Improving dynamic decision making in complex systems through ecological interface design: A research overview

Kim J. Vicente

Research on dynamic decision making in complex systems has paid little attention to the impact of interface design on human performance. Ecological interface design (EID) is a theoretical framework for designing computer interfaces for complex human–machine systems that addresses this issue. This article provides an overview of a research program on EID conducted in the author's laboratory. A detailed example showing how the principles of EID can be applied to design an interface for a simplified but representative thermal-hydraulic process-control simulation is presented. Also, the results from laboratory research, lessons learned from an industrial prototype, and the details of technology transfer to industry are reviewed. Collectively, the findings from this research program demonstrate that dynamic decision-making performance in complex systems can be significantly improved through appropriate interface

Traditionally, decision-making research has focused on understanding and evaluating human performance under discrete, static conditions (e.g., Kahneman, Slovic and Tversky 1982). This research has been criticized on the grounds of limited representativeness, and therefore of limited generalizability as well (Hogarth 1981). Because dynamics, and in particular feedback, have not been represented, traditional decision-making research is of limited value for many tasks that people face on a daily basis. A familiar example from system dynamics research is the case of managers of large organizations, who have to make decisions in the face of a turbulent, dynamic environment. A perhaps less familiar example, which has motivated the research in this article, is the case of control-room operators of complex human–machine systems (e.g. nuclear power plants), who are required to engage in real-time, interactive, dynamic decision-making. In either case, understanding the factors that affect decision making is very important, since the consequences of poor performance are potentially enormous.

Fortunately, an increasing number of decision-making researchers have been investigating dynamic decision making in complex systems (for reviews see Brehmer 1992; Sterman 1994). The vast majority of these studies paint a very pessimistic picture of people's capabilities for dynamic decision making in complex systems. However, this body of work has under-emphasized the role that the computer interface can have on human performance. This is also an issue that is of great practical importance. In management contexts, companies are in the process of moving from text-based to more flexible and powerful graphics-based user interfaces. In industrial contexts, control-room designers are moving from traditional (e.g. analog, hard-wired) displays, or text-based computer interfaces, to graphical computer technology. If either of these trends is to result in improved performance and safety, or at least to maintain current levels, then we need a deep understanding of the impact such technology can have on human performance. The thesis of this article is that dynamic decision-making performance in complex systems can be improved through appropriate interface design.

Ecological interface design (EID) is a theoretical framework for interface design for complex human–machine systems. The goal of EID is to allow operators to take advantage of their powerful perception and action capabilities,

This research was sponsored by ABB Corporate Research—Heidelberg, the Natural Sciences and Engineering Research Council of Canada, AECL Research, and the Japan Atomic Energy Research Institute. The contributions of Klaus Christoffersen, Nick Dinadis, Chris Hunter, Bill Pawlak, and Alex Pereklita to this research are gratefully acknowledged. Also, thanks to John Sterman, Andy Ford, and an anonymous reviewer for their comments on an earlier draft.

System Dynamics Review Vol. 12, no. 4, (Winter 1996): 251–279
Received November 1995
Accepted June 1996

design. This conclusion has significant implications for system dynamics research.

Kim J. Vicente is assistant professor of Mechanical & Industrial Engineering and of Biomedical Engineering at the University of Toronto, and director of the Cognitive Engineering Laboratory. He is interested in the design of interfaces for complex human–machine systems, the study of expertise, and more broadly, in the design and analysis of complex work environments. *Address*: Department of Mechanical & Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Ontario, Canada M5S 3G8. *Email*: benfica @mie.utoronto.ca

while simultaneously providing the support required for problem-solving activities (Vicente and Rasmussen 1990; 1992). This article provides an overview of research on EID undertaken in our laboratory. Although this research was conducted in the context of process-control systems, its findings have important implications for system dynamics research investigating dynamic decision making in management flight simulators (MFSs).

Ecological interface design

EID is based on two seminal concepts from cognitive engineering research, the abstraction hierarchy (AH) and the skills, rules, knowledge (SRK) framework (Rasmussen 1986). The AH is a multilevel knowledge representation framework that can be used to develop physical and functional system models, as well as the mappings between them. It is used in EID to identify the information content and structure of the interface. The SRK framework defines three qualitatively different ways in which people can process information: at the skill-based level, behavior is governed by a dynamic world model that allows people to engage in fluid perceptual-motor control; at the rule-based level, behavior is governed by rules which directly map perceptual cues in the environment to appropriate actions, without any mediating processing; finally, at the knowledge-based level, behavior is governed by a symbolic mental model which allows people to engage in analytical problem solving. The SRK framework is used in EID to identify how information should be displayed in an interface. The idea is to take advantage of operators' powerful pattern recognition and psychomotor abilities. As Sterman (1994, p. 324) has pointed out, "High survival value over millions of years caused excellent motor skills to evolve, while the cognitive skills required to understand the assumptions behind a spreadsheet or the dynamic complexity of a market have led to reproductive advantage only recently, if at all." EID takes this idea one step further by suggesting that even performance in complex dynamic systems, which may seem to demand abstract cognitive processing, may be improved if problems are represented in such a manner that allows people to deploy the powerful skills they have developed through evolution. Thus, EID recommends that information be presented in such a way as to promote skill- and rule-based behavior, allowing operators to deal with task demands in a relatively efficient and reliable manner. Knowledge-based behavior is also supported by embedding an AH representation of the work domain in the interface. This provides operators with an external visualization of system structure and dynamics which offers support during novel situations requiring adaptive problem solving.

The EID framework consists of three principles. Each is intended to support a given level of cognitive control, as follows:

1. Knowledge-based behavior—represent the work domain in the form of an AH to serve as an externalized mental model that will support analytical problem solving.
2. Rule-based behavior—provide a consistent one-to-one mapping between the work domain constraints and the cues provided by the interface.
3. Skill-based behavior—support interaction via time–space signals; the operator should be able to act directly on the display.

For a detailed justification of these principles, see Vicente and Rasmussen (1990; 1992).

Because the AH plays such an important role in EID, it is important to describe it in more detail. The AH is a multilevel representation format that describes the various layers of constraint in a work domain. Each level represents a different model of the system. For many complex human–machine systems, five levels of constraint have been found to be of use:

- the purposes for which the system was designed (functional purpose);
- the intended causal structure of the process in terms of mass, energy, information, or value flows (abstract function);
- the basic functions that the system is designed to achieve (generalized function);
- the characteristics of the components and the connections between them (physical function);
- the appearance and spatial location of those components (physical form).

Higher levels represent *functional* information about system purposes, whereas lower levels represent *physical* information about how those purposes are realized by system components.

There are two advantages to adopting the AH as a basis for interface design (Vicente and Rasmussen 1992; Bisantz and Vicente 1994). First, this approach allows one to identify, *a priori*, the information needed to cope with events that are unfamiliar to operators and that have not been anticipated by designers. This is a very important property since such unanticipated events pose the greatest threat to system safety in industrial systems. Whereas traditional approaches to interface design rarely make an attempt to deal with this problem, the AH was explicitly designed to deal with unanticipated events. Second, the AH is also a psychologically relevant problem representation. There is a significant body of empirical research from diverse domains showing that problem-solving protocols can be mapped onto an AH representation (for reviews see Rasmussen 1986; Vicente and Rasmussen 1992). Several of these studies were conducted in domains that are more similar to those studied in system dynamics, such as the decisions made by designers and managers in large-scale engineering design projects in industry (e.g., Burns and Vicente 1995). Thus, in addition to satisfying the engineering requirement of containing the information needed to cope with unanticipated events, the AH also satisfies the psychological requirement of providing a representation that is consistent with human problem-solving processes. As a result of these unique characteristics, EID is a candidate for designing graphic, computer interfaces for complex human–machine systems where unanticipated events can occur.

Literature review

Starting point

EID was originally developed in the context of process control systems, such as nuclear power plants. A review of previous research and development efforts revealed that several researchers had developed interface designs for power-plant control rooms that were based on an AH representation of the plant (Vicente 1992a; 1992b). Unfortunately, up until the time that the review was conducted (1990), no experiment had been conducted comparing a multi-level interface based on an AH representation with any other type of interface. As a result, the direct empirical evidence for the EID framework was very weak, not because the results obtained were equivocal, but because appropriate studies simply had not been conducted. This provided a starting point for our systematic research program on EID.

Research vehicle

Most of our research on EID has been conducted with two related systems, DURESS and DURESS II, both of which are thermal–hydraulic process-control microworlds. The original DURESS system was non-interactive, presenting subjects with real-time, "canned" scenarios. Subjects were able to view the behavior of the system but not to control it. DURESS II, in addition to some minor structural changes, differs from the original primarily in that it allows for real-time interaction so that subjects may actively control the system.

DURESS II is illustrated in Figure 1, and consists of two redundant feedwater streams (FWSs) that can be configured to supply water to either of the two reservoirs. Each reservoir has associated with it an externally determined demand for water that can change over time. The system purposes are twofold: to keep each of the reservoirs at a prescribed temperature (40°C and 20°C), and to satisfy the current mass (water) output demand rates. To accomplish these goals, the subject has control over eight valves (VA, VA1,

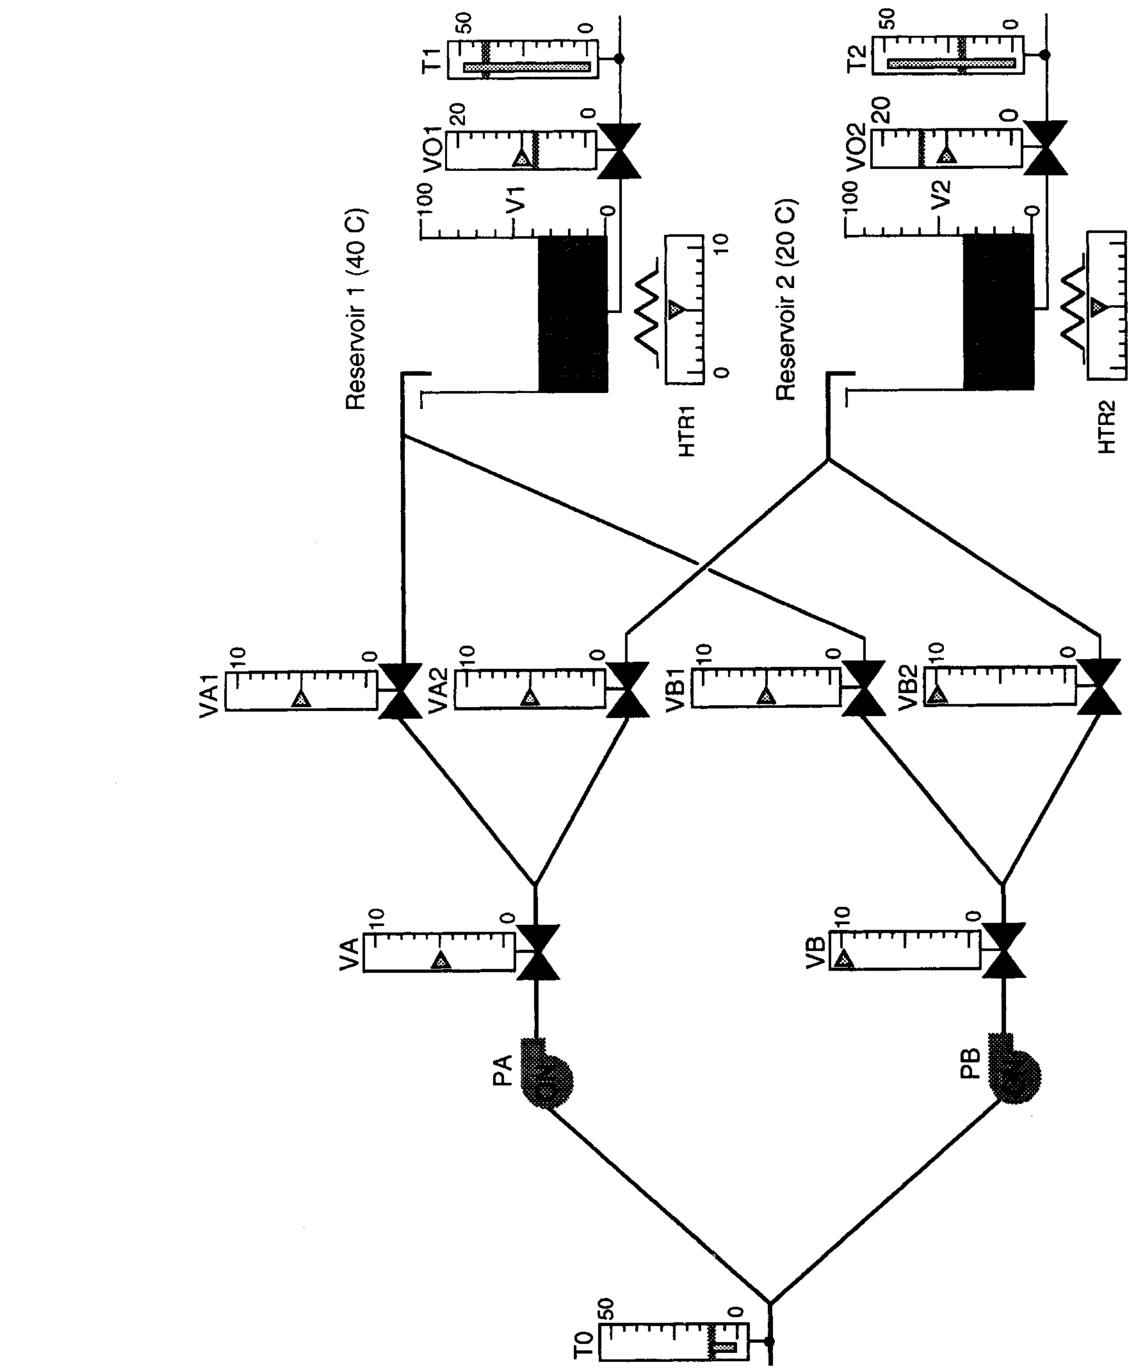

Fig. 1. *P* interface for DURESS II (from Pawlak and Vicente 1996). See text for description

VA2, VO1, VB, VB1, VB2, and VO2), two pumps (PA, PB), and two heaters (HTR1, HTR2). The time constants for the heaters and the remaining components are 15 s and 5 s, respectively.

P interface

Two different interfaces for DURESS II have been designed and compared in several experiments. The first, illustrated in Figure 1, provides a physical (*P*) representation of the system. It displays only the state of the physical components and the goal variables. The first meter on the extreme left of the display is a thermometer (T0), which measures the temperature of the water entering the system. This thermometer displays temperature as a vertical bar that increases in height as the water temperature increases. The normal inlet water temperature is 10°C, as indicated by the thin area on the T0 scale. After the thermometer, the input water stream splits and flows to two pumps (PA, PB). The maximum flowrate through each pump is 10 units/s. The pumps operate as discrete switches and are either on or off. The subject uses a mouse to click on the pump to change its state. The pumps are displayed in black (with white lettering) if they are off, and in light gray (with black lettering) if they are on. If either pump is turned on without any of the downstream valves being opened, the pump will fail after approximately 5 s.

The next set of components is the primary valves (VA, VB), which have a continuous range from 0 to 10. The valve state is set using a mouse to either drag the yellow triangular pointer to the desired setting, or to simply click on the scale at the desired point. From these primary valves, each FWS splits into two secondary valves connecting each stream to both reservoirs. The secondary valves (VA1, VA2, VB1, and VB2) operate in the same manner as the primary valves. The water then flows to the two reservoirs, where it is heated and removed, through the use of the heaters (HTR1, HTR2) and the output valves (VO1, VO2), to meet the temperature and demand goals, respectively. The reservoirs have a maximum capacity of 100 units. Reservoir levels are indicated by a scale on the side of each reservoir and by the shaded area depicting water in the reservoir. When reservoir volume exceeds the maximum capacity of 100 units, the system fails.

The heaters (HTR1, HTR2) also have a continuous range of 0 to 10. The subject can slide the triangular pointer on the heater scale to the desired setpoint, or click on the scale itself at the desired point. If there is continued heat transfer to an empty reservoir, the system will eventually fail. The water temperature in the reservoirs is displayed with thermometers (T1, T2). The goal temperature is represented as a thin area on the temperature scale. There is a tolerance of ±2°C from the setpoints (40°C and 20°C for Reservoirs 1 and

2). If the water in the reservoir boils, the system fails. Finally, the subject also controls the outlet valves (VO1, VO2) used to meet the demand goals. These operate in the same manner as the other valves, except that their maximum setting is 20. The demand for each reservoir is indicated by a thin area on the outlet valve setting scale. This goal area, ±1 unit around the desired level, moves as a function of changes in the demand. The *P* interface was developed because it is typical of how existing computer interfaces for process control systems have been designed. Thus, it serves as a meaningful baseline.

P + F interface

The second interface for DURESS II, based on the three principles of EID, will be referred to as the *P + F* interface since it provides both physical and functional representations of the system.

SUPPORTING KNOWLEDGE-BASED BEHAVIOR. The first step in designing an ecological interface is to represent the work domain in the form of an AH to identify the content and structure of the interface. Such representations provide an externalized mental model that can support effective problem solving activities. Figure 2 illustrates the AH for one reservoir. It should be noted that the system's dynamic equations, not shown, provide a formal description of the qualitative structure represented in the figure. At the top level of *Functional Purpose*, the higher-level objectives of the domain are specified: to keep the water in each temperature at criterion, and to keep enough water in each reservoir to satisfy the respective output demand. While not explicitly represented, there are also constraints on how these objectives are to be met (e.g., avoid spilling or boiling water, do not damage any components, etc.).

At the second level, *Abstract Function*, the structure of the system is described according to first principles, conservation of mass and energy. For the mass balance, there are three variables: the rate at which mass is flowing into the reservoir (source), the rate at which mass is flowing out of the reservoir (sink), and the current mass inventory. The energy topology is similar, the only difference being that there are two energy sources, the incoming water and the heater.

At the level of *Generalized Function*, DURESS II is described in terms of several standard heating and liquid-flow functions that instantiate the mass and energy topology specified above. In the energy flow, there is a heating source, a cooling source, a means for removing heat, and a means for storing heat. In the mass flow, there is a means for adding water, a means for removing water, and a means for storing water. At this level, we see that the energy and mass flows are actually coupled: the water source is also a cooling source;

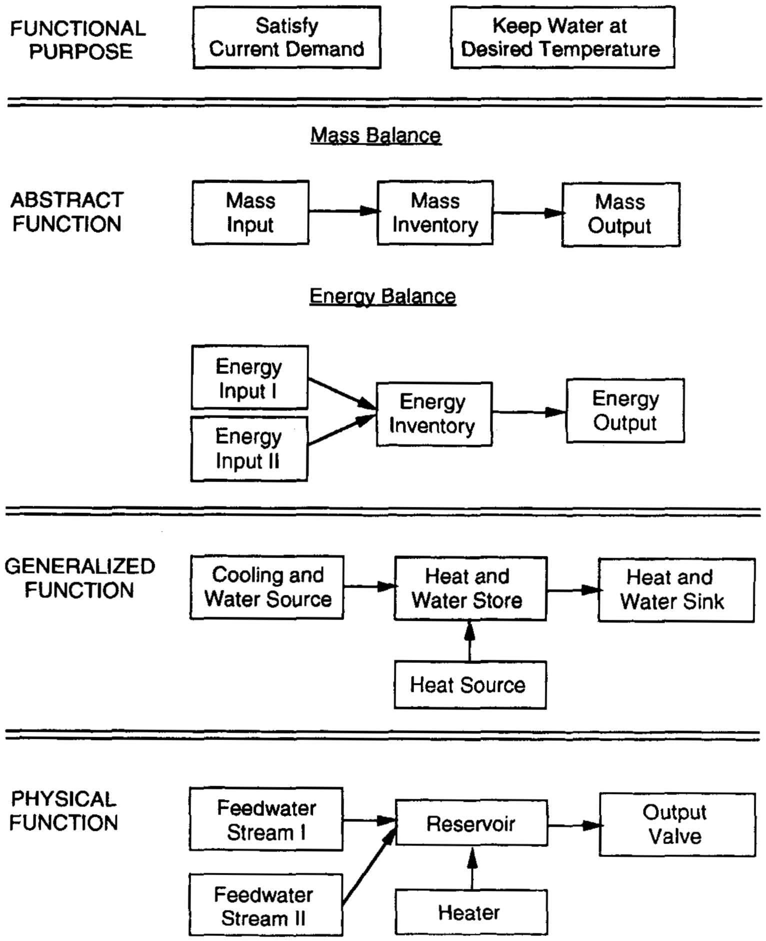

Fig. 2. AH representation of DURESS II showing the relationships within levels of the hierarchy (adapted from Vicente and Rasmussen 1990)

the means for storing water is also a means of storing energy; and the means of removing water is also a means of removing energy.

At the level of *Physical Function*, the system is described in terms of the components that realize the functions specified at the level above. These are the variables over which the operator has control. Because of space restrictions, the FWSs are represented as a whole in Figure 2, but a more detailed representation would decompose each stream into its constituent components (i.e., three valves and one pump) along a part–whole dimension (Bisantz and Vicente, 1994). While it is not explicitly shown in Figure 2, at this level we also note that the two reservoirs are coupled by the redundant FWSs. Either FWS can be configured to feed either reservoir.

The level of *Physical Form* is not shown in Figure 2 because it is similar to the level of Physical Function. The difference is that whereas Physical Function specifies the settings and causal relationships between components, Physical Form specifies the spatial layout and appearance of those very same components.

Figure 3 shows the relationships between levels of the AH representation, revealing three important structural properties. First, the mappings between Functional Purpose and Abstract Function indicate that the temperature goal is connected to both the energy and mass balances. The reason for this is that temperature is defined as energy per unit mass. In contrast, the demand goal is only connected to the mass balance. Thus, making a change in the mass flow to affect volume may have the unintended side effect of changing temperature. A second property of interest results, not from a fundamental consideration, but from the particular physical configuration that was selected for the system. In particular, there is a many-to-one mapping between Abstract Function and Generalized Function. There are several instances where a single subsystem is a means for controlling both the mass and energy balances. For instance, each FWS serves as both a source of mass and a source of energy. Thus, there is a structural coupling between the mass and energy configurations. Note that these couplings would not exist if separate subsystems were implemented to instantiate the mass and energy balances. The third structural property is similar. Figure 3 indicates that there is a many-to-one mapping between Physical Function and Generalized Function. This is because either FWS can be used to supply either reservoir. In other words, there is an indirect coupling between reservoirs via the FWSs. Thus, changing the settings of the valves or pump to control the amount of water flowing into one of the reservoirs could also inadvertently affect the flowrate to the other reservoir.

Fig. 3. AH representation of DURESS II showing the relationships between levels of the hierarchy (adapted from Vicente and Rasmussen 1990)

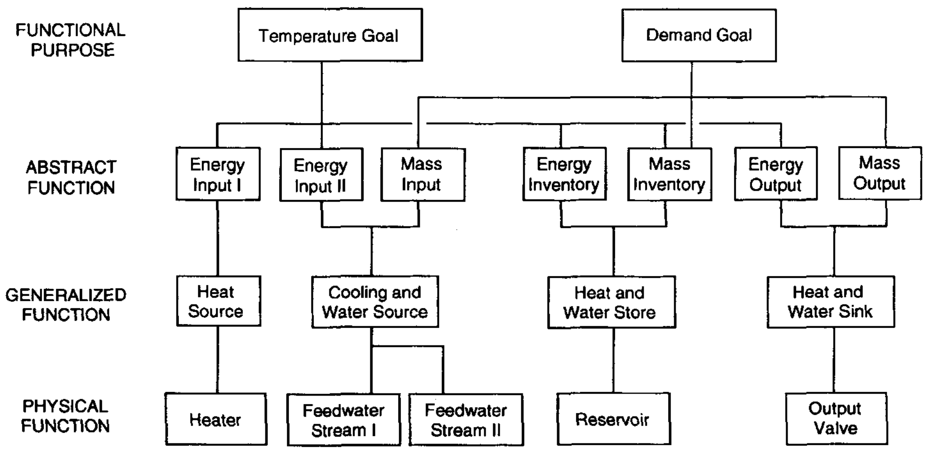

The preceding discussion reveals how the AH can provide some important insights into the structure of DURESS II. Some other properties of the hierarchy itself will be illustrated by making specific reference to Figures 2 and 3. One interesting property is that faults propagate upward in the hierarchy. With reference to Figure 3, if there is a leak in the reservoir, this causes a disturbance in the heat and water store, which in turn is reflected by changes in the mass and energy inventories, ultimately jeopardizing the achievement of the higher-level goals. Reasons for proper system functioning are derived top-down in the hierarchy. To take two examples, the heater element exists to supply heat, and the mass and energy topology has been designed to satisfy the functional purposes of the system. Finally, the relations within levels of abstraction shown in Figure 2 provide information as to what will happen when certain changes are made (e.g., increasing the mass-source flow rate will increase the mass inventory gradient).

To summarize, the AH representation of DURESS II illustrates the various layers of functional structure that are inherent in the domain. Since it describes the way the system actually functions, this representation can be viewed as a normative mental model of the system (i.e., it specifies the way the operator should think of the system during problem solving), thereby providing a basis for interface design.

SUPPORTING RULE-BASED BEHAVIOR. The next step in building an EID interface is to embed the AH structure into the interface. One can think of this phase is making visible the otherwise unobservable goal-relevant properties of the domain. According to the principles of EID, this is accomplished by mapping the relational structures of the domain, represented in this case by the system's dynamic equations, onto the visible geometric properties of the interface objects. The goal is to allow people to deal effectively with many task demands by relying on their powerful pattern recognition capabilities. Figure 4 illustrates how this has been accomplished for DURESS.

In Figure 4a, the state equation describing the mass balance is given. This equation describes a relationship between a set of variables, and thus can be considered a form of domain knowledge. With the goal of making visible the invisible, a geometric figure was developed to represent perceptually the relational structures specified by the state equation. As a result, there is an isomorphic mapping between the geometry of the trapezoid and the relationships between the variables describing the mass balance. Since the general form of the energy-balance equation is very similar to that of the mass-balance equation, a mapping similar to that shown in Figure 4a can also be developed to represent the energy balance.

Figure 5 provides a close-up view of the mass-balance display derived from

Fig. 4. Making visible the invisible: mapping the relational structures of the work domain onto the geometrical properties of the interface (from Vicente and Rasmussen 1990).

A. Mass Balance

1.) State Equation

$$\frac{dV(t)}{dt} = \frac{W_i(t) - W_o(t)}{\rho}$$

2.) Geometry

B. Temperature

1.) Algebraic Equation

$$T_w(t) = \frac{E_{tot}(t)}{V(t)\,\rho\,C_p}$$

2.) Geometry

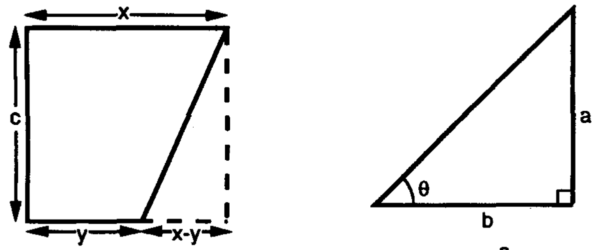

A: $$\text{Slope} = \frac{c}{x-y}$$

B: $$b = \frac{a}{\text{Tan}\theta}$$

A. 3.) Mapping

$$\frac{dV(t)}{dt} = \frac{1}{\text{slope}}$$

$W_i(t) = x$

$W_o(t) = y$

$\rho = c$

B. 3.) Mapping

$E_{tot}(t) = a$

$T_w(t) = b$

$V(t)\,\rho\,c_p = \text{Tan}\theta$

Legend: $V(t)$ = reservoir volume

$W_i(t)$ = input flow rate

$W_o(t)$ = output flow rate

$T_w(t)$ = reservoir temperature

$E_{tot}(t)$ = total reservoir energy

ρ = density
c_p = specific heat capacity
} Constants

the mapping in Figure 4a. The input is shown at the top of the graphics (MI1). The volume level is indicated by the scale on the left side of the graphic (V1) and by the shaded area of the container. The output, MO1, is shown at the bottom. The sloped line connecting the input and output relies on a funnel metaphor. For example, if the top is wider than the bottom (i.e., input > output, as in Figure 5a), then it is easy to visualize the consequence, namely, that volume should be increasing. Conversely, if the bottom is wider than the top (i.e., output > input, as in Figure 5b), then volume should be decreasing. Thus, the slope of the line represents the rate at which the volume should be changing.

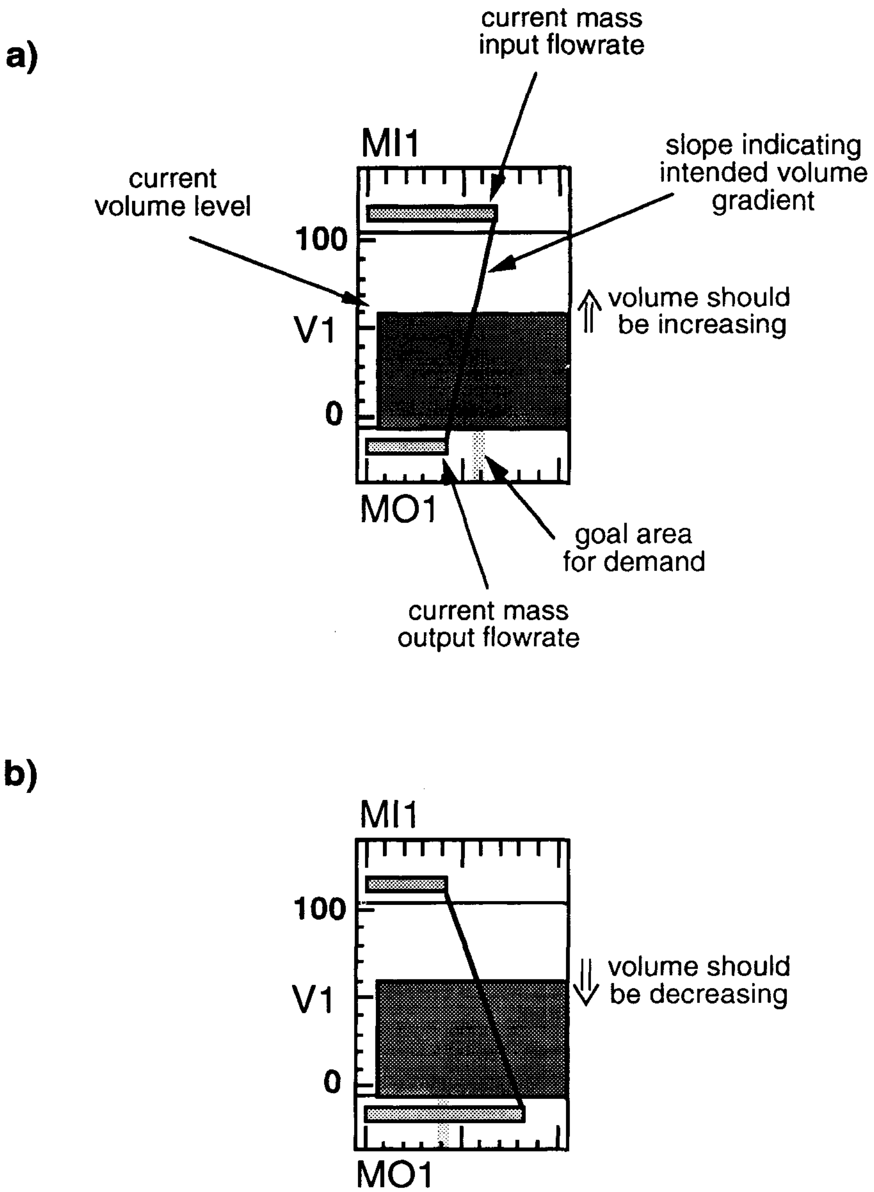

Fig. 5. Close-up view of the mass-balance display, following the mapping in Figure 4a. When input is greater than output (a), the slope is positive indicating that the volume should be increasing. When input is less than output (b), the slope is negative indicating that the volume should be decreasing (from Pawlak and Vicente 1996)

The next display, illustrated in Figure 6, shows the energy balance. Since it is based on essentially the same mapping function, the energy-balance display functions in a very similar manner to the mass-balance display. As shown in Figure 6a, the primary difference is that the energy input to the reservoir (EI1) is partialled out according to its two contributors. The rate at which energy is

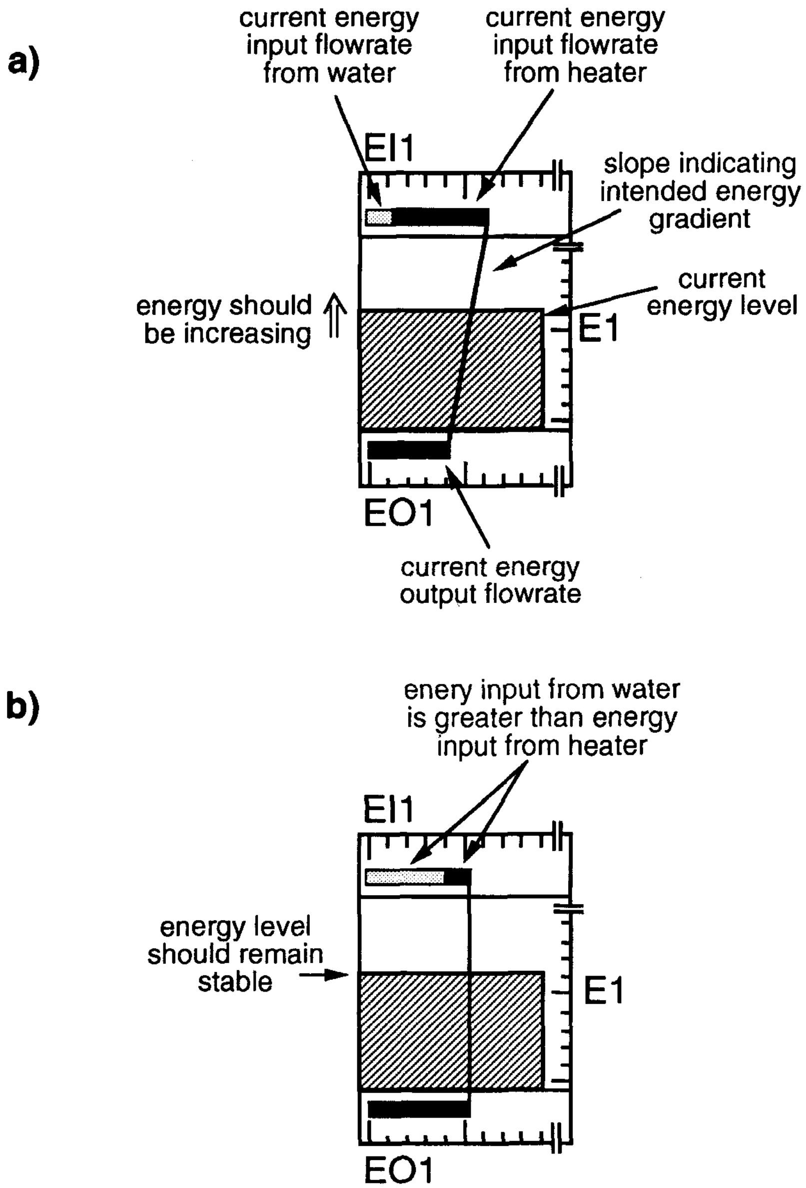

Fig. 6. Close-up view of the energy-balance display, following the mapping in Figure 4a. This display functions in the same manner as the mass-balance display, except that there are two energy input sources, the water coming into the reservoir (lightly shaded bar) and the heater (darkly shaded bar). Figure 6a shows the case where the heater is adding energy at a greater rate than the incoming water. Also, in this case, the energy inventory should be increasing because the energy input is greater than the output, as indicated by the positive slope. Figure 6b shows the case where the incoming water is adding energy at a greater rate than the heater. Also, in this case, the energy inventory should be stable because the energy input equals the energy output, indicated by the perpendicular line (from Pawlak & Vicente 1996)

added by the FWS (i.e., temperature of the water, in °C, multiplied by its flowrate) is shown as a lightly shaded bar, and the rate at which energy is added by the heater is shown as a dark bar. The energy inventory is represented by the area, E1. Again, a funnel metaphor is used, so that, for example, if the total energy input flowrate is greater than the output (as in

Figure 6a), then energy inventory should be increasing. Figure 6b shows the same display in a different state. Now, the incoming water is adding energy at a greater rate than the heater. Another difference is that the total energy input flowrate now equals the energy output flowrate, so the sloped line is perpendicular, indicating that the energy inventory should remain stable.

If we return to Figure 4, the algebraic equation relating the temperature of the water in the reservoir, the energy contained therein, and the reservoir volume is shown in Figure 4b. Again, a geometric figure was developed to illustrate the relationships specified by the equation. The result, in this case, is an isomorphic mapping between the trigonometric relationships describing the geometry of the triangle and the algebraic equation describing the behavior of the system.

A close-up of the display based on this mapping is provided in Figure 7. The horizontal line with a ball on the end emanates from the current volume level (V1). Change in the height of this line always accompanies any change in mass inventory (i.e., the bar will always be at the same height as the water level, V1). The diagonal line in the center display rotates about its leftmost endpoint (connected to the top left of the T1 box) and always passes through the ball on the end of the horizontal line. The slope of the diagonal line represents the function that maps the relationship between volume and energy onto temperature. This mapping is indicated by the line emanating from the current energy inventory level (E1) that comes across and reflects off the diagonal line at a right angle down onto the current value on the temperature

Fig. 7. Close-up view of the display showing the relationship between volume level, energy inventory level, and reservoir temperature, based on the mapping in Figure 4b (from Pawlak & Vicente 1996)

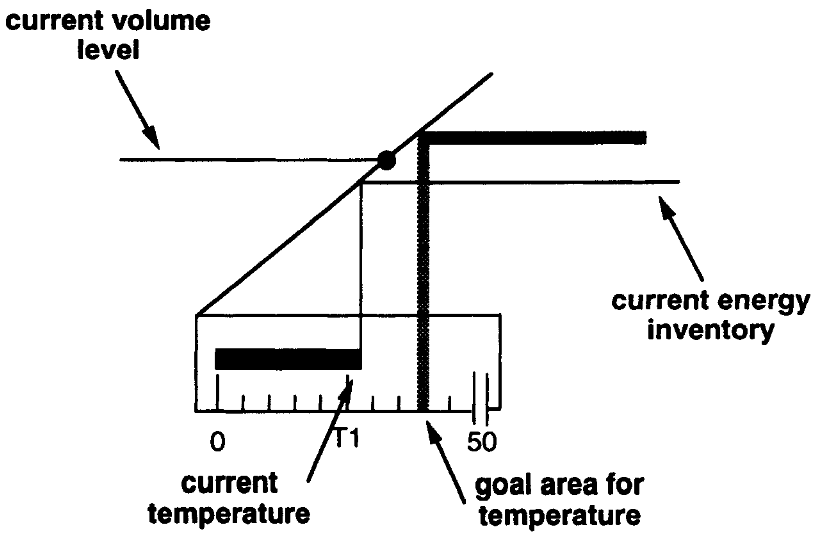

scale (T1). The goal temperature is indicated by the thin shaded area on the temperature scale. This goal area reflects up from the temperature scale, off the diagonal line, and onto the energy inventory scale. As a result, the energy level required to achieve the goal temperature is directly visible.

Figure 8 shows how this display reveals changes in system state. In Figure 8a, the energy inventory is being held constant while the volume is increased slightly. The increase in volume causes the horizontal line emanating from the volume level to go up. This change in vertical position, in turn, causes the diagonal line in the center display to rotate counterclockwise, increasing the slope of the diagonal line. The energy inventory level is at the same height, but it now intersects the diagonal line at a different point. As a result, the new temperature is lower than the old temperature, which makes intuitive sense since the energy per unit volume is lower than before, as a result of the increase in volume. Note also that the energy level required to achieve the same goal temperature is now greater than before because of the increase in volume.

Figure 8b shows how this display changes when volume is held constant and energy inventory level is decreased. In this case, the change is much simpler. Because volume is constant, the height of the horizontal bar is fixed, which in turn fixes the slope of the diagonal line in the center. The only change that occurs when energy inventory is decreased is that the horizontal line emanating from the far right goes down, which means that it intersects the diagonal at a different point, indicated by the dashed line in Figure 8b. Again, the new temperature is lower than the previous one because there is now less energy in the same amount of water.

The key behind the mappings described in Figure 4, and the resulting displays illustrated in Figures 5 to 8, is that they translate the constraints inherent in the domain into constraints on the interface geometry. This is equivalent to embedding knowledge about the domain in the interface. The great advantage of providing an externalized representation of the domain constraints in the interface is that it unburdens the operator's memory. For example, rather than having to remember and retrieve the relationship describing the mass balance from memory, the operator can rely on the interface to provide this information. Making visible the invisible, then, allows operators to control the system by relying more on the powerful capabilities of perception and action than on comparatively more effortful and error-prone inferential processes.

The product of the design process up until this point is shown in Figure 9. Information at every level of the AH is included. Beginning at the top level of Functional Purpose, the demand and temperature setpoints are represented as thin goal areas on the mass output (MO1, MO2) and temperature scales (T1,

Fig. 8. The dynamics of the display in Figure 7. Figure 8a reveals how the display changes when energy inventory is held constant while volume is increased, leading to a decrease in temperature. Figure 8b reveals how the display changes when volume is held constant while energy inventory is decreased, leading to a decrease in temperature (from Pawlak & Vicente 1996)

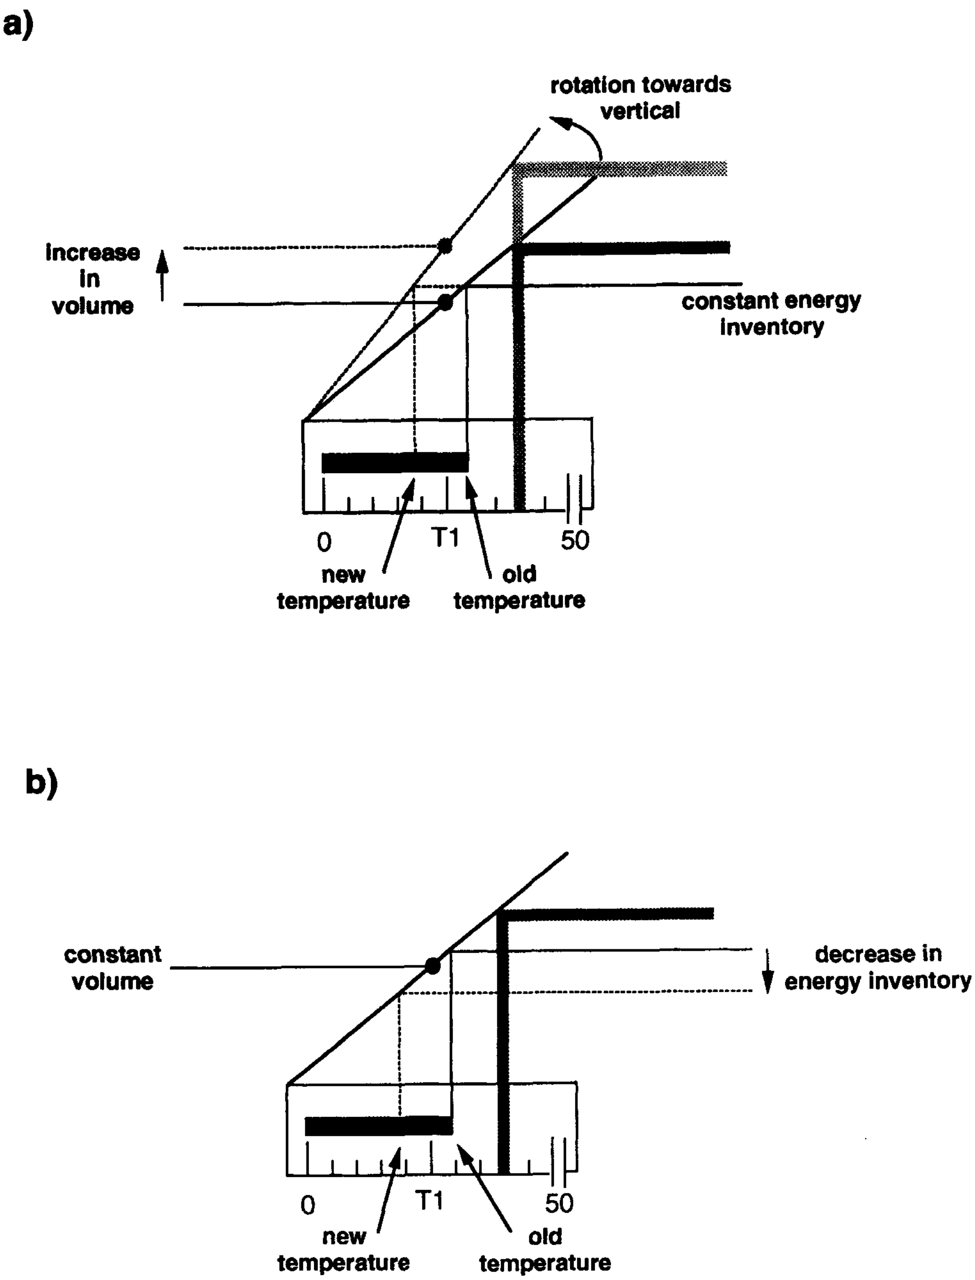

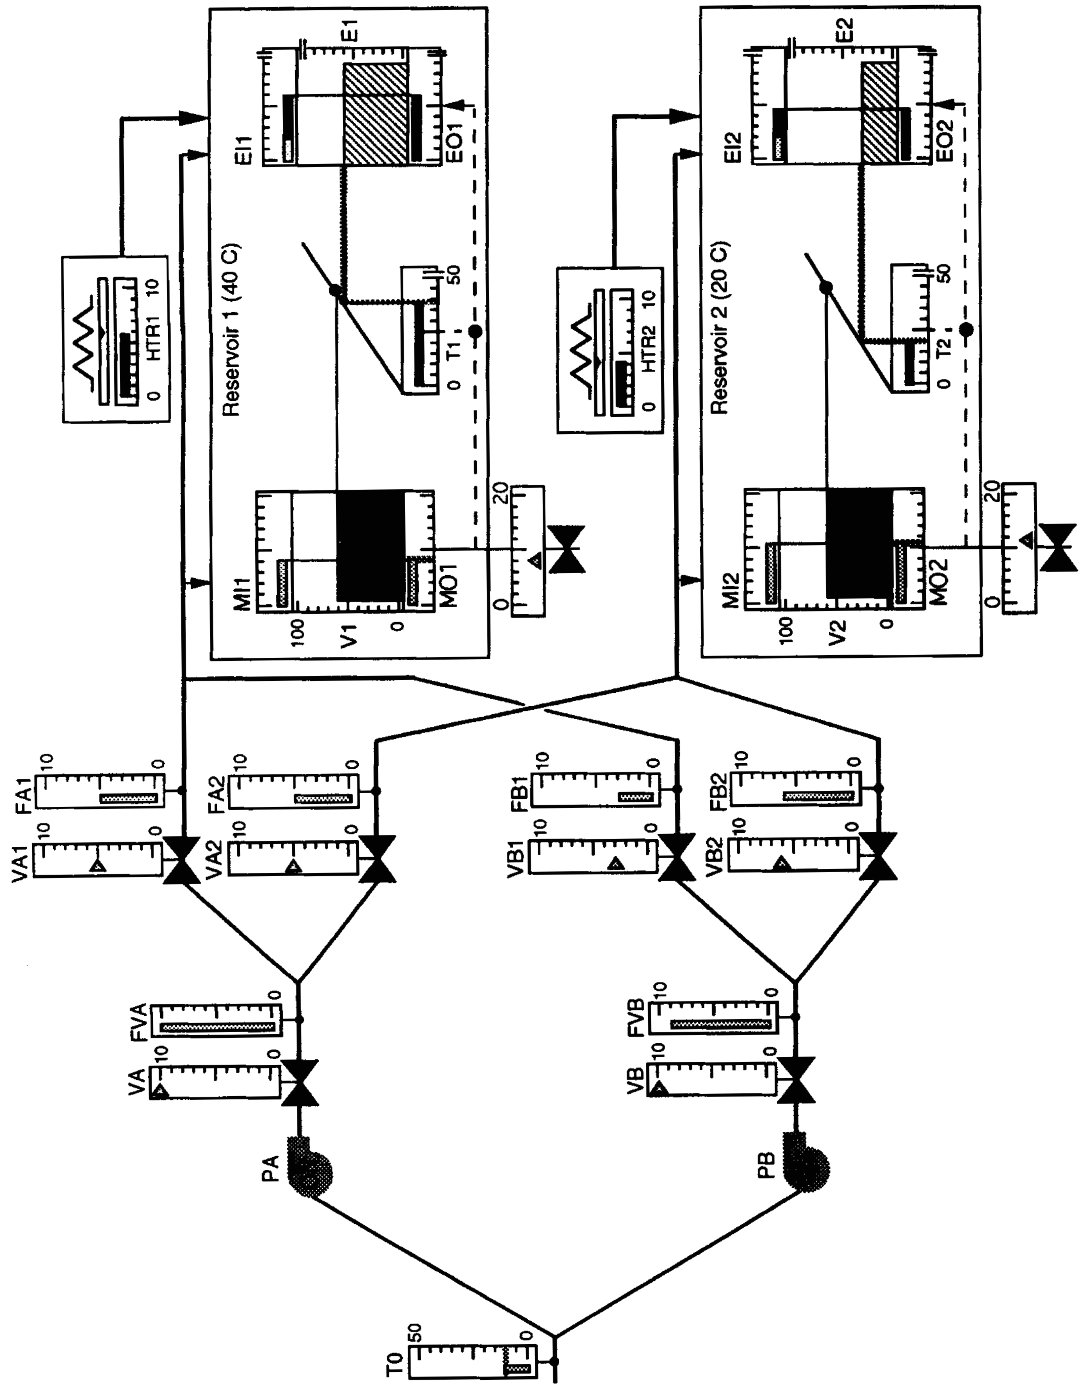

Fig. 9. $P + F$ interface for DURESS II. This figure shows the system in steady state, and therefore provides a baseline against which to compare Figures 10 and 11 (from Pawlak & Vicente 1996)

T2) in the interface. The level of Abstract Function is represented by the boxed groups of graphics on the right. This is the most innovative portion of the $P + F$ interface, providing additional higher-order functional information in the form of first principles (i.e., mass and energy conservation laws). Within each boxed group, the rectangular graphic on the left represents the mass balance for the reservoir (depicted in Figure 5), and the rectangular graphic on the right represents the energy balance (depicted in Figure 6). The display in the center of each group with the diagonal line (depicted in Figures 7 and 8) shows the relationship between volume, energy inventory, and temperature. Note that off-scale markers are represented in the output temperature scales and the energy input, inventory, and output scales as well. These were added to the interface by creating a gap in the scale at the off-scale point, thereby allowing subjects to discriminate the maximum value from off-scale. At the Generalized Function level, the flowrates in each FWS (e.g., FVA, FPA, FA1, FA2) and the heat transfer rates (e.g., HTR1) are displayed as bar scales. At the level of Physical Function, the valve settings (e.g., VB) and heater settings (e.g., HTR2) are indicated by the small triangular pointers on the respective scales. Since the pump settings (e.g., PB) are discrete (either ON or OFF) they are directly labeled on the pumps themselves. At the final level of Physical Form, the relative spatial layout of the components and the connections between them are also represented.

To summarize, the interface in Figure 9 has exploited the mappings described in Figure 4 to make the process dynamics visible to the operator. Whereas the interface in Figure 2 included physical information (i.e., the levels of Functional Purpose, Physical Function, and Physical Form), the $P + F$ interface in Figure 9 also includes functional information as well (i.e., the levels of Generalized Function and Abstract Function), thereby representing the entire AH.

Supporting skill-based behavior. The final principle of EID is concerned with action. EID suggests that, whenever possible, commands should be communicated by directly acting on the display—the familiar idea of direct manipulation. The intent here is to close the perception–action loop, thereby taking advantage of skilled motor control. As an example, to change the setting of the components in the EID interface in Figure 9 (e.g., VA, or HTR1), it is sufficient to "drag" the triangular setpoints to the new value using a mouse. For the pumps, a change in state is achieved by merely clicking on the pump icon itself. This direct-manipulation design is preferable to abstract, and often awkward, command languages, commonly found in some word processors.

AN EXAMPLE. An impression of the dynamics of the $P + F$ interface can be obtained by viewing Figures 9 to 11, which show the progression of events before and after a block in valve VA1. In Figure 9, the system is at steady state. Figure 10 illustrates the system five seconds after the onset of the fault, and shows a number of significant changes. FA1, the flow through the blocked valve, has started to decrease, as has FVA. This reduction in inflow, in turn, causes the mass and energy inputs to Reservoir 1 (MI1 and EI1, respectively) to decrease. As a result, the input flowrates for both mass and energy are now less than their respective output flowrates, thereby causing the two lines connecting the input and output flows to deviate from vertical. Using the funnel metaphor on which these balance displays are based, the bottom is now wider than the top, so the two slanted lines indicate an expected negative gradient, which should cause the levels to decrease. The decrease in FA1 also causes a temporary surge in FA2 which propagates to a slight increase in MI2 and EI2. In this case, the balance lines are slightly tilted in the opposite direction (i.e., inputs slightly greater than outputs), indicating an expected small positive gradient. Figure 11 shows the state of the system twenty seconds after the fault, by which point VA1 has become completely blocked. This fact is clearly shown by the fact that VA1 is open but there is no flow through the valve (i.e., FA1 = 0). As a result, FVA has decreased as well, now supplying only the rate being drawn through VA2. Because FWS A is no longer supplying water to Reservoir 1, MI1 and EI1 have decreased considerably. Examining EI1, the fact that the lightly shaded bar (representing the contribution of the incoming water) is much smaller than before and the dark bar (representing the contribution of the heater) is the same length as before indicates that the reduction in EI1 is caused by a reduction in the flow of water, not a reduction in the transfer of heat. This reduction in inflow causes the sloped lines for the mass and energy balances to be even more negative than in Figure 10. As expected, the volume (V1) and the energy inventory (E1) have decreased as a result. Furthermore, because there is no longer any cold water coming into Reservoir 1 and the volume is decreasing, the temperature (T1) has increased and is now slightly outside the goal region.

In summary, Figures 9 to 11 give an idea of how the $P + F$ interface reacts dynamically to a fault. A number of changes indicating propagation in time and space are illustrated. These changes are intended to aid performance, but it is equally possible that they can confuse people and thereby have no effect, or even a negative effect, on performance. Thus, it is important to compare empirically the P and $P + F$ interfaces to evaluate their relative worth.

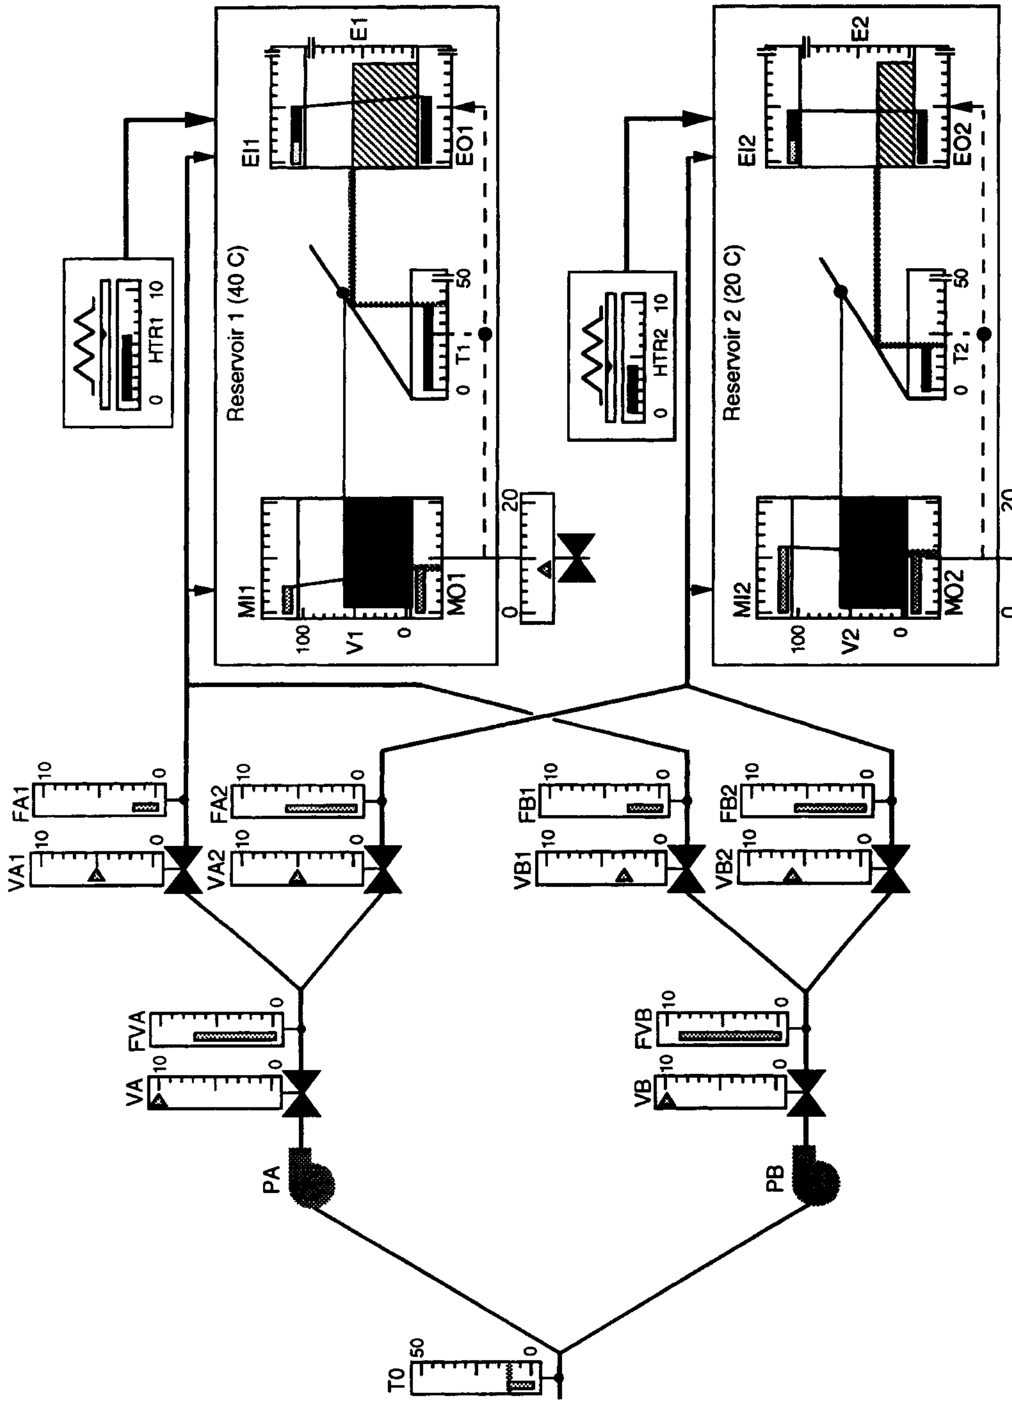

Fig. 10. *P* + *F* interface during a transient showing the dynamic content of the interface. This snapshot was taken five seconds after the beginning of a block of valve VA1 (from Pawlak & Vicente 1996)

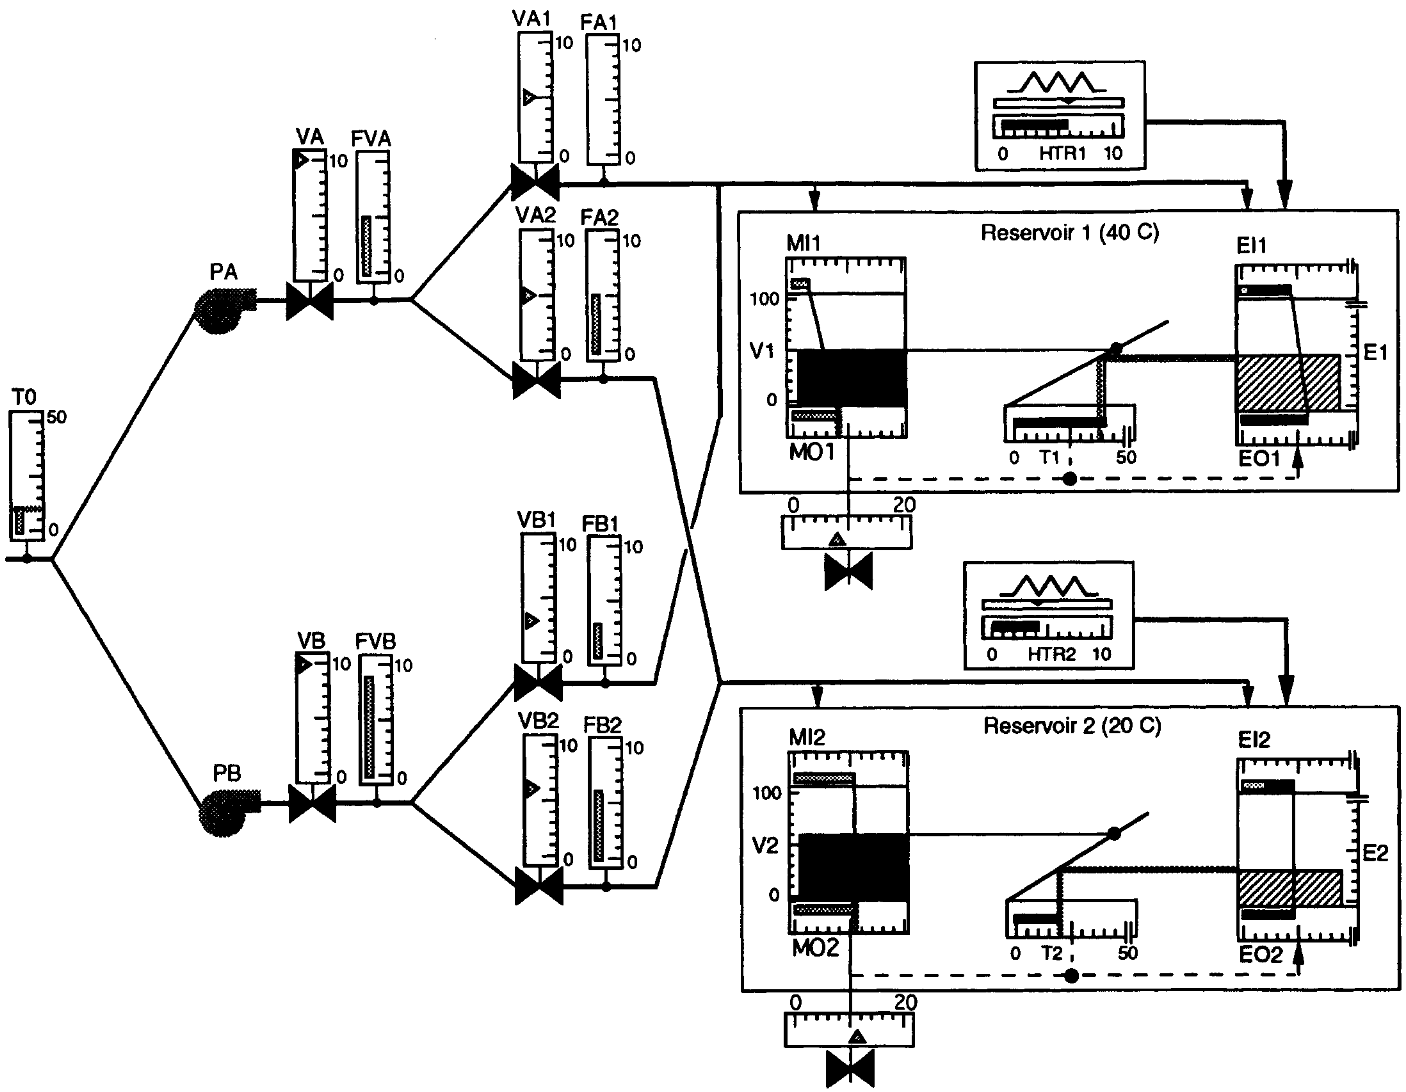

Fig. 11. $P + F$ interface during a transient showing the dynamic content of the interface. This snapshot was taken twenty seconds after the start of the fault, at which point valve VA1 had become completely blocked (from Pawlak & Vicente 1996)

Laboratory findings

The first experimental evaluation of EID was conducted in the context of the non-interactive version of DURESS (Vicente, Christoffersen and Pereklita 1995, Experiment 1). The goal of that study was to compare the earlier versions of the P and $P + F$ interfaces (very similar to those in Figures 1 and 9, respectively) in terms of how well they support problem-solving behavior. Theoretical experts and novices were required to diagnose and remember the values of dynamic, "canned" scenarios using one of the two interfaces. The experimental evidence indicated that the $P + F$ interface provided better support for problem solving than the P interface. The second evaluation of EID, also conducted with DURESS (Vicente et al. 1996, Experiment 2), examined subjects' ability to diagnose real-time canned scenarios, but this time only using the $P + F$ interface. In addition, verbal protocols were collected as the subjects tried to diagnose the nature of the events presented to them. A process-tracing analysis was conducted by mapping subjects' verbalizations onto a two-dimensional problem-space representation of DURESS, defined by an AH and a part–whole hierarchy (Rasmussen 1986). Process measures of performance were correlated with product measures to determine whether there were any statistically significant relationships between subjects' cognitive strategies and their diagnosis accuracy. The results indicate that, the greater the extent to which subjects adopted the top-down "zooming-in" strategy, which the AH is intended to support (Rasmussen 1986), the more accurate their diagnosis performance was. This experiment was the first to demonstrate empirically and reliably the problem-solving advantages associated with reasoning in an AH problem space. Nevertheless, the results are limited by the fact that the subjects did not have extensive practice and did not interactively control the system.

A third study was designed to overcome the limitations of the previous two (Pawlak and Vicente 1996). The experiment was conducted with the interactive DURESS II simulation so that subjects could control the system components. Slightly revised versions of the P and $P + F$ interfaces, illustrated in Figures 1 and 9, were compared. Subjects were given extensive practice at controlling the system (one hour per weekday for four weeks) with one of the two interfaces before their performance on normal events and unfamiliar faults was evaluated. The results revealed that, under normal conditions, there was no performance difference between the $P + F$ and P interfaces. However, dual-task performance results indicate that the P interface relies more on verbal resources, whereas the $P + F$ interface requires more spatial resources. A process-tracing analysis of the fault trials showed that the $P + F$ interface led to faster fault detection and more accurate fault-diagnosis performance. Moreover, the $P + F$ subjects exhibited a more sophisticated and effective set of fault-management strategies, similar to those observed in field studies of experienced operators in complex human–machine systems. In addition, a deficiency of the $P + F$ interface was identified, suggesting a need for integrating trend information with emergent feature displays. The primary contribution of this study, then, was to compare the P and $P + F$ interfaces under a more *representative* (Brunswik 1956) set of experimental conditions.

The most recent experimental evaluation was intended to investigate the long-term influences of EID on operator performance and knowledge (Christoffersen, Hunter and Vicente, 1996a; 1996b; 1996c). More specifically, a longitudinal experiment lasting six months was conducted to compare the P and $P + F$ interfaces for DURESS II under a variety of conditions, including normal trials, routine faults, and non-routine faults. Just as in the previous study, subjects controlled the system every weekday (not including holidays) for about one hour per day. Product measures (e.g., time, actions) and process measures (e.g., verbal protocols) of performance were collected. In addition, several knowledge elicitation measures were occasionally administered to determine how the subjects' knowledge organization evolved over time. At the end of the experiment, subjects switched interfaces and had to control DURESS II under normal and routine fault conditions using the new interface.

The primary findings of the study were threefold. First, on normal trials, there was very little difference in the average performance of the interface groups. The group using the P interface, however, consistently showed more variability in their performance, occasionally taking much longer than usual to complete the required tasks. This effect was found to hold even after five months of practice. Moreover, a transfer manipulation, requiring subjects to control the system with the interface that the other group had been using, showed that the $P + F$ group (now using the P interface) became more variable than the P group (now using the $P + F$ interface). These results clearly show that the enhanced consistency observed over several months of practice was due to the $P + F$ interface, not to the subjects. Second, for both routine and non-routine faults, the $P + F$ interface was found to lead to better fault management performance, especially with respect to diagnosis accuracy. These effects seemed to stem from strategy differences between interface groups, which were in turn a result of the interaction between the subjects' knowledge and the information provided in the interfaces. Third, subjects using the $P + F$ interface, who actively explored the system and reflected on the feedback provided, achieved levels of adaptation and performance not observed with subjects using the P interface. However, one $P + F$ subject, who adopted a surface approach to learning, exhibited a very shallow knowledge base and poor performance (although no worse than that attained by subjects using the P interface with a comparable level of motivation). Thus, it appears that there are certain preconditions that have to be satisfied if the benefits of an EID interface are to be fully enjoyed.

Industrial prototype

Although the results obtained from the laboratory studies with DURESS and DURESS II have been very encouraging, it is important to determine if the EID framework scales up to industrial systems. Accordingly, a project was undertaken to design a prototype EID interface for the feedwater subsystem of an ABB conventional power plant (Dinadis and Vicente 1996). The study showed that the EID framework can be meaningfully applied to industrial systems that are much larger in scale than the DURESS system. Another significant outcome of the feasibility study was a proof of principle that EID can be integrated with other interface design concepts, such as visual momentum (Woods 1984) to address navigation issues and perceptual organization principles (Mumaw, Woods, and Eastman 1992) to address visual-form issues.

Technology transfer to industry

Research on EID has already led to varying degrees of technology transfer to the nuclear and process-control industries. A very modest form of transfer has occurred with AECL Research and Honeywell. Both of these companies have adapted the $P + F$ interface for DURESS into prototypes that are intended to illustrate state-of-the-art interface design concepts, which may eventually find themselves in advanced control rooms of the future. More substantial technology transfer has occurred with Toshiba in Japan (Monta et al. 1991), who have adopted EID as the basis for designing their advanced control room for a next-generation boiling-water reactor plant and have incorporated and adapted specific features of the $P + F$ interface for DURESS II (e.g., the mass balance graphics in Figure 5, and an adapted version of the display in Figure 7) into some of their displays. This application is notable since it has been conducted at the scale of a full-scope nuclear power plant simulator. Finally, Mitsubishi Heavy Industries in Japan has demonstrated a very strong interest in EID, contracting Battelle to initiate a five year research program, solely on EID (Lee and Sanquist 1995).

In summary, although many issues remain to be investigated before one can confidently recommend that industrial control rooms be based on the EID framework, the research conducted so far has generated some very promising findings leading to technology transfer to industry. In particular, this research program has shown that dynamic decision-making performance in complex systems can be significantly influenced by the interface that users have available to perform the task. This fact has important implications for system dynamics researchers investigating dynamic decision making in management flight simulators.

Implications for system dynamics research

In contrast to the research on EID, system dynamics research has frequently led to the conclusion that people are severely impaired in their ability to cope effectively with complex systems (e.g., Sterman 1989). This conclusion has also been reached by researchers outside of the system dynamics community (e.g., Dörner 1987; Brehmer, 1992). One potential explanation for this conflicting set of findings is that the subjects did not have enough practice with the simulation to adapt to its dynamics. The results of Paich and Sterman (1993) rule out this interpretation, at least for one microworld. Nevertheless, an interaction between interface design and practice cannot be ruled out at this point. A second possibility is that the dynamics of the systems investigated by these researchers are more complex than those of the systems investigated under the rubric of EID. There is plenty of evidence indicating that more complex system dynamics negatively impact performance (e.g., Diehl and Sterman 1995), but no controlled comparison seems to have been made between the types of systems studied by system dynamics researchers and a process control system, such as DURESS II. Thus, it is not known to what extent this set of factors accounts for the conflicting set of results described above, although it is very likely that it accounts for at least some of the observed differences in performance.

A third, as yet unexplored, explanation is that the findings paint a very unflattering picture of people's capabilities to engage in dynamic decision making in complex systems which may, in part, be due to the impoverished interfaces used in those experiments. In fact, some authors explicitly refer to opaqueness (Brehmer 1992) or lack of transparency (Dörner 1987) as characteristics of dynamic decision making problems. However, the research reviewed above shows that opaqueness is a property of an interface, *not* an inherent property of complex systems.

There are reasons to suspect that the interfaces used by system dynamics researchers are opaque, and therefore more similar to the *P* than to the *P + F* interface for DURESS II. One of the important differences between the *P* and *P + F* interfaces is that the former only presents data, whereas the latter tries to show relationships between data as well. To take an arbitrary example, the interface developed by Sterman (1987) for the Strategem-2 microworld also focuses primarily on raw data. This can be seen if one compares Figures 3 and 1 in Sterman (1987), which show, respectively, the interface and the relationships that actually govern the simulation. The key system relationships (e.g., the lags and the positive feedback loop) are not perceptually specified in the interface in a salient manner. This may explain why subjects treated the system as if there was no feedback loop, and exhibited poor performance (Sterman, 1989). This possibility has been recognized recently in the system dynamics community (Diehl and Sterman 1995, p. 213) but it has not been pursued. We are conducting research in our laboratory to test this hypothesis.

In the meantime, there are several differences between process control systems and management-decision domains that must be considered in evaluating the generalizability of the research on EID to the problems of greatest concern to the system dynamics community. Some of these issues must await experimental evaluation before they can be confidently resolved, but it may be helpful to discuss them here. First, in MFSs the simulation is controlled by the learner rather than proceeding in real time. This does not pose any problem for design, since the $P + F$ interface for DURESS II could just as well be used for a discrete, batch simulation. Whether the performance advantages of EID described above will also be observed under these conditions is an open issue, however. Second, managers make decisions at a much slower pace than process-control operators, which suggests that skilled perceptual-motor control may not be achievable. This does not seem to pose an insurmountable problem, since, even in process plants, operator control actions are discrete and sporadic, rather than continuous, because of the slow dynamics of the plant. Moreover, there does not seem to be any reason why information (e.g., monthly sales reports) cannot be presented so as to exploit effectively the power of perception, regardless of the dynamics of the environment. Third, the problems with which managers tend to be concerned are nonlinear and of higher dimension than DURESS II. There does not seem to be an insurmountable problem in generalization here either, because displays that have some of the features of the $P + F$ interface have been empirically demonstrated to lead to enhanced performance for systems with nonlinear dynamics (Vicente et al. 1996) and of high dimensionality (Beuthel et al. 1995). It is important to note, however, that the impact of EID on human performance in systems with positive feedback loops has yet to be evaluated. Fourth, it is possible that interfaces based on EID cannot be developed for managers outside the laboratory because the available models of actual organizations and their environments are not as reliable as those that are available for process-control plants. This is certainly a potential limitation, but its investigation can be deferred until the value of principles such as EID has been evaluated in MFSs, where system models are known because they form the basis for the simulation.

In closing, it is important to address head-on a counter argument which may cause some readers to undervalue the relevance of the work presented in this paper to the system dynamics community. More specifically, it could be argued that the interfaces that are typically used in system dynamics research on MFSs are as good as, or better than, existing interfaces used by managers. This might lead one to believe that such research is actually more representative than using an interface based on the principles of EID. This is a perfectly legitimate approach if one's goal is to *understand* the current decision-making competency of managers. However, if one's goal is instead to *improve* the decision-making competency of managers, then this approach is not nearly as productive. By accepting existing interfaces as an unalterable given, one might be led to the unqualified conclusion that people are poor decision makers in complex systems. As mentioned above, this seems to be the conclusion that has been voiced in the system dynamics community. But the research reviewed above suggests that those disappointing claims may be partly a function of the interfaces that have been used to date in that line of research.

An exaggerated analogy would be to assess people's abilities to drive a car by conducting an experiment at night, in poor weather with bad visibility, with a very dirty windshield, and with no wiper fluid. People may not be very proficient under such situations, but that does not necessarily mean that people are bad drivers. All one can conclude is that performance is not ideal under conditions with limited visibility. Moreover, if one has control over the factors that affect poor driving performance, then it may be possible to create the conditions for better performance. An analogous interpretation should be offered for existing research on MFSs; people do not seem to do very well with the interface at hand. However, the interface made available to managers is a potential point of design leverage. Only future research can determine whether people are inherently limited in their ability to control systems with lags, nonlinearities, and positive feedback loops, or whether interface design principles such as EID can be leveraged to improve dynamic decision-making performance in complex systems. As with all other false dichotomies in the behavioral sciences, the answer to this question will probably wind up being "a little bit of both".

References

Beuthel, C., B. Boussoffara, P. Elzer, K. Zinser, and A. Tißen. 1995. Advantages of Mass-data-displays in Process s&c. *Proceedings of 6th IFAC/IFIP/IFORS/IEA Symposium on Analysis, Design, and Evaluation of Man-Machine Systems*, 439–444. Cambridge, MA: MIT.

Bisantz, A. M., and K. J. Vicente. 1994. Making the Abstraction Hierarchy Concrete. *International Journal of Human–Computer Studies* 40: 83–117.

Brehmer, B. 1992. Dynamic Decision Making: Human Control of Complex Systems. *Acta Psychologica* 81: 206–223.

Burns, C. M., and K. J. Vicente. 1995. A Framework for Understanding Interdisciplinary Interactions in Design. *Proceedings of DIS '95: Symposium on designing interactive systems*, 97–103. New York: ACM.

Brunswik, E. 1956. *Perception and the Representative Design of Experiments.* Berkeley, CA: University of California Press.

Christoffersen, K., C. N. Hunter, and K. J. Vicente. 1996a. A Longitudinal Study of the Effects of Ecological Interface Design on Skill Acquisition. *Human Factors.* 38: 523–541.

———. 1996b. A Longitudinal Study of the Effects of Ecological Interface Design on Fault Management Performance. *International Journal of Cognitive Ergonomics.* In press.

———. 1996c. A Longitudinal Study of the Effects of Ecological Interface Design on Deep Knowledge. Manuscript available from author.

Diehl, E., and J. D. Sterman. 1995. Effects of Feedback Complexity on Dynamic Decision Making. *Organizational Behavior and Human Decision Processes* 62: 198–215.

Dinadis, N., and K. J. Vicente. 1996. Ecological Interface Design for a Power Plant Feedwater Subsystem. *IEEE Transactions on Nuclear Science* 43: 266–277.

Dörner, D. 1987. On the Difficulties People Have in Dealing with Complexity. In *New Technology and Human Error* ed. J. Rasmussen, K. Duncan, and J. Leplat 97–109. Chichester: John Wiley & Sons.

Hogarth, R. M. 1981. Beyond Discrete Biases: Functional and Dysfunctional Aspects of Judgemental Heuristics. *Psychological Bulletin* 90: 197–217.

Kahneman, D., P. Slovic, and A. Tversky. 1982. *Judgement under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.

Lee, J. D., and T. F. Sanquist. 1995. *Review of Ecological Interface Design Research: Applications of the Design Philosophy and Results of Empirical Evaluations.* Seattle, WA: Battelle.

Monta, K., Y. Takizawa, Y. Hattori, T. Hayashi, N. Sato, J. Itoh, A. Sakuma, and E. Yoshikawa. 1991. An Intelligent Man–Machine System for BWR Nuclear Power Plants. *Proceedings of AI 91: Frontiers in Innovative Computing for the Nuclear Industry*, 383–392. Jackson, WY: ANS.

Mumaw, R. J., D. D. Woods, and M. C. Eastman. 1992. *Interim Report on Techniques and Principles for Computer-based Display of Data* (STC Report 92-1SJ4-CHICR-R1). Pittsburgh, PA: Westinghouse Science and Technology Center.

Paich, M., and J. D. Sterman. 1993. Boom, Bust, and Failures to Learn in Experimental Markets. *Management Science*, 39: 1439–1458.

Pawlak, W. S., and K. J. Vicente. 1996. Inducing Effective Control Through Ecological Interface Design. *International Journal of Human-Computer Studies.* 44: 653–688.

Rasmussen, J. 1986. *Information Processing and Human–Machine Interaction: An Approach to Cognitive Engineering.* New York: North Holland.

Sterman, J. D. 1987. Testing Behavioral Simulation Models by Direct Experiment. *Management Science* 33: 1572–1592.

Sterman, J. D. 1989. Misperceptions of Feedback in Dynamic Decision Making. *Organizational Behavior and Human Decision Processes* 43: 301–335.

———. 1994. Learning in and about Complex Systems. *System Dynamics Review* 10: 291–330.

Vicente, K. J. 1992a. Multilevel Interfaces for Power Plant Control Rooms I: An Integrative Review. *Nuclear Safety,* 33: 381–397.

———. 1992b. Multilevel Interfaces for Power Plant Control Rooms II: A Preliminary Design Space. *Nuclear Safety* 33: 543–548.

Vicente, K. J., K. Christoffersen, and A. Pereklita. 1995. Supporting Operator Problem Solving through Ecological Interface Design. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-25: 529–545.

Vicente, K. J., N. Moray, J. D. Lee, J. Rasmussen, B. G. Jones, R. Brock, and R. Djemil. 1996. Evaluation of a Rankine Cycle Display for Nuclear Power Plant Monitoring and Diagnosis. *Human Factors.* 38: 506–521.

Vicente, K. J., and J. Rasmussen. 1990. The Ecology of Human–Machine Systems II: Mediating 'Direct Perception' in Complex Work Domains. *Ecological Psychology* 2: 207–250.

———. 1992. Ecological Interface Design: Theoretical Foundations. *IEEE Transactions on Systems, Man, and Cybernetics* SMC-22: 589–606.

Woods, D. D. 1984. Visual Momentum: A Concept to Improve the Cognitive Coupling of Person and Computer. *International Journal of Man–Machine Studies* 21: 229–244.