics was the development of intelligent learning environments for physical domains and complex systems. This paper demonstrates how a synergistic combination of qualitative physics and other AI techniques can be used to create an intelligent learning environment for students learning to analyze and design *thermodynamic cycles*. Pedagogically this problem is important because thermodynamic cycles express the key properties of systems which interconvert work and heat, such as power plants, propulsion systems, refrigerators, and heat pumps, and the study of thermodynamic cycles occupies a major portion of an engineering student's training in thermodynamics. This paper describes CyclePad, a fully implemented learning environment which captures a substantial fraction of a thermodynamics textbook's knowledge and is designed to scaffold students who are learning the principles of such cycles. We analyze the combination of ideas that made CyclePad possible, comment on some lessons learned about the utility of various techniques, and describe our classroom experimentation in progress.

INTRODUCTION

One of the central motivations for research into qualitative physics has been its potential for the construction of intelligent learning environments. By providing computational ac-

counts of human reasoning about the physical world, ranging from what the person on the street knows to the extensive expertise of scientists and engineers, qualitative physics should

Interactive Learning Environments
Volume 5, pages 19-32
Copyright © 1998 by Swets & Zeitlinger
All rights of reproduction in any form reserved.
ISSN: 1049-4820

Kenneth D. Forbus, Qualitative Reasoning Group, The Institute for the Learning Sciences, Northwestern University, 1890 Maple Avenue, Evanston, IL, 60201 USA; e-mail: forbus@ils.nwu.edu. Peter B. Whalley, Department of Engineering Science, Oxford University, Parks Road, Oxford, OX13PJ, UK; e-mail: whalley@vax.ox.au.uk

provide representation languages and reasoning techniques that can be applied to helping people make the transition from novice to expert reasoning about physical systems. Indeed, some of the earliest work in the field was directly aimed at instructional problems (e.g. Forbus & Stevens, 1981; Brown et al., 1982). Over the last decade there have been several important efforts aimed at using qualitative physics to help teach diagnosis, troubleshooting, and operation of complex physical systems (e.g. White & Frederiksen, 1990; Massey et al., 1988; Govindaraj, 1987), but little effort has been focused on using qualitative physics in classroom settings, to help undergraduates learn principles of a domain (one rare exception is: Roschelle, 1993).

In this paper we describe a system, called CyclePad, that has been built to help engineering undergraduates appreciate and therefore learn important principles of thermodynamics. CyclePad provides a conceptual CAD environment where students can design and analyze power plants, refrigerators, and other thermodynamic cycles. It relies on a synergistic combination of existing AI techniques: compositional modeling to represent and reason with modeling assumptions, qualitative representations to express the intuitive knowledge of physics needed to detect impossible designs, truth-maintenance to provide the basis for explanations, and constraint reasoning and propagation to provide efficient mathematical reasoning. It incorporates a substantial fraction of the knowledge in a typical engineering thermodynamics textbook (Whalley, 1992), and has been tested on over two dozen examples of problems involving steady-state, steady flow systems where numerical answers or single-parameter sensitivity analyses are required.

Section 2 describes the pedagogical problems that motivated the design of CyclePad, including a brief overview of what thermodynamic cycles are and how they work. Section 3 demonstrates CyclePad's operation from a user's perspective. How CyclePad works is the subject of Section 4. Section 5 describes some preliminary experiences we have had with students

using CyclePad, and Section 6 summarizes the lessons we have learned so far in building it. Section 7 outlines our plans for future work.

THE TASK: TEACHING THE DESIGN OF THERMODYNAMIC CYCLES

A thermodynamic cycle is a system within which a working fluid (or fluids) undergoes a series of transformations in order to process energy. Every power plant and every engine is a thermodynamic cycle. Refrigerators and heat pumps are also examples of thermodynamic cycles. Thermodynamic cycles play much the same role for engineering thermodynamics as electronic circuits do for electrical engineering: A small library of parts (in this case, compressors, turbines, pumps, heat exchangers, and so forth) are combined into networks, thus potentially generating an unlimited set of designs for any given problem. (Practically, cycles range from four components, in the simplest cases, to networks consisting of dozens of components.) One source of the complexity of cycle analysis stems from the complex nature of liquids and gases: Subtle interactions between their properties must be harnessed in order to improve designs. Cycle analysis answers questions such as the overall efficiency of a system, how much heat or work is consumed or produced, and what operating parameters (e.g., temperatures and pressures) are required of its components. An important activity in designing cycles (or indeed in many engineering design problems) is performing sensitivity analyses, to understand how choices for properties of the components and operating points of a cycle affect its global properties.

To illustrate, consider the sequence of power plant designs in Fig. 1. Figure 1(a) shows a simple Rankine cycle, which pumps a working fluid (as liquid water) into a boiler to produce steam. In the turbine the high-pressure steam expands, thus performing work. Heat is extracted from the steam in the condenser so that the working fluid is again water. Finally, this water is pumped into the boiler (which requires

work to be done in the pump) thus beginning the whole cycle again. These processes happen continuously in steady flow, in every part of the system. As it happens, this cycle is not very efficient, that is, a rather low proportion of the heat supplied in the boiler is converted into work in the turbine. Raising the temperature (and hence the pressure) of the boiler increases efficiency, but due to material properties of the components, there are upper bounds on operating pressures and temperatures. Figure 1(b) shows a more efficient design, which uses a second turbine to extract more energy from the steam. The purpose of the reheater is to ensure that the steam does not become very "wet", i.e., to begin to condense, because water droplets moving at high speed may damage the blades of the turbine. The extra energy required to reheat the steam is more than balanced by the additional work gained from the second turbine. One can do even better, however. Figure 1(c) shows a *regenerative feedwater cycle* where some of the steam from the outlet of the high-pressure turbine is routed back to the water feeding the boiler. The boiler then is adding heat to water which is starting at a higher temperature, and this increases the efficiency.

The analysis and design of thermodynamic cycles is the major task which drives engineering thermodynamics, aside from applications to chemistry. In thermodynamics education for engineers, cycle analysis and design generally appears towards the end of their first semester, or is even delayed to a second course, since understanding cycles requires a broad and deep understanding of the fundamentals of thermodynamics. Even the most introductory engineering thermodynamics textbooks tend to devote several chapters to cycle analysis, and in more advanced books the fraction devoted to cycles rises sharply. Indeed, some textbooks focus exclusively on cycle analysis (e.g., Haywood, 1985). Aside from their intrinsic interest, the conceptual design of thermodynamic cycles provides a highly motivating context for students to learn fundamental principles more deeply than they would otherwise.

A variety of problems arise when teaching students how to design and analyze thermodynamic cycles:¹

1. Students tend to get bogged down in the mechanics of solving equations and carrying out routine calculations. This leads

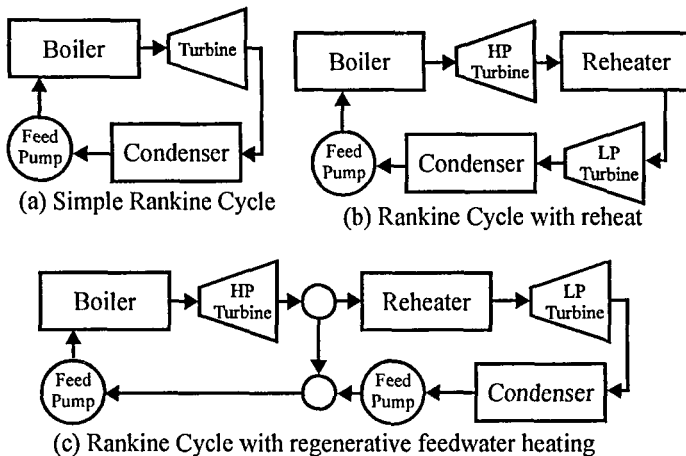


Figure 1. Sequence of Conceptual Designs for a Power Plant

them to avoid exploring multiple design alternatives and to avoid carrying out trade-off studies (e.g., seeing how efficiency varies as a function of turbine efficiency versus how it varies as a function of boiler outlet temperature). Yet without making such comparative studies, many opportunities for learning are lost.

2. Students often have trouble thinking about what modeling assumptions they need to make, such as assuming that a heater operates isobarically, leading them to get stuck when analyzing a design.
3. Students typically don't challenge their choices of parameters to see if their design is physically possible (e.g., that their design does not require a pump that produces rather than consumes work).

CyclePad was designed specifically to help students learn engineering thermodynamics by providing an intelligent learning environment that handles routine calculations, facilitates sensitivity analyses, helps students keep track of modeling assumptions, and detects physically impossible designs.

OVERVIEW OF CYCLEPAD

CyclePad can be viewed as a CAD system for the conceptual design of thermodynamic cycles, although it provides substantially more explana-

tion capabilities than existing CAD software. CyclePad performs steady state analyses of steady-flow thermodynamic cycles. The restriction to steady-state is standard for this kind of analysis, since issues of how to start up and shut down the plant, or how easy it will be to monitor, maintain, or troubleshoot are issues of concern only after the basic design has been shown to be sound with respect to the goals for it (e.g., amount of work produced, efficiency, etc.). The restriction to steady-flow systems means that CyclePad cannot currently be used to analyze internal combustion engines, such as Otto or Diesel cycles. Although we plan to extend CyclePad to analyze such systems, steady flow cycles constitute the majority of the cycle-related material taught to engineering students. (For example, in Whalley (1992) four out of five chapters on cycles concern steady flow cycles, in Haywood (1985) it is 9 out of 10 chapters, and El-Wakil (1984) focuses only on steady-flow systems.)

When a user starts up CyclePad, they find a menu of component types (e.g., turbine, compressor, pump, heater, cooler, heat exchanger, throttle, splitter, mixer) that can be used in their design. Table 1 illustrates the components available and the modeling assumptions that can be made about them. Components are connected together by *stuffs*, which represent the properties of the working fluid at that point in the system. (Stuffs serve the same role as nodes in electronic circuits.) The interface helps the user put together a design by highlighting what parts remain unconnected and providing simple cri-

Table 1. CyclePad components and their associated modeling assumptions

Turbine	adiabatic, isentropic
Compressor	adiabatic, isentropic, isothermal, polytropic, non-polytropic
Pump	isothermal
Heater	isobaric
Cooler	isobaric
Heat Exchanger	concurrent, countercurrent
Throttle	none
Splitter	saturated, non-saturated
Mixer	none

tiques of the structure. Once the structural description of the cycle is finished (e.g., there are no dangling connections or stuffs), CyclePad allows the user to enter an analysis mode, where the particular properties of the system, such as the choice of working fluid, the values of specific numerical parameters, and modeling assumptions can be entered. Figure 2 illustrates.

CyclePad accepts information incrementally, deriving from each user assumption as many consequences as it can. At any point questions can be asked, by clicking on a displayed item to obtain the set of questions (or commands) that make sense for it. In addition to numerical parameters and structural information, all modeling assumptions made about a component are displayed with it, and clicking on a component

shows the modeling assumptions that can legitimately be made about that component, given what is known about the system so far. The questions and answers are displayed in English, and include links back into the explanation system, thus providing an incrementally generated hypertext:

Q: Why is $Nu_{thermal}(CYCLE) = 46.00550\%$?

A: $Nu_{thermal}(CYCLE) = 46.00550\%$ was found by using
 $Nu_{thermal}(CYCLE) = \frac{\text{net work}(CYCLE)}{Q_{in}(CYCLE)}$
 on
 $Q_{in}(CYCLE) = 4069.555 \text{ kW}$
 $\text{net work}(CYCLE) = 1872.219 \text{ kW}$

$\text{net } Q(CYCLE) = 1,384 \text{ kW}$
 $\text{net work}(CYCLE) = 1,384 \text{ kW}$
 $Nu_{thermal}(CYCLE) = 42.82\%$
 $\text{work_ratio}(CYCLE) = 0.99$

TUR-1 works isentropically
 TUR-1 works adiabatically
 $\text{work}(TUR-1) = 1,395 \text{ kW}$
 $Q(TUR-1) = 0 \text{ kW}$
 $\text{mass_flow}(TUR-1) = 1.00 \text{ kg/sec}$

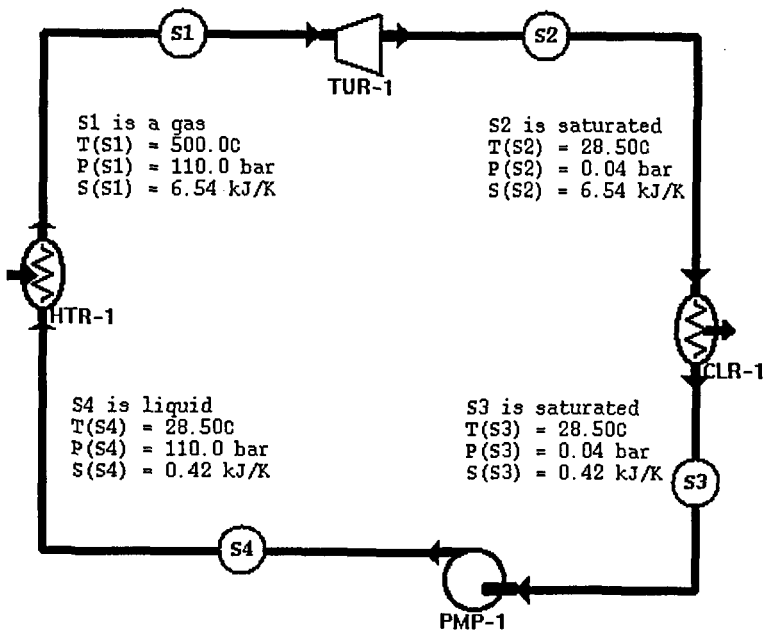


Figure 2. CyclePad Uses an Annotated Schematic Interface Model

- Q: Why is net work(CYCLE) = 1872.219 kW?
- A: net_work(CYCLE) = 1872.219 kW
was found by using
net_work(CYCLE) = work_in(CYCLE)+work out(CYCLE)
on
work in(CYCLE) = -15.68446 kW
work out(CYCLE) = 1887.903 kW
- Q: Why is work_out(CYCLE) = 1887.903 kW?
- A: work out(CYCLE) = 1887.903 kW
was found by using
work out (CYCLE) =
work(TURBINE2) +work(TURBINE1)
on
work(TURBINE2) = 1372.131 kW
work(TURBINE1) = 515.7723 kW
- Q: Why does work out(CYCLE) = work(TURBINE2)+work(TURBINE1) hold?
- A: work_out(CYCLE) =
work(TURBINE2)+work(TURBINE1)
because we assumed:
CYCLE is a thermodynamic cycle
the work-flows-out of CYCLE
consist of
{work(TURBINE2),
work(TURBINE1)}
- Q: How has TURBINE1 works isentropically been used?
- A: TURBINE1 works isentropically
was used to derive
TURBINE1 works adiabatically
 $T(S3) = T_{out_i}(TURBINE1)$
 $s(S2) = s(S3)$

In addition to numerical assumptions, selecting a component provides commands for making or retracting modeling assumptions concerning that component. For example, clicking on a new turbine yields a menu of commands which offers the options of assuming the turbine is adiabatic or isentropic. Such modeling assumptions can introduce new constraints which may help carry an analysis further and new parameters (e.g., the efficiency of the turbine) that must be set.

When CyclePad uncovers a contradiction, it changes the interface to provide tools to resolve the problem by presenting the source of the contradiction (e.g., an impossible fact becoming believed, or conflicting values for a numerical parameter) and the set of assumptions underlying that contradiction. The hypertext dialog facilities can be used with this display to figure out which assumption(s) are dubious and change them accordingly.

Sensitivity analyses are an important tool for building a student's intuitions about how physical principles contribute to the way an artifact works. CyclePad performs sensitivity analyses by asking a student to select a dependent parameter first. It then offers a choice of potentially relevant independent parameters, based on its analysis of the cycle. Given an independent and dependent parameter, CyclePad derives an equation (including table references, if needed) that describes the dependent parameter in terms of the independent parameter, and plots the results. Figure 3 illustrates.

We have tested CyclePad on over two dozen examples to date, ranging from simple ideal gas problems to the analysis of a combined gas turbine/steam Rankine cycle system. We believe that the current version of CyclePad can solve all of the problems in Whalley (1992) concerning steady-state analyses of steady-flow cycles that require numerical answers or sensitivity analyses involving a single parameter. (We are continuing to test it on new examples, drawn from other textbooks as well.) CyclePad is very efficient. The combined gas turbine/steam Rankine cycle is the most complex system in Whalley (1992), consisting of ten components. Good students take between 20 minutes and one hour to solve this problem. CyclePad does somewhat more work in analyzing this problem than a good student would, instantiating 219 equations involving 362 parameters, whereas a solution can be found using only 52 equations. However, CyclePad is still faster, taking just over two minutes on a workstation, versus just over ten minutes on a PowerBook 165c. We believe that the combination of the speed at which CyclePad carries out the routine calculations,




Figure 4. A CyclePad Sensitivity Analysis. This Graph Shows How Thermal Efficiency Varies for Different Choices of Outlet Pressure for the Turbine for the Cycle of Figure 2

its explanation facilities, and its consistency-checking facilities, will make it a valuable tool for students learning thermodynamics.

HOW CYCLEPAD WORKS

The overall structure of CyclePad was inspired in part by EL (Stallman & Sussman, 1977), an experimental system for DC analysis of analog electronic circuits. EL was one of the first systems to use constraint propagation and dependency networks to organize its reasoning, and introduced the idea of dependency-directed backtracking. In this section we see how Cy-

clePad exploits the advances made by the field since EL, by examining each of the AI ideas that contributes to CyclePad's operation, and the reasons for these particular design choices.

The Role of Compositional Modeling

As noted previously, steady-state analyses are required for the conceptual design of thermodynamic cycles. By restricting ourselves to steady-flow systems, it is also the case that the process structure (i.e., the collection of physical processes acting in each component) is fixed for all time. These restrictions allow us to organize the

domain theory around the components which comprise a cycle and the properties of the working fluid at particular locations (i.e., the connections between components).

The modeling language used in CyclePad is similar to other implementations of compositional modeling. For example, the following is part of CyclePad's model of a heater.

A sample of CyclePad's knowledge base

```
(defEntity (Abstract-tax ?self ?in
  ?out)
  (thermodynamic-stuff ?in)
  (thermodynamic-stuff ?out)
  (total-fluid-flow ?in ?out)
  (== (mass-flow ?in)
    (mass-flow ?out))
  (parameter (mass-flow ?self))
  (parameter (Q ?self))
  (parameter (spec-Q ?self))
  (heat-source (heat-source
    ?self))
  ((parts :cycle) has-member
    ?self)
  (?self part-names (in out))
  (?self IN ?in) (?in IN-OF ?self)
  (?self OUT ?out)
  (?out OUT-OF ?self))

(defAssumptionClass
  ((abstract-Hx ?hx ?in ?out))
  (isobaric ?hx)
  (:not (isobaric ?hx)))

(defEntity (Heater ?self ?in ?out)
  (abstract-Hx ?self ?in ?out)
  (?self instance-of heater)
  (heat-flow (heat-source ?self)
    (heat-source ?self)
    ?in ?out)
  ((heat-flows-in :cycle)
  has-member (Q ?self))
  ( (Q ?self) 0.0))

(defEquation Hx-law
  ((Abstract-Hx ?hx ?in ?out))
```

```
(:= (spec-h ?out)
  (+ (spec-h ?in)
    (spec-Q ?hx))))
```

```
(defEquation spec-Q-definition
  ((Abstract-Hx ?hx ?in ?out))
  (:= (spec-Q ?hx)
    (/ (Q ?hx) (mass-flow. ?hx))))
```

CyclePad's knowledge base currently consists of 29 conceptual entities, 5 physical processes, 9 assumption classes, 98 equations, 40 pattern-directed rules and 41 background facts about thermodynamics.

Modeling assumptions are organized into *assumption classes* (Addanki et al., 1989; Falkenhainer & Forbus, 1991). Assumption classes are always associated with particular classes of components. The relevance of one assumption class can depend on the particular choices made for another assumption class. For example, it only makes sense to consider whether a compressor is isentropic if it is already known (or assumed) to be adiabatic.

The Role of Constraint Reasoning and Propagation

A design is not finished until numerical values have been chosen for its parameters. This is one reason why the overwhelming majority of thermodynamics textbook problems require numerical answers.² This fact, plus the relative simplicity of the equations involved, has meant that constraint propagation has sufficed for CyclePad.

In compiling CyclePad's knowledge base, equations are automatically converted into antecedent constraint rules that propose values for the n th variable in an equation whenever the other $n-1$ variables are known. Redundant equations are introduced when needed to overcome simultaneities. This automatic translation simplifies development. Equations in their original form are still represented in the knowledge base, however, and are used in two ways. First, they are part of the dependency structure for any results calculated via constraint propagation,

for explanatory accuracy. Second, they can be inspected via the query system, so that students can find out what equations mention a specific parameter, and what equations might be used to calculate a desired value.

Property tables comprise a critical source of information for CyclePad. Property tables are woven into the constraint propagator via pattern-directed rules, operating under the same protocol as the rules compiled for equations. Due to the inherent loss of accuracy in interpolation, it is important, unlike equations, to *avoid* using tables in every logically possible fashion. Given a superheated vapor, for instance, knowing the pressure and temperature suffice to determine everything else (e.g., the specific enthalpy, specific entropy, etc.). If one redundantly computes from, say the specific enthalpy and specific entropy what the pressure and temperature will be, it is very likely that the newly estimated values will trigger a contradiction, given the accumulated inaccuracies in the interpolation process. Consequently, an important design choice in implementing tables is selecting which directions of access are likely to prove most productive for the kinds of analyses being made.

The Role of Qualitative Physics

In CyclePad qualitative physics provides the medium for representing constraints on what is physically possible. Occurrences of physical processes inside components are explicitly represented. Each process occurrence includes ordinal constraints that are tested against numerical values by CyclePad's constraint propagation mechanism. Below is shown CyclePad's current knowledge about physical processes.

Physical processes provide qualitative constraints

```
(defProcessEpisode
  (fluid-flow ?in ?out)
  (same-substance ?in ?out))

(defProcessEpisode
```

```
(total-fluid-flow ?in ?out)
(fluid-flow ?in ?out)
(= (mass-flow ?in)
   (mass-flow ?out)))

(defProcessEpisode
  (heat-flow ?src-start ?src-end
   ?dst-start ?dst-end)
  ( (T ?src-start) (T ?dst-start))
  (:not ( (T ?src-start)
          (T ?dst-end)))
  (:not ( (T ?dst-end)
          (T ?src-end))))

(defProcessEpisode
  (compression ?in ?out ?worker)
  ( (P ?out) (P ?in))
  ( (spec-shaft-work ?worker) 0))

(defProcessEpisode
  (expansion ?in ?out ?receiver)
  ( (P ?out) (P ?in))
  ;; Does work
  ( (spec-shaft-work ?receiver)
    0))
  ;;For (?proc ?in ?out)
  ;; a process occurring
  ;; in component ?c and
  ;;process assumption ?I,
  (?I ?c) => (?I (?proc ?in
    ?out))
  (isobaric (?proc ?in ?out))
  => (P ?in) = (P ?out)
  (isothermal (?proc ?in ?out))
  => {T ?in} = {T ?out}
  (isentropic (?proc ?in ?out))
  => (S ?in) = (S ?out)
  (isentropic ?c) => (adiabatic ?c)
  {adiabatic ?c} => (Q ?c) = 0.0
```

The Role of Truth Maintenance

We used an LTMS (McAllester, 1980) in CyclePad because it offered the best tradeoff between inferential power and economy. (In fact, CyclePad's inference engine is a customized version of the LTRE system from (Fortus & de Kleer, 1993)). We ruled out a JTMS because

Horn clauses are too clumsy for many of CyclePad's inferential needs, including biconditionals (used in definitional consequences of modeling assumptions, e.g., a compressor is operating isentropically exactly when its isentropic efficiency is 1.0) and TAXONOMY constraints (Hayes, 1985) (used in implementing assumption classes). The ability of an ATMS to provide rapid switching between very different contexts was not required: While frequent additions and retractions of assumptions are made in carrying out an analysis, typically these changes are a small fraction of the working set of assumptions in force.

A critical role for the LTMS dependency network is as an input for explanation generation. Explanations in CyclePad are in terms of *structured explanations*, an abstract layer between the reasoning system and the interface that casts the consequences of the inference system in terms relevant to the user. This includes summarization (e.g. Gruber & Gautier, 1993), as in removing any reference to implementation-dependent information such as the constraint propagation mechanism from an argument. It also includes making explicit implicit dependencies, such as the variables whose values must be known before the constraint rule implementing a particular equation will fire, when explaining what assumptions might lead to more progress.

PRELIMINARY RESULTS WITH STUDENT USERS

For formative evaluation purposes, we have tried out CyclePad with engineering undergraduates at Oxford and Northwestern University. The first students, at Oxford, attempted to use the forms and table-based interface that we used for system development purposes. It was clear from their responses that a much better interface was required. We developed a new graphical interface (as illustrated above), based roughly on the STEAMER icon model (Stevens et al., 1981; Hollan et al., 1984). STEAMER's schematics included distinct icons to provide "gauges" that could be tapped into the underly-

ing simulation. This metaphor was fine for teaching principles of operation but is not quite right for design. The metaphor we use instead is that of an *annotated schematic*, where dynamic displays of property values and modeling assumptions are associated with each component and stuff, updated whenever the assumptions underlying the analysis change. The properties displayed for each component (and stuff) can be changed according to the student's current information needs, either by component class or on a component-by-component basis.

CyclePad with this new interface was then made available to Northwestern students taking the advanced thermodynamics course in the Mechanical Engineering Department, taught by Prof. Siavash Sohrab. Students were taught how to use CyclePad in a special laboratory session, using X-terminals to access it running on an RS/6000 workstation in our laboratory. Roughly one half of the students taking the course attended the CyclePad session, and we gathered information on their use of the system on problems of their own choosing. We learned several things from this experience. First, the students were able to use the system to help them with homework problems. Second, the explanation system is currently too local—students would get lost working from conclusions back to assumptions, and also in working backwards from goal quantities to properties that might be assumed to derive them.

Based on these experiences, we redesigned the interface yet again, and ported the system to Windows, so that we could easily distribute the software on machines that are widely available in college laboratories and owned by students. We are now also working with instructors at the US Naval Academy who are introducing CyclePad in their thermodynamics classes, which provides additional perspectives in teaching thermodynamics and to an even larger pool of students for our educational experiments. Advanced thermodynamics courses already tend to take a design focus, so introducing CyclePad in that context has been relatively straightforward compared to introductory courses, where the use of design tasks is more novel. Currently,

in collaboration with our instructor-colleagues, we are designing new homework exercises for introductory thermodynamics courses that exploit CyclePad's capabilities to give students experience with design.

LESSONS LEARNED SO FAR FROM DEVELOPING CYCLEPAD

CyclePad represents one of the first attempts to use recent ideas developed by the qualitative physics community in an intelligent learning environment. While CyclePad is still in the process of serious field-testing, we believe that we have already learned several generally useful lessons in building it.

Compositional Modeling Scales Up

Previous uses of compositional modeling have either focused on large but purely qualitative domain theories, or small quantitative theories. CyclePad demonstrates that the ideas of compositional modeling can be used to organize a substantial body of quantitative and qualitative knowledge so that it can be used effectively.

Automatic model formulation, which typically has been the focus of previous compositional modeling work, is less relevant for this application. Nevertheless, the mechanisms of assumption classes and logical constraints between modeling assumptions provide a valuable service in helping the user organize an analysis. In fact, one of the skills being taught in using CyclePad is model formulation. A boiler, for instance, is typically approximated as a heater for the purposes of cycle analysis. A flash chamber is modeled as a splitter whose working fluid is saturated and with particular assumptions about the dryness of the outlets. A multi-stage turbine is modeled as a sequence of turbines and splitters. CyclePad helps users analyze models, so they can figure out if their choice of idealization makes sense, but currently CyclePad does not provide direct assistance with formulating an idealized model from an informal specification.

Regarding Constraint Reasoning

In this task numerical constraint propagation suffices. There are however natural extensions of CyclePad's analytic abilities for which algebraic manipulation would be useful. For instance, some insights about how a cycle works are best captured via equations.³ We plan to extend CyclePad to derive such equations on demand. Our experience with the constraint rules compiler in CyclePad, and other work on thermodynamics problem solving (Skørstad & Forbus, 1989), suggests that relatively simple algebraic capabilities will suffice for this extension.

We draw two additional conclusions regarding constraint manipulation. First, the commercial world has developed many powerful symbolic algebra packages, such as Matlab, Mathematica, Maple, and Macsyma, which in some cases are excellent off-the-shelf solutions to particular problems. However, we suspect that many educational applications will be like CyclePad: Simple algebraic facilities are all that is required, and thus the complexity (and expense) of integrating commercial symbolic algebra packages can be avoided. Second, we found that special-purpose constraint languages (e.g., Steele & Sussman, 1980) were too restrictive for our purposes. Given the need to reason about modeling assumptions and the need to integrate information from property tables, it was much easier to implement a simple constraint propagator inside a pattern-directed inference system than it was to interface a special-purpose constraint manipulator. Aside from applications where scaling up to extremely large system descriptions (e.g., VLSI CAD) is a key requirement, it is hard to see any situation where using such languages makes sense.

Regarding Qualitative Physics

The combination of steady-state analysis, the restriction to steady-flow systems, and the use of idealized components dramatically simplified the representation of physical processes, since

the occurrence of particular physical processes could simply be stipulated inside a component.

CyclePad's focus on quantitative analysis also means that the major inferential role for qualitative physics is ruling out physically impossible designs. We believe similar simplifications will hold in many other applications, since well-designed artifacts explicitly represent the important physical changes in terms of the kinds of components and connections that comprise a schematic, and many science and engineering educational applications involve quantitative knowledge heavily.

On the other hand, certain extensions to CyclePad's capabilities will require substantially more qualitative representations and reasoning. For instance, CyclePad currently does not try to explain how components work, nor does it provide assistance for understanding the physical rationale underlying design changes. To formalize such arguments will take richer qualitative representations, as well as the ability to reason with property diagrams (e.g., Pisan, 1994) and teleology (e.g., Everett, 1995). Fortunately, the automatic instantiation of physical process descriptions from a domain theory in a restricted ontology is an inexpensive and well-understood operation.

Regarding Explanation Generation and TMSs

The use of a structured explanation system as an abstraction layer between interface and reasoning system was extremely helpful in developing CyclePad, since it allowed us to optimize each independently. We also found, as suggested by Reiter & Mellish (1993), that sophisticated natural language generation techniques were inappropriate for this task. The ability to automatically generate hypertext in response to a user's questions obviates the need for discourse planning, and the fixed nature of the task means that issues such as selecting the appropriate level of detail in an explanation can be postponed. Hypertext allows users to select how much they want to know about a topic, and since the hypertext is only generated on demand,

many navigation problems common in static hypertexts are avoided.

ATMS technology (de Kleer, 1986) has been widely used in qualitative reasoning systems because of its ability to rapidly switch between alternate interpretations. As noted previously, this ability is unnecessary in CyclePad, and we suspect that this will be true for most educational applications. The major surprise we had was that we had to introduce garbage collection into the TMS. Students make and retract many assumptions during an analysis, most of which are never reused, so that the monotonic growth in memory usage of a traditional TMS is a bad bargain. Instead, we designate a subset of facts (i.e., statements about numerical values) as statements which should be erased when no longer believed. The algorithms to do this correctly are slightly subtle (see Everett & Forbus, 1996 for details), but they allow us to keep the advantages of a TMS for explanation generation while avoiding memory problems.

DISCUSSION

CyclePad demonstrates that qualitative physics has advanced enough to support new applications of AI to educational problems. Compositional modeling provides representational tools and techniques that can be used to encode a substantial body of knowledge about engineering thermodynamics, with constraint propagation providing analytic capabilities and qualitative representations providing the intuition needed to detect student blunders. Automatically generated hypertext explanations enable the user to explore the consequences of his or her assumptions, and figure out what modeling assumptions are needed to make further progress.

The preliminary results from using CyclePad with students have been sufficiently encouraging that we are carrying out larger-scale experiments with thermodynamics classes at Oxford University, Northwestern University, and the US Naval Academy. The extensions described below are motivated in part by what we think is

necessary to make the systems like CyclePad an integral part of thermodynamics education.

There are several straightforward extensions to CyclePad that are in progress. First, we are extending it to handle non-steady flow cycles, such as Otto and Diesel cycles. Second, we will add some algebraic capabilities, so that CyclePad can help students derive algebraic expressions that capture important tradeoffs in specific systems.

We view CyclePad as part of a *virtual laboratory* for exploring thermodynamic cycles. A virtual laboratory is a software environment consisting of a set of parts, corresponding to physical parts or important abstractions in the domains of interest, tools for assembling collections of these parts into designs, and facilities for analyzing and testing designs. By working in this software environment, students can "build" their designs and try them out without expense or danger. In simpler domains some commercial software exists that can be viewed as virtual laboratories (e.g., Interactive Physics for simple dynamics and Electronics Workbench). A novel contribution of qualitative physics is the ability to generate explanations. For educational applications, explanation generation is vital, to help students see what aspects of a situation are important and to tie what they are observing back to fundamental principles. One of our next steps is to extend CyclePad's explanation facilities, by adding *coaches* (Burton & Brown, 1982; Lesgold et al., 1992; Schank & Nohan, 1991) to help students, both to guide them through the analysis process (including the representation of real devices in terms of ideal components) and to suggest improvements to a student's design. We believe that software environments that give students more design experiences, especially conceptual design where the ratio of principles to details is very high, can help them gain a deeper appreciation of science and learn to be better engineers.

Acknowledgments: The initial development of CyclePad was supported by a grant from the Science and Engineering Research Council in the UK and by grants from the Office of Naval Research and NASA

in the U.S. Further development of CyclePad is being supported in the US by the National Science Foundation. The Office of Naval Research is supporting the collaboration with the US Naval Academy. We thank Prof. Siavash Sohrab allowing us to set up CyclePad laboratory sessions for his thermodynamics class, Yusuf Pisan for many enlightening bug reports, suggestions, and for running the CyclePad laboratory sessions, Will Wright for interface ideas, and John Everett and Andy Bachmann for developing the new CyclePad interface.

NOTES

1. These observations are based on the experience of the second author, who teaches engineering thermodynamics to undergraduates.
2. In a typical textbook we surveyed, 90% of the exercises required numerical answers.
3. For example, figuring out that for a gas turbine cycle the maximum specific work output is achieved when the pressure ratio is the square of its maximum possible value (Whalley, 1992).

REFERENCES

- Addanki, S., Cremonini, R., & Penberthy, J.S. (1989). Reasoning about assumptions in graphs of models. *Proceedings of IJCAI-89*.
- Brown, J. S., Burton, R., & de Kleer, J. (1982). Pedagogical, natural language, and knowledge engineering techniques in SOPHIE I, II, and III. (pp. 227-282). In Sleeman, D. and Brown, J.S. (Eds.), *Intelligent Tutoring Systems*, Academic Press.
- Burton, R., & Brown, J. S. (1982). An investigation of computer coaching for informal learning activities. In Sleeman, D. and Brown, J.S. (Eds.), *Intelligent Tutoring Systems*, Academic Press, 79-98.
- de Kleer, J. (1986). An assumption-based truth maintenance system. *Artificial Intelligence*, 28: 127-162.
- El-Wakil, M. (1984). *Powerplant Technology*, McGraw-Hill.
- Everett, J. (1995). A Theory of Mapping from Structure to Function Applied to Engineering Thermodynamics. *Proceedings of IJCAI-95*.
- Everett, J., & Forbus, K. (1996) Scaling up Logic-Based Truth Maintenance Systems via Fact Garbage Collection. To appear in *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96)*.

- Falkenhainer, B., & Forbus, K. (1988). Setting up large-scale qualitative models. *Proceedings of AAAI-88*.
- Falkenhainer, B., & Forbus, K. (1991). Compositional Modeling: Finding the Right Model for the Job. *Artificial Intelligence*, 51, 95-143.
- Forbus, K., & de Kleer, J. (1993). *Building Problem Solvers*, MIT Press.
- Forbus, K., & Stevens, A. (1981). Using Qualitative Simulation to Generate Explanations. *Proceedings of the Cognitive Science Society*.
- Govindaraj, T. (1987). Qualitative approximation methodology for modeling and simulation of large dynamic systems: Applications to a marine steam power plant. *IEEE transactions on systems, man, and cybernetics*, vol SMC-17, no. 6.
- Gruber, T., & Gautier, P. (1993). Machine-generated explanations of engineering models: A compositional modeling approach. *Proceedings of IJCAI-93*.
- Hayes, P. (1985). Naive Physics 1: Ontology for Liquids. In Hobbs, J. and Moore, R. (Eds.) *Formal Theories of the Commonsense World*, Ablex, Norwood, NJ.
- Haywood, R. W. (1985). *Analysis of Engineering Cycles: Power, Refrigerating and Gas liquefaction Plant*, Pergamon Press.
- Hollan, J., Hutchins, E., & Weitzman, L. (1984). STEAMER: An interactive inspectable simulation-based training system. *AI Magazine*, 5(2), 15-27.
- Lesgold, A., Eggan, G., Katz, S., & Rao, G. (1992). Possibilities for Assessment Using Computer-Based Apprenticeship Environments. In Regian, J., & Shute, V. *Cognitive Approaches to Automated Instruction*. Erlbaum.
- Massey, L., de Bruin, J., & Roberts, B. (1988). A Training System for System Maintenance. In Psofka, J. Massey, L., and Mutter, S. *Intelligent Tutoring Systems: Lessons Learned*. Erlbaum.
- McAllester, D. (1980). *An outlook on truth maintenance*. MIT AI Lab memo AIM-551.
- Nayak, P. (1992). *Automated modeling of physical systems*. Ph.D. dissertation, Computer Science Department, Stanford University.
- Pisan, Y. (1994). *Visual reasoning about physical properties via graphs*, submitted for publication.
- Reiter, E., & Mellish, C. (1993). Optimizing the costs and benefits of natural language generation, *Proceedings of IJCAI-93*.
- Roschelle, J. (1993). Collaborative Conceptual Change: Jointly acting social and cognitive processes. *Proceedings of CogSci-93*.
- Schank, R. C., & Nohan, M.Y. (1991). Empowering the student: New Perspectives on the Design of Teaching Systems. *The Journal of the Learning Sciences*, 1(7-35).
- Skorstad, G., & Forbus, K. (1989). Qualitative and quantitative reasoning about thermodynamics, *Proceedings of the Cognitive Science Society*.
- Stallman, R. M., & Sussman, G. J. (1977). Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis, *Artificial Intelligence*, 9, 135-196.
- Steele, G., & Sussman, G. J. (1980). CONSTRAINTS: A language for expressing almost-hierarchical descriptions, *Artificial Intelligence*, 14:1-39.
- Stevens, A., Roberts, B., Stead, L. Forbus, K., Steinberg, C., & Smith, B. (1981). *STEAMER: Advanced computer-aided instruction in propulsion engineering*, BBN Technical report No. 4702.
- Whalley, P. (1992). *Basic Engineering Thermodynamics*, Oxford University Press.
- White, B., & Frederiksen, J. (1990). Causal model progressions as a foundation for intelligent learning environments. *Artificial Intelligence*, 42, 99-157.